Making It Explicit

Reasoning, Representing, and Discursive Commitment

Robert B. Brandom

To Wilfrid Sellars and Richard Rorty

without whom most of it would not even be implicit

... both a new world,
and the old made explicit ...

We shall not cease from exploration
And the end of all our exploring
Will be to arrive where we started
And know the place for the first time.

<div align="right">T. S. ELIOT, "Four Quartets"</div>

Contents

PART TWO

5

The Expressive Role of
Traditional Semantic Vocabulary:
'True' and 'Refers'

It was when I said,
"There is no such thing as the truth,"
That the grapes seemed fatter.
The fox ran out of his hole.

You . . . You said,
"There are many truths,
But they are not part of a truth."
Then the tree, at night, began to change.

WALLACE STEVENS, "On the Road Home"

I. FROM INFERENCE TO TRUTH, REFERENCE, AND REPRESENTATION

1. Discursiveness, Sapience, and Sentience

We are the creatures who say 'we'—who can explicitly take or treat someone as one of us. Adopting this practical attitude is adopting a discursive normative stance: attributing propositionally contentful commitments and entitlements, and attributing to those to whom we attribute them in turn the practical recognition of such deontic statuses by their corresponding acknowledgment and attribution of them. Sapience of the sort distinctive of us is a status achieved within a structure of mutual recognition: of holding and being held responsible, of acknowledging and exercising authority. The specifically *discursive* character of that normative social structure—what makes it appropriate to interpret the statuses we institute by our deontic scorekeeping activities as having their significance determined by *propositional* contents—consists in the *inferential* articulation of those recognitive practices. We are the ones who give and ask for *reasons* for what we say and do.

Propositionally contentful commitments are picked out in the first instance as those that can both serve as and stand in need of reasons; it is playing an appropriate role in the practices of giving and asking for reasons that confers propositional contents on commitments and the performances

that express them. Offering a reason is making a claim—performing a speech act with the force or significance of an assertion, the undertaking (by overt, explicit acknowledgment) of a doxastic commitment. Practices in which performances are accorded assertional significance deserve for that reason to be called *linguistic* practices. That what one interprets a community as producing and consuming are *reasons* is not something that can be achieved simply by a stipulation to that effect on the part of the interpreter. Rather, the interpretation must attribute to the community practices that incorporate a suitable structure of *inheritance* of *entitlement* to the commitments the practitioners are understood as undertaking and attributing. Such a structure arises out of the interaction of two different dimensions: the intercontent, intrapersonal *consequential* inheritance of entitlement, and the intracontent, interpersonal *testimonial* inheritance of entitlement. Each deontic attitude on the part of a scorekeeper is the attribution to (or acknowledgment by) an interlocutor of a deontic status having a certain content. What it means for commitments with *different* contents to be undertaken by or attributed to the *same* interlocutor, and for commitments undertaken by or attributed to *different* interlocutors to have the *same* content, is to be understood in terms of this broadly inferential structure of justification and communication.

We are sentient creatures as well as sapient ones, but our sentience is different from that of those who cannot give and ask for reasons. Described in the language of physiology, our sensing may be virtually indistinguishable from that of nondiscursive creatures. But we not only sense, we also perceive. That is, our differential response to sensory stimulation includes noninferential acknowledgment of propositionally contentful doxastic commitments. Through perception, when properly trained and situated, we find ourselves passively occupying particular positions in the space of reasons.

We are practical creatures, as well as linguistic ones, but our purposive activity is different from that of those who cannot give and ask for reasons. Described in the language of physiology, our motor activity may be virtually indistinguishable from that of nondiscursive creatures. But we not only produce performances, we perform actions. The performances we produce include noninferential responses to acknowledgments of propositionally contentful[1] practical commitments. Through action, when properly trained and situated, we can respond to the particular positions we occupy in the space of reasons by actively altering the nondiscursive environment.

Our mammalian cousins, primate ancestors, and neonatal offspring—who are sentient and purposive but not discursive creatures—are interpretable as perceiving and acting only in a derivative sense. An interpreter can make sense of what they do by attributing propositionally contentful intentional states to them, but the interpreter's grasp of those contents and of the significance of those states derives from mastery of the richer practices of giving and asking for reasons for the sort of doxastic and practical discursive com-

mitments that are not attributed to these simpler folk. The activities they are interpreted as engaging in do not suffice to confer anything recognizable as propositional contents on their states, attitudes, and performances. Our discursive practices make us semantically autonomous in a sense in which their nondiscursive practices do not.

2. Making and Taking Reasons, Seeking and Speaking Truths

An account has been offered of practices that are, it is claimed, sufficient to confer genuinely propositional contents on our states, attitudes, and performances. For a sentence to express such a content is for its proper use to be determined by the norms implicit in assertional practices such as those described in Chapters 3 and 4. Even if that account is accepted, however, it is not obvious from what has been said so far why various features of the discursive scorekeeping model ought to be treated as *necessary* for propositional contentfulness. It might be granted that treating a community as engaging in the sort of linguistic social practices described in Chapters 3 and 4 ought to count as interpreting them as instituting doxastic and practical deontic statuses with assertible and therefore propositional contents—as in this sense doing what we discursive interpreters do—while insisting that the linguistic and social dimensions of these attributed practices are not needed for the states creatures are taken to have to deserve to count as propositionally contentful. Responding to this sort of objection requires looking more closely at the notion of propositional content.

The account elaborated in Part 1 of this work focuses exclusively on one way of understanding what it is to exhibit such contents: in terms of inferential articulation, above all playing a premissory role in reasoning. As was acknowledged at the outset of this investigation, however, there is another, perhaps more familiar route to understanding the sort of sapience being called on to do demarcational duty. It begins with the concept of *truth*, rather than that of *inference*. We are believers, and believing is taking-true. We are agents, and acting is making-true. To be sapient is to have states such as belief, desire, and intention, which are propositionally contentful in the sense that the question can appropriately be raised under what circumstances what is believed, desired, or intended would be *true*. Understanding such a content is grasping the conditions that are necessary and sufficient for its truth.

Something has been said about why understanding ourselves as makers and takers of reasons entails understanding ourselves as seekers and speakers of truths; acknowledging a doxastic commitment is taking-true, and producing an assertional performance is putting a claim forward as true. This discussion at best offers only a starting point, however, from which an account of the relation between the theoretical concepts of inference and truth might depart. It needs to be supplemented by an account of the sort of socially articulated inferential role something must play in order to have a

propositional content in the sense of having truth conditions. Theorists who begin by associating sentences or states with truth conditions—whether by bare stipulation or, more ambitiously, by exhibiting practical proprieties of use or function sufficient to confer such contents (so the association is established by the creatures being interpreted rather than the interpreter)—can then derive from that association various norms governing inference. For the inference from p to q can be understood as correct (in the sense of commitment-preserving) just in case the truth conditions of p are a subset of the truth conditions of q.[2] An account such as the present one, which pursues the converse direction of explanation by understanding propositional contents in the first instance in terms of proprieties of inference, accordingly owes a corresponding story about assessments of truth and the association of truth conditions with sentences and the beliefs they express.

In order to turn a construal of the practical attitude of taking-true (as undertaking an inferentially articulated doxastic commitment) into an account of the sort of grasp of the propositional content of a state or utterance that consists in associating truth conditions with it, it is necessary to understand the expressive role played by 'true' in its various uses. What must an interpreter be doing for it to be correct to take what is being associated with the utterances and intentional states of those interpreted to be *truth* conditions? An answer to this question has to explain how an expression (possibly in an alien language) needs to be used—what practical proprieties its employment must be subject to—in order for it properly to be understood as meaning 'true'. In the context of the present project, such an account is subject to two further criteria of adequacy.

First, the vocabulary used to specify the expressive role played by 'true' and its cognates must be that provided by the inferentially and socially articulated deontic scorekeeping model of discursive practice. That is, it is necessary to show how that model can be extended to include truth talk. For the only locutions officially acknowledged to be intelligible within the confines of this work are those that can be understood as having just the contents that are conferred on them by the role they play in discursive practice. So the issue is what one would need to do to add the expressive power provided by 'true' to the game of giving and asking for reasons as so far specified.

Second, the inferentialist order of explanation is committed to showing why expressions, performances, and intentional states having semantic contents that are propositional in the sense of being assertible—here construed roughly as both being capable of serving as a reason and potentially standing in need of reasons—therefore are also appropriately talked about as having contents that are propositional in the sense of having truth conditions. What is the connection between these two ways of picking out the propositional? Both of these explanatory demands can be addressed by showing what it is about the inferentially and socially articulated deontic scorekeeping prac-

tices (which according to the story being told here confer propositional contents) that 'true' has the job of expressing explicitly (that is, in assertible, so propositional, form).

3. The Representational Dimension of Discursive Practice

The need for an account of truth talk in broadly inferential terms is only part of a larger methodological requirement: that the *representational* dimension of intentionality and discourse be addressed. The concept of representation lies at the center of the theoretical idioms typically used today to discuss intentional and semantic contentfulness. It has been at least since Descartes forged his worldview in the bifurcated mold, one side of which consists of representings and the other of representeds—the former conceived as discursive, on the model of algebraic equations, the latter as geometric, on the model of the figures those equations determine. Four aspects of that representational idiom deserve special attention.

First, in addition to the use of 'true' to formulate a semantic approach to sentences in terms of truth values and truth conditions, the tradition to which the later Frege gave birth uses 'refers' or 'denotes' to express the representational relations terms and predicates stand in. Tarski's discussion of truth definitions for formalized first-order languages[3] provides a paradigm of how the use of 'true' and of 'refers' or 'denotes' ought to be understood to be related to each other.

Second, in tandem with this semantic focus on truth and reference is a distinction between extensional and intensional contexts. Central among the latter (which are the source of special explanatory difficulties in the representationalist context) are propositional-attitude-ascribing locutions, paradigmatically ". . . believes that . . ." An account must be offered of the different semantic behavior of these two classes of locutions—a difference that can be roughly rendered in the traditional terminology as the distinction between contexts for which only what is represent*ed* matters for the inferential significance of sentences exhibiting those contexts, on the one hand, and contexts in which represent*ings* themselves are somehow involved in determining such inferential significances, on the other.

Besides these technical semantic notions crafted by self-conscious theorists in the representationalist tradition, there are locutions in ordinary language that are used to express claims concerning what people are talking or thinking *about*. Prime among these ordinary representational locutions is the use of 'of' in *de re* ascriptions of propositional attitudes, as in "Royce believed *of* the author of *Dreams of a Spirit-Seer* that he was a great philosopher" or "What James said about tough- and tender-minded thinkers is true *of* Kant, but not true *of* Hegel." The nontechnical notion of intentional *aboutness* that figures in the fundamental pretheoretical distinction between what people are talking or thinking about and what they say or think about it is to be

understood in terms of what is expressed by the use of these *de re* specifications of the contents of intentional states and the utterances that express them. It is from these roots in ordinary conversational practice that more technical representational concepts are elaborated.

Finally there is the crucial idea of *objective* representational proprieties of judging and inferring. It is a feature of our assessments of the application of at least some concepts that we take it that besides correctness of discursive attitude in the sense of entitlement to the commitment undertaken—the concept-user being *justified* in the application, either noninferentially by observable circumstance, inferentially by concomitant commitments and entitlements, or deferentially by available testimony—there is also a kind of correctness that is determined by how things are with what is represented. Whether the application of a concept is correct or incorrect in this sense is independent of the attitudes of the one applying the concept. ("Full many a gem of purest ray serene / The dark unfathom'd caves of ocean bear: / Full many a flower is born to bloom unseen, / And waste its sweetness in the desert air.")[4] Making sense of this fundamental characteristic of our linguistic practice means funding a notion of discursive *success* (both doxastic and practical) that transcends our attitudes (our takings and tryings) and answers instead to the *objects* those attitudes address. Such success is determined by the properties of those objects and by the relations they stand in. A certain sort of social and inferential articulation of attitudes must be shown to institute proprieties and confer contents such that what it is correct to conclude or to claim and what one has actually done depends on how the objects referred to, talked about, or represented in one's discursive attitudes actually are. Talk of objective representational conceptual contents involves both representation of objects and objective proprieties of representation. It remains to be seen what these are and how they are related.

On the face of it, acknowledgment and analysis of conceptual proprieties that are objective in this sense would seem simply to be incompatible with a social-phenomenalist approach to norms. For what is distinctive of such an approach is precisely its treating the norms that govern our discursive conduct as instituted ultimately by our *attitudes*. How, then, can such an approach hope to make intelligible a notion of attitude-*transcendent* discursive norms? The answer takes the form of specifying a certain sort of *stance*—a constellation of attitudes that can be understood as taking or treating certain sorts of assessments *as* answering to facts concerning the objects represented rather than to anyone's attitudes toward or claims about them. Understanding what it is to adopt such a representational stance toward (or interpretation of) what we do is understanding the *implicit* practical attitudes that are expressed *explicitly* by our use of representational locutions—both the technical semantic tropes involving 'true' and 'refers' and the *de re* ascriptions of propositional attitude used in ordinary language to make clear what has been achieved in the way of communication or action. Indeed, it is just

the aspects of conceptual content made explicit by these representational locutions that turn out to require that the practices conferring those contents be specifically *linguistic social* practices of the sort outlined in the previous chapter.

4. Inference, Substitution, and Anaphora

The various representational tropes just considered all involve associating semantic contents with *sub*sentential expressions—above all, singular terms and predicates. The understanding of reference to objects, of the distinction between extensional and intensional occurrences of expressions, of *de re* locutions, and of the notion of the truth of at least some claims as a normative status or propriety that answers to the properties objects have (and the relations they stand in) rather than those anyone *takes* them to have (or stand in) all involve grasping—either implicitly as a practitioner or explicitly as a theorist—conceptual contents that are not propositional. No account of this sort of conceptual content has yet been offered here. The semantic contentfulness of sentential expressions has been derived from the role they play in inferences, as premises and conclusions; the semantic contentfulness of expressions in the grammatical categories of singular terms and predicates cannot consist in their playing inferential roles in this sense. The notion of conceptual content must accordingly be broadened beyond the propositional in order to give it application beyond the category of sentences.

The inferentialist order of semantic explanation was given its modern form by the young Frege, whose definition of *begriffliche Inhalt* motivates one strand of the model developed here. (The social, practical, and specifically linguistic dimensions of that model—and to some extent the understanding of the normative one as well—are rooted elsewhere.) One of Frege's most important ideas is his strategy for extending broadly inferential notions of content to subsentential expressions, which cannot serve as premises and conclusions of inferences. The key theoretical concept he introduces for that purpose is *substitution*. His idea is that the way in which sentences are related to one another when one results from the other by substituting one subsentential expression for another confers an *indirectly* inferential role on the occurrence of subsentential expressions. Roughly, subsentential expressions can be sorted into equivalence classes that can be thought of as having the same conceptual content in an extended sense. For they can be assimilated insofar as substitution of one for another does not alter some feature of the inferential role of the sentences they are substituted into—paradigmatically, insofar as such substitution does not turn any materially good inferences those sentences are involved in into materially bad inferences. It is this methodology for carving up sentences into semantically significant subsentential parts by noting inferential invariants under substitution that he elabo-

rates into the theory of *functions,* which throughout his later work he takes to be one of his greatest intellectual contributions.

In Chapter 6 Frege's notion of substitution is used to investigate the fine structure of the inferential roles conferred on sentences (and the states whose contents they express) by discursive scorekeeping practices of the sort already discussed. The key notion is that of substitution *inferences:* those whose conclusions are substitutional variants of their premises. The scorekeeping significance of the occurrence of singular terms and predicates is explained by their association with substitution-inferential *commitments.* These are commitments to the propriety of a particular range of substitution inferences: those that exhibit a certain pattern preserving the deontic status of doxastic or practical commitment in the passage from premise to conclusion. In these terms it is possible to explain what it means for propositionally contentful sentences to represent objects as exhibiting properties and standing in relations—and as a consequence to offer an account of the distinction between extensional and intensional sentential contexts in which such subsentential expressions may occur. The official theoretical semantic apparatus is accordingly deepened and made more powerful by adding to the notion of *inference* that of *substitution.*

Both the sentential inferential story about conceptual contents and the subsentential substitutional one rely essentially on the *repeatability* of intentional states and their linguistic expressions. Sharpening the analytic focus, Chapter 7 explores the structures by which conceptual repeatables are constructed out of unrepeatables. Descending to this level is necessary, to begin with, in order to understand the conceptual content (and so pragmatic scorekeeping significance) of deictic or demonstrative responses to the nondiscursive environment within which discursive practice is conducted. Such unrepeatable but conceptually articulated responses make a crucial contribution to empirical cognition, in particular in its relation to practical action. Indexicals (the paradigm of expression-types, tokenings of which uttered by different individuals on different occasions are grammatically guaranteed to have different significances) are only a particularly obtrusive special case of a much more general phenomenon, however—one that is fundamental to the very possibility of interpersonal communication and comprehension. Even if attention is restricted to nonindexical expression types, the ubiquity of differences in collateral beliefs or commitments across interlocutors ensures that lexically cotypical utterances issuing from different mouths will often have very different deontic significances.

It will emerge that talk of the *representational* dimension of discourse is in fact addressed to the *social-inferential* mechanisms by means of which communication can nonetheless be secured across such doxastic gaps. These mechanisms are examined in deontic scorekeeping terms in Chapters 7 and 8. Chapter 8, which discusses various forms of *ascription* of conceptually contentful deontic statuses and attitudes, looks at the expressive power of the locutions used to make various aspects of these mechanisms proposition-

ally *explicit*. As always, such explicitating vocabulary is intelligible only against the background of an understanding of the *implicit* practices that such vocabulary makes it possible to express in explicit (assertible) form.

The most important underlying structure is that of *anaphora*, the paradigm of which is the relation a pronoun stands in to its antecedent. In Chapter 7, treating one expression tokening as anaphorically dependent on another (whether they are uttered by the same or by different interlocutors) is explained as taking the substitution-inferential significance of producing the dependent tokening to depend in a certain way on the substitution-inferential significance of producing its antecedent tokening. In other words, anaphora is construed as a special mechanism for the *inheritance* of substitution-inferential commitments. The third component of the theoretical semantic idiom employed here is accordingly *anaphora*. The result is a three-tiered semantic structure: *inference, substitution,* and *anaphora* (ISA for short).

5. 'True' and 'Refers'

This chapter looks at the expressive function of the technical semantic vocabulary that has been employed by theorists outside the minority inferentialist tradition pursued here. On the side of sentences, the central locution is 'true', while on the side of subsentential expressions, it is 'refers' or 'denotes'. The discussion of the way these expressions work plays a dual role in the exposition of the account of discursive practice presented here: retrospectively as offering some of the justification for the decision to focus on *inferential*, rather than *representational*, semantic primitives (announced in Chapter 2), and prospectively as motivating the descent to *substitutional* inferential substructures in Chapter 6 (culminating in the treatment of *anaphora* in Chapter 7). For on the first, or retrospective, point, it is argued here that once the *expressive* function of 'true' and 'refers' is properly understood, it is seen to be incompatible with the *explanatory* function those locutions have been accorded in the dominant semantic theoretical tradition.[5] Insofar as this case can be made out, it makes urgent the demand for alternatives to these semantic primitives—precisely the demand that the inferentialist semantic program pursued here seeks to satisfy.

On the second, or prospective, point, the account of the expressive role of 'true' and 'refers' that is arrived at in this chapter construes them as functioning *anaphorically*. They are anaphoric proform-forming operators; the paradigmatic use of 'true' is to construct a special kind of pro*sentence*, while the paradigmatic use of 'refers' is to construct a special kind of pro*noun*. The expressive role of traditional representationalist semantic vocabulary can be rendered only in the broadly inferentialist idiom used to specify the model of discursive practice offered here, then, inasmuch and insofar as that model can be extended so as to encompass the attribution by deontic scorekeepers of specifically *anaphoric* relations between the significance of various expres-

sion tokenings. This is a promissory note that is redeemed two chapters later, after the substitutional raw materials required have been brought onboard. Once that has been done, an account will have been offered in systematic terms of why it is useful to talk about intentional contentfulness in terms of truth and reference, even though the practices that confer the contents these locutions express should be understood rather in terms of reasons and proprieties of inference.

The previous two chapters presented a model of social linguistic practices. A fundamental criterion of adequacy for such a social-practical model of discursiveness or intentionality is that it be possible to show in the idiom it provides that the normative significances those practices institute are sufficient to confer genuinely propositional contents on states, attitudes, performances, and expressions that are suitably caught up in them. To do that, it is necessary to explain how the social-inferential articulation of those practices can amount to objective representational content. Such an explanation in turn requires an account, in terms of the social and inferential articulation of the pragmatic significances instituted by deontic scorekeeping, of what we are saying when we make claims about what our claims and beliefs represent or are about. The philosophical elaboration of these representational notions uses the technical vocabulary of truth and reference. But what are we saying when we say that a claim is true, or that a term refers to an object?

It is natural, and not obviously inappropriate, to understand these notions themselves in representational terms. To do so is to take '. . . is true' to denote a property of claims and beliefs, and so of the sentences that express them, and to understand '. . . refers to . . .' as denoting a relation between words and the world, between linguistic or intentional items and nonlinguistic, nonintentional ones. Now a great many of our beliefs *are* true, and a great many of our terms *do* refer to objects. But it is necessary to be careful about what assumptions are made concerning what is being said when these claims are made. For exactly what we are saying when we make these true claims— indeed, what we are *doing* when we do so—is of critical importance for determining what *explanatory* role such claims are suited for.

The discussion of representational locutions in the rest of this work moves on two tracks: one having to do with technical representational terms employed by semantic theorists to make explicit various important features of the conceptual contents they address, the other with the locutions of ordinary language practitioners use to make explicit the representational dimension of the doxastic and practical commitments they attribute and undertake. Together these discussions seek to make clear what the representational purport implicit in our discursive practices consists in. This chapter looks to specify the actual expressive role played by the technical terms 'true', 'refers', 'denotes', and their cognates. In order to do this precisely and clearly, it seeks to show how vocabulary playing that expressive role can be introduced into and understood in terms of the fundamental model of discursive practice

offered in Chapter 3 and Chapter 4. It also argues that once this expressive role is properly understood, it becomes clear that representational locutions are not suited to play the role of primitives in a semantic theory. Nailing down this point requires looking at their actual expressive role, which is not capturable in the sort of representational semantic theory that uses these notions as primitives; it is therefore necessary to work in another one. The conclusion one arrives at by such an analysis is that some other sort of semantic theory is needed. So the broadly inferential approach introduced in Chapter 3 serves both as the use-language in which the argument is conducted and as the answer to the question that is raised by seeing the inadequacy of truth and reference locutions to serve as fundamental semantic explanatory vocabulary—namely, if semantics ought not to be done in terms of truth conditions for sentences and reference and denotation for singular terms and predicates, what semantic primitives ought to be employed?[6]

The rhetorical situation at this point is complex, and there is a danger that the conclusion that the traditional semantic primitives are inappropriate for explanatory purposes (though quite in order for expressive ones) might be taken to be a result that depends on the inferentialist semantic metalanguage being employed. In fact this thought about the use of 'true' and 'refers' can be motivated independently, and that is the tack taken here. So the analysis of traditional representational semantic vocabulary in this chapter could have been presented before inferential approaches to semantics were introduced in Chapter 2, in order to motivate the search for some new semantic primitives—since most theorists use truth and reference as much because they cannot think of an alternative as for any other reason. But this order of explanation would give too much prominence to the notions of truth and reference; the model being pursued here can be motivated apart from the troubles of this semantic tradition. The expressive role of traditional semantic vocabulary is explained here in terms of *anaphora*, and the argument that this is the right way to understand 'true' and 'refers' is conducted independently of the considerations of the previous chapters. The notion of anaphora that is appealed to is in principle equally available to the representationalist, who will think of it just as grammatically guaranteed coreference.[7] This neutral analysis is then supplemented by showing, in Chapter 7, how anaphora is to be understood according to the inferentialist model of Chapter 3.

II. TRUTH IN CLASSICAL PRAGMATISM

1. Stereotypical Pragmatism

The popular conception of the theory of truth of classical pragmatism is summed up in the slogan "The truth is what works." According to this view, the pragmatists were trying to give a theory of truth in the sense of offering necessary and sufficient conditions for possession of that property.

Their innovation is then seen to consist in taking the possession of this property by a belief to consist in a relation not simply to what is believed but also to what is desired. Working, or being satisfactory, involves a further argument place beyond the standard representational or correspondence notion, for it is relative to preferences, purposes, interests, needs, or some such satisfiable practical states.[8]

A theory of truth, on this line, is generically a pragmatic one if it treats truth as the property of conducing to the satisfaction of some state associated with the believer—paradigmatically desire. Specific versions of this genus of explanation will be distinguished by how they understand the state, its subject, and the sort of satisfactoriness involved. Thus within the pragmatic genus truth might be identified with properties as various as evolutionary adaptiveness for a species and optimality for felt-preference maximization by a time-slice of an individual agent.

This sort of understanding of truth as a property of utility for some end—a matter of how useful, in some sense, it is to hold the belief that is a candidate for truth—may be called 'stereotypical pragmatism'. It is important to notice what sort of a theory it is. Pointing out the apparent appropriateness on some occasions of questions such as, I believe that the theory works (for instance, makes correct predictions) but how do I know it is true? already shows that this sort of pragmatism is very implausible if it is conceived as elucidating our concept of truth.[9] As Dewey was well aware, views of this stripe can best be maintained as revisionary proposals—not as accounts of what we mean by 'true' but as suggestions that we stop using that concept and get along instead with the pragmatist's notion of utility.

Any assessment of the merits of such a proposal depends on an account of what the role of the concept of truth is, what explanatory uses the property of truth is wanted for. For only in that context can it be argued that some utility notion would better serve those ends or play that role. The significance of the classical pragmatists in the present story derives from their contribution to that antecedent question—the question of what expressive and explanatory work is and ought to be done by the truth concept. Although their account of the role of truth talk cannot, in the end, be counted as correct, it nevertheless provides the central idea around which an adequate account can be constructed. The answer that eventually emerges as to the role of 'true' makes it hard to see how stereotypical pragmatism, even as a revisionary proposal, can amount to anything other than changing the subject, sharing only a homonym with ordinary truth talk.

2. Five Theses of Classical Pragmatism

There is no question that the classical American pragmatists at times commit themselves to what has just been called 'stereotypical pragmatism about truth'. But there is a deeper and more interesting explanatory

strategy that the pragmatists pursue as well. According to this way of setting out their account, concern for what 'works' or is satisfactory is only the final move in an innovative rethinking of the nature of truth and belief. The early moves are worthy of attention, even though the final one does not in the end prove satisfactory. The essential point of a theory such as James's is to treat calling something true as doing something more like praising it than like describing it.[10]

Five separable theses can be distinguished in the elaboration of this approach. First is the performative, antidescriptive strategy, emphasizing the *act* of calling something true rather than the descriptive content one thereby associates with what is called true. Next is an account of that act as the personal taking up of a certain sort of normative stance or attitude. Taking some claim to be true is endorsing it or committing oneself to it. Third is a particular understanding of that stance or attitude. Endorsing a claim is understood as adopting it as a guide to action, where this in turn is understood in terms of the role the endorsed claim plays in practical inference, both in first-person deliberation and in third-person appraisal. Fourth, and least important, is the view that an advantage of understanding the appropriateness or correctness of adopting an attitude of endorsement in terms of its role in guiding action consists in the possibility for some sort of not merely subjective measure of that appropriateness, namely, the success of the actions it leads to. This is the only strand of the argument acknowledged or embraced by stereotypical pragmatism.

Finally and, it will be argued, most significantly, the theory claims that once one has understood acts of *taking-true* according to this four-part model, one has understood all there is to understand about truth. Truth is treated, not as a property independent of our attitudes, to which they must eventually answer, but rather as a creature of taking-true and treating-as-true. The central theoretical focus is on what one is *doing* when one takes something to be true—that is, our use of 'true', the acts and practices of taking things to be true that collectively constitute the use we make of this expression. It is then denied that there is more to the phenomenon of truth than the proprieties of such takings. Theories of this general sort may be called 'phenomenalist', in recognition of the analogy with the paradigmatic subjective phenomenalism concerning physical objects, whose slogan was "esse est percipi." The significance of such a move in the context of the present approach should be obvious, for taking-true is just asserting or judging. The classical pragmatist line of thought accordingly holds out the possibility of understanding the use of 'true' in terms of what we are doing when we make a claim, putting forward a sentence *as* true.

According to this decomposition of their view into five theory-features, the pragmatists start with the idea that in calling something true, one is *doing* something, rather than, or in addition to, *saying* something. Instead of asking what property it is that we are describing a belief or claim as having

when we say that it is true, they ask about the *practical significance* of the *act* we are performing in attributing that property. We accomplish many things by talking, and not all of them are happily assimilated to describing how things are. One ought not to conclude that because truth ascriptions are expressed in the same subject-predicate grammar that descriptions are, they must for that reason be understood to function as descriptions do. The pragmatic approach, centering on the act of calling something true rather than the content one thereby characterizes it as displaying, has much to recommend it. It has been seized upon by a number of authors who would not go on to accept the account of the act in question that the pragmatists offer. For, stripped of those further commitments, the recommendation is for a *performative* analysis of truth talk. In Fregean terms it is the suggestion that 'true' is a force-indicating, rather than a sense-expressing, locution.

Wittgenstein notoriously warned against thoughtless assimilation of sentence use to fact-stating, and of term use to referring. In the wake of Austin's discussions, theorists such as Strawson offered accounts of 'true' as a performative.[11] Its use was to be assimilated to other sorts of commitment-undertaking, in a way parallel to that expressed by the explicit performative 'I promise . . .' In the same spirit, other contemporary accounts were offered of 'good' as expressing a kind of commendation, as taking up an attitude or expressing one's own relation to something, rather than as describing it by attributing some objective property. This is the sort of assimilation James had been urging in saying that truth is "what is good in the way of belief." Such remarks are often misinterpreted as claiming descriptive equivalence, or coextensiveness of the predicates 'true' and 'what it is good for us to believe'. On such a reading, the allegedly uncontroversial claim "It is good for us to believe the truth,"—that is, the truth is among the things it is good for us to believe—is turned on its head. Necessary conditions are treated as sufficient, and truth is defined as whatever it is good for us to believe. James intent was rather to mark off 'true', like 'good', as a term whose use involves the taking up of a nondescriptive stance, the undertaking of a commitment that has eventual significance for action. In the vocabulary of this work, it is adopting an implicitly normative practical attitude.

What motivates such a performative analysis, for the pragmatists no less than for later theorists, is the special relation that obtains between the force or practical significance of an act of taking-true (which one might, before the performative possibility has been broached, uncritically have called an act of 'describing as' true) and the force or significance of a straightforward assertion. In asserting "It is true that *p*," one asserts that *p*, and vice versa. The force or significance of the two claims is the same. On the face of it, this redundancy or transparency of force, the fact that adding the operator 'It is true that . . .' to what one is going to assert does not change the force or significance of that assertion, might be explained in either of two different ways. One might take it that the content that is expressed in a truth ascrip-

tion is special, and that the redundancy of force of truth claims arises out of features of the property a claim or belief is said to exhibit when it is described as true. One must then go on to offer an account of why attributing that property has the consequences that it does for the force of one's attribution. This can be compared to treating claims using 'good' or 'ought' as describing properties of actions, and then needing a theory explaining the special motivational role that attributions of these properties must be taken to have for the attributor. The pragmatic theories being considered adopt the more direct path of taking the transmitted force of truth claims as the central phenomenon, one that is merely obscured by the misleading grammar of property ascription.

Dewey's assertibilist theory of truth develops these ideas along explicitly performative lines using the model of utterances of 'I claim (or assert) that p'.[12] The claim that the force of freestanding utterances of this type and of 'It is true that p' are equivalent is especially liable to misinterpretation as the claim that the contents expressed by these utterances are the same. As will be seen below, it is easy to show that that is not so. In any case, as a revisionist, Dewey did not even claim equivalence of force, though that was the dimension along which he assessed the relationship between his views and the tradition. He has accordingly often been 'refuted' on the basis of misunderstandings of theories that he did not subscribe to in the first place.

To this performative, antidescriptive explanatory commitment, the pragmatists add a particular sort of account of the act of taking-true as adopting a *normative stance* toward the claim or belief. In treating something as true, one is praising it in a special way—endorsing it or committing oneself to it. The stance is normative in involving what the claim to which one has taken up a truth-attitude is good for or *appropriately* used for. For treating something as the truth is plighting one's troth to it, not just acknowledging that it has some property. Truth undertakings are taken to be personal in that the proprieties of conduct one thereby commits oneself to depend on one's other commitments, commitments to choose (representing preference, desire, interest, need, and so on) as well as commitments to say (assert and believe). One is expressing or establishing one's own relation to a claim, in taking it to be true, rather than recognizing some independent property that claim already had. Again the model of promising is important. This important emphasis on the normative character of cognitive undertakings was their acknowledgment of the central Kantian legacy that has been rejuvenated for us by Wittgenstein. Its expression is often obscured (Peirce is, as so often, an exception) by the pragmatists' further commitment to the sort of naturalism about the norms involved that gives rise to the attribution to them of stereotypical pragmatism.

Their understanding of the sort of commitment undertaken in taking-true is as a commitment to rely on the belief or claim in question in guiding practical activity. This in turn is understood as a commitment to using the

claim as a premise in practical inferences, whose conclusions are not further claims but actions—that is, performances under a description that is privileged by its relation to deliberation and appraisal. Relative to the truth-taking commitment, one ought to reason practically in one way rather than another. The proprieties of practical inference concerning whether to bring an umbrella are different for one who takes-true the claim that it is raining than for one who does not. The force of such proprieties is normative in that, although they may be ignored, the significance or force of the agent's commitment is to the effect that they *ought* not to be. It is these prudential 'ought's that appraisal of actions assesses.

Thus the stance or attitude that one adopts in treating something as true is to be understood by its role in orienting action when activated by a contextualized attempt to satisfy the desires, preferences, and so forth that one finds oneself with. The discussion of action in the previous chapter indicates that such a naturalistic, instrumental picture of the 'ought's involved in action and belief is not the only candidate. But the fact that the classical pragmatists go on to offer a naturalistic account of norms along these lines should not distract attention from their appreciation that the phenomenon of taking-true they address concerns *normative attitudes*, and that attributing truth is for them attributing a *normative status*.

Pragmatism in the stereotypical sense arises when one conjoins the ideas of a performative analysis of taking-true, of the relevant kind of performance as undertaking a personal commitment, and of the commitment as specifying the appropriate role of a claim in action-orienting deliberation, with the further idea that the measure of the correctness of the stance undertaken by a truth-attributor is the success of the actions it guides. Although the point will not be pursued here, it may be noted in passing that it is not easy to say what relation to the success of individual actions some substantive property of beliefs ought to stand in for it to play the sort of explanatory role that classical pragmatism envisions. It is by no means the case that for each belief in each deliberation, an agent is more likely to succeed in accomplishing the desired result if that belief is true than if it is false. The success of a plan can be thwarted by collateral false beliefs, and in the context of such beliefs, it may actually be helpful to have a further, compensating false belief. Someone who desires to cross a ravine and has formed the plan of felling a tree if one tall enough can be found near the edge will be more likely to succeed if equipped with the false belief, of the only likely candidate tree, that it is more than forty feet tall, than with the true belief that it is only thirty feet tall, in the case where that individual also believes falsely of the twenty-five-foot-wide ravine that it is forty feet across. Although the possibility of this sort of case is no reason to prefer in general to act on false beliefs, it does show that success is associated with a property of *sets* of beliefs, rather than individual beliefs. Where this point is acknowledged, the desired pragmatist principle is often formulated as the claim that "if all our beliefs were true, then the actions we undertook on those beliefs would satisfy our desires."[13]

But even this does not seem right. To ensure success, it would be necessary to ban not only error but ignorance, for facts of which the agent is unaware can lead to failure, even if all the agent's beliefs are correct. Of course this will not be true if those beliefs include one to the effect that there are no circumstances of which the agent is unaware that would thwart the plan being adopted. But including such beliefs threatens the principle with triviality in the same way that including among the beliefs stipulated to be true the belief that the plan will succeed would threaten it. So it is not clear how the vague pragmatist pronouncement that "the truth is what works" can be made more precise, without either falsifying or trivializing it. Fortunately, for the purposes of this chapter it is not necessary to resolve this issue.

3. Phenomenalism

Nonetheless, pragmatism is normally identified with the claim that the measure of the correctness of the stance undertaken by a truth-attributor is the success of the actions it guides. But the explanatory role played by this most notorious of the pragmatists' tenets ought to be understood in the light of the larger strategy for relating the concepts of truth and belief that it subserves. From a methodological point of view, perhaps the most interesting feature of the pragmatic approach is its commitment to phenomenalism about truth.[14] Only in the context of a phenomenalist explanatory strategy can commitments of the first three sorts be seen as illuminating the notion of truth. For what they really supply is a theory of *taking*-true. It is in the overarching commitment to the effect that once one understands what it is to take or treat something as true, one will have understood as well the concept of truth that the phenomenalism of this strategy consists.

The force-redundancy approach to truth emphasizes the practical equivalence of taking something to be true and believing it, so another way of putting the point is the following: instead of starting with a metaphysical account of truth, such as that of the correspondence theorists, in opposition to which the pragmatists defined themselves, and employing that in one's account of beliefs, which are then conceived as representations that could be true, that is, have the property previously defined, the pragmatists go the other way around. They offer an account of believing or taking-true, characterized by the three sorts of commitments already canvassed, that does not appeal to any notion of truth. Being true is then to be understood as being *properly* taken-true (believed). It is this idea that is built on here, jettisoning the details of the classical pragmatist account of belief or taking-true, and substituting for it the account of assertion and doxastic discursive commitment introduced in Chapter 3.

From this point of view, what is of most interest about the classical pragmatist stories is not stereotypical pragmatism but the dual commitment to a *normative* account of claiming or believing that does not lean on a supposedly explanatory antecedent notion of truth, and the suggestion that

truth can then be understood phenomenalistically, in terms of features of these independently characterized *takings*-true. The sort of explanatory strategies here called 'phenomenalist' in a broad sense treat the subject matter about which one adopts a phenomenalistic view as *supervening* on something else, in a way whose paradigm is provided by classical sensationalist phenomenalism about physical objects. The slogan of this narrower class of paradigmatically phenomenalist views is, "To be is to be perceived." The characteristic shift of explanatory attention enforced by these approaches is from what is represented to representings of it. The representeds are explained in terms of the representings, instead of the other way around. Talk ostensibly about objects and their objective properties is understood as a code for talk about representings that are interrelated in complicated but regular ways. What the naive conservatism implicit in unreflective practice understands as objects and properties independent of our perceptual takings of them now becomes radically and explicitly construed as structures of or constructions out of those takings. Attributed existence, independence, and exhibition of properties are all to be seen as features of attributings of them.

The general structure exhibited by this sort of account is that the facts about *having* physical properties are taken to supervene on the facts about *seeming* to have such properties. Or in the vocabulary to be preferred here, the facts about what things *are K*s, for a specified sortal *K*, supervene on the facts about what things are *taken* to be *K*s. According to such an explanatory strategy, one must offer first an independent account of the takings—one that does not appeal in any way to what it is to be a *K* in order to explain what it is to take something to be one. Thus classical phenomenalism concerning physical properties such as *red* found itself obliged to account for states of the attributing subject in which things *look-red* or *seem-red* without invoking the redness that is attributed in such takings. Once that obligation is satisfied, it can further be claimed that there are no facts about what things are red, or what it is for things to be red, over and above all the (possible) facts about what things look or seem red.

Classical subjective phenomenalism regarding physical objects and properties notoriously failed in both of these component explanatory tasks. Cartesian mental acts seemed ideal candidates for the takings in question. This ontological category had been given an epistemic definition in terms of the privileged access (in the sense of transparency and incorrigibility) subjects have to the class of takings that includes perceptual seemings. That something could not *seem* red to a subject who did not by virtue of that very taking *know* that it seemed red, and that something could not merely *seem* to seem red without *really* seeming red, made this class of takings appear well suited to provide the independently characterized base of a supervenience relation. Their special epistemic status seemed to guarantee for these subjective takings or attributings the possibility of a characterization independent of what they take or attribute. For one knows all about these states

just by having or being in them, apart from any relation to anything but the knowing subject and the known mental state.

4. 'Looks' Talk and the Errors of Subjective Phenomenalism

But this way of understanding 'looks' is a mistake. As various authors have shown (the locus classicus is Sellars's "Empiricism and the Philosophy of Mind"), 'looks'-talk does not form an autonomous stratum of language: a game one could play, though one played no other. When one understands properly how the 'seems' operator functions, one sees that the incorrigibility of such claims essentially arises from their *withholding* of the endorsements involved in unqualified claims about how things actually are. As was indicated already in Chapter 4, noninferential reporting practices involve two distinguishable components: a reliable differential responsive disposition and the capacity to exercise that disposition by *endorsing* a claim. It is this second component that puts the response into the inferentially articulated space of applications of *concepts*, by bringing it into the game of giving and asking for reasons. Sellars's account of 'looks' or 'seems' is that these locutions are introduced *after* such a practice is under way, as a way of dealing with systematic sources of perceptual error that one becomes aware of through having to withdraw unreflective noninferential claims on the basis of their incompatibility with commitments one is otherwise entitled to (noninferentially, inferentially, or by testimony). One can then *express* the responsive disposition to call something red—a disposition to endorse a claim (undertake a doxastic commitment) *if* one did not know better (have reason to think that commitment is liable to be flawed)—without actually endorsing the claim one is acknowledging a certain temptation to endorse. This is how one begins to become reflective about the implicit appropriate circumstances of reporting. Thus in saying how things *look*, one is withholding an endorsement and so is naturally incorrigible.

In arguing for this diagnosis, Sellars points to two sorts of uses of 'seems' or 'looks' that it makes sense of and that are unintelligible if one construes them as minimal, incorrigible-because-unconceptualized reports on immediate experience. One is *generic* looks claims: the hen looks to have a number of spots, but there is no particular number of spots she looks to have; the polygon looks many-sided, but it does not look to be 998-sided, and does not look to be 999-sided, and so on. For Sellars, these cases are matters of what commitment one is prepared to endorse and defend. One can be willing to endorse the claim that an object is red without being willing to go further and endorse the claim that it is scarlet—and similarly for all the other shades of red. One withholds endorsement from the more specific claim, while endorsing the generic one. On traditional accounts, one must either jettison the transparency and incorrigibility of our knowledge of our own experiences or sense data or else treat these things as unlike ordinary objects in some-

times being *merely* generic. Such strenuous metaphysical exertions are obviated by Sellars's analysis.

The other use he points to is the difference between:

> It looks as though there is a tree over there.
> There is something that looks to be a tree over there.
> There is a tree over there.

The different scope of the 'looks' operator corresponds to the different scope of the commitments undertaken and withheld. In the first case, one endorses nothing about the claim, merely evincing the noninferential disposition to apply concepts one in the event refuses to apply. In the second case, one endorses the existence of something over there and evinces one's temptation to call it a tree, while resisting that temptation and not endorsing that characterization. In the third case, one endorses the whole claim. One is 'incorrigible' exactly as far as one withholds endorsement. But of course one cannot withhold endorsement unless one can grant it, and to do that one must be in the game of making ordinary corrigible noninferential reports. The very incorrigibility that recommended 'seems' statements as a basis in terms of which everything epistemically less certain could be understood turns out to be an expression of the parasitic relation that these withholdings of endorsement have to the risky practices of endorsement from which they derive their meaning, by contrast to which they exhibit their special status. Whatever may be their role in the order of *justification*, in the order of *understanding* 'seems-red' presupposes 'is-red'.

It may be noted in passing that Sellars's analysis applies to first-person uses of 'looks' (or more generally 'seems'). His endorsement-withholding diagnosis is confirmed by consideration of third-person, or attributing, uses. For in that case an ambiguity arises: the claim "It (merely) seems to S that *p*" may be used either to attribute S's withholding of an endorsement (the state S would express in propria persona by uttering "It seems to me that *p*") or to attribute the undertaking of a commitment to S (the state S would express by claiming that *p*) while expressing the *attributor's* withholding of endorsement of the claim that *p*. In this sense, 'seems' is used to express a hybrid deontic state dual to that expressed by 'true': the first corresponds to "*S incorrectly* believes that *p*," and the second to "*S correctly* believes that *p*"—in the sense of 'correct' that has to do with commitment, rather than entitlement.

Although Sellars does not say so, an entirely parallel analysis applies on the side of noninferential practical discursive *exits*. Just as a doxastic commitment whose success one is not prepared to endorse can be denominated a mere *seeming*, so a practical commitment whose success one is not prepared to endorse can be denominated a *trying*. Just as while it can seem to one that something is red without its actually being red, but it cannot merely seem to one that it seems to one to be red without its actually seeming so,

so while one can try to lift the weight without actually lifting it, one cannot merely try to try to lift it, without actually trying to do so. Neither seemings nor tryings iterate nontrivially. Just as the first point can be put (with substantial danger of misleading in a Cartesian direction) by saying that the distinction between appearance and reality does not apply to appearances, so the second can be put (subject to the same danger) by saying that the distinction between willing and succeeding does not apply to willings. The danger is that one will be tempted to think of seemings as cognitions about which one cannot be in error, and of tryings as actions one cannot fail to accomplish. The mind can then be thought of as consisting in these two sorts of representings that cannot fall short of what is represented (which in turn can be understood only by the coincidence in this case of representing and represented): states such that just by having them one counts on the one hand as knowing something and on the other as doing something. On this rendering, the mind extends just as far as its total cognitive and practical dominion; its limits are scribed by the inapplicability of the concepts of error and failure.

The cost of treating these degenerate cases as representational paradigms is to render unintelligible in the ordinary fallible cases the relation between doxastic or practical representings and the represented states of affairs known or brought about by them. These temptations are best avoided by correctly diagnosing the source of the noniterability of the 'seems' or 'tries' operator, which is the phenomenon that originally motivates this disastrous metaphysics of the mental. For then grasp of what is expressed by both 'seems' and 'tries' talk is seen to depend on grasp of what is expressed in ordinary fallible 'is' talk; one cannot withhold endorsements one cannot undertake or attribute, and a further disavowal of an endorsement once disavowed is without effect. So understood, neither the cognitive infallibility of seemings nor the practical infallibility of tryings is eligible to serve a foundational role. A subject conceived as contracted to these activities alone cannot be coherently thought of as grasping or accomplishing anything, hence not as a subject at all.

Because in the order of understanding, grasp of what is expressed by concepts of the form 'seems-K' presupposes grasp of what is expressed by corresponding concepts of the form 'is-K', the classical phenomenalist basis of takings as subjectively certified seemings cannot be secured with the autonomy from the properties taken to be exhibited that is requisite for the subsequent framing of phenomenalist supervenience explanations. Those explanations have troubles, however, apart from those regarding their basis. Generic phenomenalism has been characterized here in terms of supervenience. The sense intended is that one vocabulary supervenes on another just in case there could not be two situations in which true claims (that is, facts) formulable in the supervening vocabulary differed, while the true claims formulable in the vocabulary supervened on do not differ. More neutrally (but in the end equivalently) put, once it is settled what one is committed to as

expressed in the one vocabulary, then it is settled what one is committed to as expressed in the other.

Classical subjective phenomenalism about physical objects and properties typically makes stronger, reductionist claims, involving further commitments beyond supervenience. These regard the equivalence of sentences (or in the most committive cases, individual terms and predicates) in physical object talk to sentences (or terms and predicates) constructible in the language of takings-as-seemings. Again, the sort of equivalence in question might vary from the extreme of definitional or translational equivalence down to mere coextensiveness. In none of these forms are phenomenalist claims of this reductionist variety plausible today (as Quine argues in "Epistemology Naturalized").[15] Attempts to work out these reductive phenomenalist strategies have shown that the conditions under which there are reliable connections between how things seem perceptually and how they are can themselves be stated only in terms of how things are. The inference from things seeming red to their being red depends on there being *in fact* no filters, strange lights, retina-altering drugs, and so on. That there not *seem* to be such is far from sufficient.[16]

These explanatory failures of phenomenalism in the narrow sense ought not to be taken to impugn the prospects of phenomenalist strategies in the broad sense. For those difficulties arise from the way its general phenomenalist commitments are specialized: in applying to perceivable physical properties, in offering an account of the relevant sort of takings as incorrigible subjective perceptual seemings, and in insisting on reduction rather than just supervenience as the relation between them. Phenomenalism in general is a structure that antirealist accounts of many different subject matters may exhibit. It elaborates one way of taking seriously what Dummett calls the issue of "recognition transcendence." To detail a specific version of this sort, one must specify three things: what it is that one is taking a phenomenalist approach to (for example, physical objects, mental activity, semantic properties, the past, and so forth), how one conceives the takings or attributings on which talk of such things is taken to supervene, and how in particular the supervenience relation is conceived. Corresponding to each specific phenomenalist claim of this sort will be a class of claims that qualify as *realist* in the sense of denying the phenomenalist's "nothing but" account of the subject matter in question. For the classical pragmatist the facts about what is true supervene on the facts about taking-true, that is, on the actual action-guiding roles of beliefs. In order to appreciate the significance of the pragmatists' phenomenalist strategy, one must first consider the development of the basic idea that truth locutions are force-indicating, rather than content-specifying. The subsequent trajectory of this idea will be reconstructed, before returning to the issue of phenomenalism about truth. In the end, the phenomenalism of the classical pragmatists about truth will appear as a special case of the way in which normative statuses have been taken to be instituted by

normative attitudes, so that talk of commitments can be traded in for talk of undertaking and attributing commitments.

Before going on to see what is wrong with pragmatic phenomenalism about truth, it is worthwhile to recall briefly what can make it attractive. The account of knowledge claims offered in Chapter 3 depends on this approach. It can now be seen that that account should be seen as propounding a certain kind of phenomenalism about knowledge, which depends on a corresponding phenomenalism about truth. Its primary focus is not on knowledge itself but on *attributions* of knowledge, attitudes toward that status. The pragmatist must ask, What are we *doing* when we say that someone knows something? According to a phenomenalist reconstruction of the classic justified-true-belief account of knowledge (developed above in 4.1.2), in taking someone to know something, one first of all *attributes* a *commitment*, that is, takes someone to believe. One further *attributes entitlement* to that commitment, that is, takes the committed subject to be justified.

What, then, is the function of the truth condition on knowledge? Conventionally, treating taking the claim that the subject is committed to as true is understood as attributing some property to it, characterizing it or describing it. But it has already been pointed out that the pragmatist's account of taking the claim to be true is as *acknowledging* or *undertaking* a commitment to it. The truth condition does not qualify the entitled commitment that is attributed but simply indicates that the attributor of knowledge must endorse it. This is a deontic attitude that differs in its *social* perspective. Attributions of knowledge have the central linguistic status that they do because in them commitment to a claim is *both* attributed and undertaken. This phenomenalist distinction of social perspective, between the act of attributing and the act of undertaking a commitment, is what is *mis*taken for the attribution of a descriptive property (for which an otiose metaphysics then appears to be required).

A pragmatic phenomenalist account of knowledge will accordingly investigate the social and normative significance of acts of attributing knowledge. The account of taking-true being considered is what makes possible such a way of thinking about knowledge claims. The account of asserting introduced in Chapter 3 can be extended to such an account of knowledge, via the notion of knowledge *claims*, only if something can be made of the phenomenalist strategy of classical pragmatism about truth. There are some serious obstacles, however, to pursuing this strategy.

5. Embedded Uses of 'True' and the Failure of the Pragmatists' Analysis

On the pragmatic line being considered, it is the practical significance or force of asserting that defines taking-true, and this sense of

taking-true accounts for our use of 'true'. In spite of all that there is to recommend such a hypothesis, this conjunctive thesis cannot be correct as it stands. A familiar point of Frege's shows the inadequacy of the basic pragmatic claim. Truth talk cannot be given a purely pragmatic rendering because not all uses of '. . . is true' have assertoric or judgmental force. The force-based approach can at most account for a subset of our uses of truth locutions. Frege drew attention to the use of sentences as components of other sentences. Assertion of a sentence containing another sentence as a component is not in general assertion of the embedded sentence. That is, the embedded sentence does not occur with assertional force, does not express something the assertor of the containing sentence is thereby committed to.

As the antecedent of a conditional, 'It is true that p,' for instance, cannot have the significance of a taking-true if that is understood as the expressing of assertional force. In this sense one does not take-true the claim that p in asserting, 'If it is true that p, then it is true that q'. Of course this is just the point that was urged in Chapter 3 against traditional accounts that identify assertion as a kind of predication or representation. In those cases as in this one, the embedded uses show that what is at issue has to do with the content expressed, rather than the force attached to the speech act in which that content is expressed.

The pragmatic approach, then, offers an account only of the freestanding uses of sentences formed with '. . . is true', not the embedded ones. This is the same rock on which, as Geach has shown,[17] performative accounts of the use of 'good' have foundered. It is precisely because one cannot embed, say, questions and imperatives as antecedents of well-formed conditionals (in which they would occur without their characteristic force) that their significance as askings and commandings is associated with their force—and so is not to be understood as a feature of the descriptive content they express. If the essence of calling something good consisted in doing something rather than saying something, then it should not be possible to say things like, "If that is good, then one ought to do it." That one can sensibly say things like this shows that 'good' has descriptive content that survives the stripping away by embedding of the force associated with freestanding describings.

So an embedding test can be treated as criterial for broadly descriptive occurrences of expressions. According to this test, 'It is true that p' has nonperformative uses that the pragmatists' approach does not account for. And it is not open to the pragmatist simply to distinguish two senses of truth claims, one freestanding and the other embedded, and proceed from ambiguity. For on such a line one would be equivocating in inferring from the freestanding 'It is true that p' and the conditional 'If it is true that p, then it is true that q', in which it occurs embedded, that it is true that q, by detachment. So the pragmatic theory must be rejected and the phenomena it points to otherwise explained.

This sort of objection surfaces in many forms. Those who incorrectly take

Dewey to have offered an analysis of 'true', rather than a candidate replacement notion, must thereby treat his assertibilism as the assertion of an equivalence of content between the sentences 'It is true that p' and the explicit performative 'I (hereby) claim that p'. The most such a made-up thinker would be entitled to claim is that the force of the freestanding utterance of these sentences is the same. The stronger theory is refuted by noticing that 'It is true that p' and 'I claim that p' behave differently as embedded components. For instance, they are not intersubstitutable as the antecedents of conditionals, saving the inferential role of the resulting compound.

So an account such as is often attributed to Dewey is subject (as Putnam has pointed out in different terms)[18] to a version of Moore's naturalistic fallacy argument. Not everyone who is committed to the conditional 'If it is true that p, then it is true that p' is committed also to the conditional 'If I claim that p, then it is true that p'. If we like, we can put this point by saying that there is nothing self-contradictory about the claim 'It is possible that I claim that p, and it is not true that p'.[19] The naturalistic fallacy point is thus just another way of putting the objection from embedding.

III. FROM PRAGMATISM TO PROSENTENCES

1. Redundancy

Pointing to the sentential embedded uses of '. . . is true' shows the inadequacy of the pragmatists' attempt to make do with a notion of taking-true as asserting. Analyzing and identifying uses of truth locutions by means of *redundancy of force* (that is, by a formal property of the pragmatic significance of acts of asserting freestanding truth claims) is not a sufficient explanatory strategy. It is not that freestanding force redundancy is not a central phenomenon of truth talk. But not all uses of truth locutions take this form. More is required for an account of the use of 'true' than can be provided simply by an account of taking-true as asserting or undertaking an assertional commitment. The pragmatic account cannot for this reason be the whole truth.

Rather than simply discarding that approach, it is possible to amend it so as to retain the pragmatic account for the freestanding uses to which it properly applies. For there is a more general redundancy view that has the force redundancy of freestanding truth-takings as a consequence. Embedded uses can be explained by a notion of *redundancy of content* according to which (apart from niceties having to do with type/token ambiguities) even in embedded contexts 'It is true that p' is equivalent to p. For even their embedded occurrences are equivalent as antecedents of conditionals, in the sense that anyone who is committed to 'If it is true that p, then q' is thereby committed to 'If p then q' and vice versa. Furthermore, intersubstitutability

of 'it is true that p' and p in *all* occurrences, embedded or not, is sufficient to yield force redundancy in freestanding uses as a consequence. If two asserted contents are the same, then the significance of asserting them in the same pragmatic context should be the same. On such a content redundancy view, the pragmatists have simply mistaken a part for a whole.

Redundancy views such as Ramsey's accordingly provide a generalization of the pragmatist's point, one that permits an answer to the otherwise decisive refutation offered by the embedding objection. Accounts that generalize to the intersubstitutability of 'Snow is white' and 'It is true that snow is white' are clearly on the right track. They show what is needed to supplement the pragmatists' account in order to deal with some embedded occurrences. But they do not yet account for all the contexts in which the taking-true locution '. . . is true' occurs. Such simple redundancy accounts will not offer a correct reading of sentences like 'Goldbach's conjecture is true'. For this sentence is not interchangeable with 'Goldbach's conjecture'. For instance, the former, but not the latter, appears as the antecedent of well-formed and significant conditionals. So content redundancy, while relaxing the limitations constraining the original pragmatic account, will not apply correctly in all the contexts in which truth locutions occur.

2. Disquotation

Such cases show that the content redundancy view must in turn be revised to include the operation of some sort of disquotation or unnominalizing operator. In the cases to which the simple content redundancy theory applies, the additional operation will be transparent. But in the case of sentences such as 'Goldbach's conjecture is true', the claim with respect to which the truth-taking is content redundant must be determined by a two-stage process. First, a sentence nominalization is discerned. This may be a description such as 'Goldbach's conjecture', a quote-name such as 'Snow is white', a 'that'-clause sortal such as 'the claim that snow is white', or other sort of nominalization. Next, a sentence is produced that is nominalized by the locution picked out in the first stage. This is a sentence expressing Goldbach's conjecture, named by the quote-name, one which says *that* snow is white, and so on. It is this sentence that is then treated by theory as intersubstitutable with the truth-attributing sentence, whether occurring embedded or freestanding.

A content redundancy account with disquotation or unnominalization is more satisfactory and deals with more cases than a simple content redundancy account does—just as content redundancy accounts represent improvements of theories acknowledging only redundancy of force. But even disquotational views will not account for all the uses of '. . . is true' that might be important. They will not deal correctly, for instance, with occurrences such as 'Everything the policeman said is true', in which a *quantified*

sentence nominalization is employed. For here what is nominalized is a whole *set* of sentences, and there need in general be no single sentence that is equivalent to all of them. A further refinement of content redundancy accounts is required if they are to be able to deal with this range of cases.

3. Anaphora and Prosentences

The most sophisticated version of the redundancy theory, one capable of handling quantificational truth idioms, is the remarkable anaphoric analysis undertaken by Grover, Camp, and Belnap in their essay "A Prosentential Theory of Truth."[20] For the original redundancy and disquotational theories, each use of '. . . is true' is associated with some sentence on which it is redundant or with which it shares its content. Whatever else this may mean, it at least includes a commitment that the intersubstitution of the sentence containing 'true' and its nonsemantic equivalent, in some privileged range of contexts, preserves assertional and inferential commitments.

The difficulty in extending this intersubstitutional account to the quantificational case is that there the use of the sentence containing 'true' is determined not by a single sentence but by a whole *set* of sentences, those expressing whatever the policeman has said. Of course disquotation or unnominalization may produce sets of sentences as well, as more than one sentence may express Goldbach's conjecture. But in such cases the sentences must all share a content or be redundant on each other—that is, must be intersubstitutable with each other in the relevant contexts—whereas there is no requirement that any two sentences that express things the policeman has said be in any other way equivalent. So what is it that the sentence containing 'true' shares its content with, or is redundant upon in the sense of intersubstitutability? What is distinctive of the anaphoric development of redundancy theories is its use of the model of *pronouns* to show how, in spite of this difficulty, the quantificational cases can be treated both as redundant in the same way nonquantificational cases are and as deriving their content from a whole *set* of nonintersubstitutable sentences.

It has been noticed that pronouns serve two sorts of purposes.[21] In the *lazy* use, as in 'If Mary wants to arrive on time, she should leave now', they are replaceable by their antecedents, merely avoiding repetition. In the *quantificational* use of pronouns, as in 'Any positive integer is such that if it is even, adding it to one yields an odd number', such replacement clearly would change the sense. 'If any positive number is even, adding any positive number to one yields an odd number' is not a consequence one becomes committed to by undertaking the original claim. In such cases, the semantic role of the pronoun is determined by a set of admissible substituends, which is in turn fixed by the grammatical antecedent (here 'any positive number'). Asserting the original sentence commits one to each of the results of replacing the

pronoun 'it' in some occurrence by some admissible substituend, that is, some expression that refers to a positive number.[22]

The prosentential theory of truth is what results if one decides to treat '... is true' as a syncategorematic fragment of pro*sentences* and then understands this new category by semantic analogy to other proforms, in particular to pronouns functioning as just described. So 'Snow is white is true' is read as a prosentence of laziness, having the same semantic content as its anaphoric antecedent, perhaps the token of 'Snow is white' that it contains. The prosentence differs from its antecedent in explicitly acknowledging its dependence upon an antecedent—as 'She stopped' differs from 'Mary stopped' when the pronoun has some token of the type 'Mary' as its antecedent. Otherwise, the lazy uses are purely redundant.[23] The advance on earlier conceptions lies in the availability on this model of *quantificational* uses of prosentences containing 'true'. 'Everything the policeman said is true' is construed as containing a quantificational prosentence, which picks up from its anaphoric antecedent a set of admissible substituends (things the policeman said). Expanding the claim in the usual way, to 'For anything one can say, if the policeman said it, then it is true', explicitly exhibits 'it is true' as the quantificationally dependent prosentence. Each quantificational instance of this quantificational claim can be understood in terms of the lazy functioning of prosentences, and the quantificational claim is related to those instances in the usual conjunctive way.

By analogy to pronouns, prosentences are defined by four conditions[24]:

1. They occupy all grammatical positions that can be occupied by declarative sentences, whether freestanding or embedded.
2. They are generic, in that *any* declarative sentence can be the antecedent of some prosentence.[25]
3. They can be used anaphorically either in the lazy way or in the quantificational way.
4. In each use, a prosentence will have an anaphoric antecedent that determines a class of admissible sentential substituends for the prosentence (in the lazy case, a singleton). This class of substituends determines the significance of the prosentence associated with it.

There are many philosophical virtues to explicating each occurrence of 'true' as marking the use of a prosentence in this sense. Quite varied uses, including embedded ones, of expressions involving 'true' in English are accounted for by means of a unified model. That model is in turn explicated by appeal to the familiar and closely analogous pronominal anaphoric reference relation. Not only is the semantics of such uses explained, but their pragmatic features are as well—namely, acknowledgment of an antecedent and the use of truth locutions to endorse or adopt someone else's claim. Tarski's biconditionals are appropriately underwritten, so the necessary condition of adequacy for theories of truth that he establishes is satisfied.

A feature dear to the hearts of the prosententialists is the metaphysical

parsimony of the theory. For what in the past were explained as attributions of a special and mysterious *property* (truth) to equally mysterious bearers of truth (propositions) are exhibited instead as uses of grammatical proforms anaphorically referring only to the sentence tokenings that are their antecedents. A further virtue of the prosentential account is that anaphora is a relation between *tokenings*. Consequently the use of tokenings of types such as 'That is true' as a response to a tokening of 'I am hungry' is construed correctly—just as 'he' can have 'I' as its antecedent without thereby referring to whoever uttered 'he'. An incautiously stated content redundancy theory would get these indexical cases wrong. Finally, the uses of 'true' falling under the elegant, anaphorically unified treatment include quantificational ones such as 'Everything the oracle says is true', which are recalcitrant to more primitive redundancy and disquotational approaches.

The classical pragmatists' insistence that in calling something true, one is not describing it is respected. For one does not describe a cat when one refers to it pronominally by means of an 'it'. This point is further broadened to accommodate embedded uses where the account of the describing alternative as endorsing does not (as emerged above) apply. 'True' functions anaphorically and not descriptively even in such cases. And anaphoric inheritance of content explains equally why freestanding or force-bearing uses of 'It is true that p' have the pragmatic significance of endorsements of the claim that p. The prosentential account shows how the pragmatists' insights can be preserved, while accounting for the uses of 'true' that cause difficulties for their original formulation. It is accordingly a way of working out the content redundancy rescue strategy.

4. Prosentence-Forming Operators

The treatment of quantificational prosentences represents an advance over previous redundancy theories. As the theory is originally presented, however, the treatment of lazy prosentences in some ways retreats from the ground gained by disquotational developments of redundancy theories. The explanatory costs associated with the original theory arise because it treats *most* occurrences of 'true' as quantificational. Thus the official version of 'The first sentence Bismarck uttered in 1865 is true' construes it as a quantified conditional of the form 'For any sentence, if it is the first sentence Bismarck uttered in 1865, then it is true', in which 'it is true' is a prosentence of quantification.

One of the strengths of the prosentential account is its capacity to use the logical structure of quantification to explain the use of complicated sentences such as 'Something John said is either true or has been said by George'. There should be no quarrel with the author's treatment of these sentences that "wear their quantifiers on their sleeves." And it is clear that any sentence that has the surface form of a predication of truth of some sentence nominalization can be construed as a conditional propositional quantifica-

tion. But it is not clear that it is a good idea to assimilate what look like straightforward predications of truth to this quantificational model. To do so is to reject the disquotational treatment of these lazy prosentences, which has no greater ontological commitments and stays closer to the apparent form of such sentences. Otherwise almost all sentences involving 'true' must be seen as radically misleading as to their underlying logical form. The account of truth talk should bear the weight of such divergence of logical from grammatical form only if no similarly adequate account can be constructed that lacks this feature. It would be preferable to follow the treatment of sentence nominalizations suggested by disquotational generalizations of redundancy theories.

In fact there is no barrier to doing so. The original motivations of the prosentential account carry over directly to a disquotational or unnominalizing variant. According to such an account, 'The first sentence Bismarck uttered in 1865' is a sentence nominalization, a term that picks out a sentence tokening. In this case it describes the sentence, but it could be a quote-name, demonstrative, 'that'-clause sortal, or any sort of nominalization. Its function is just to pick out the antecedent on which the whole prosentence formed using 'true' is anaphorically dependent, and from which it accordingly inherits its content. Ontological commitment is just to sentence tokens and to anaphoric dependence, which prosententialists require in any case.

A brief rehearsal of the considerations leading the authors of the prosentential theory to do things otherwise will show that their reasons ought not to count against the adoption of a disquotational variant of the prosentential account. They say: "This account differs radically from the standard one since on (what we have called) the subject-predicate account 'that' in 'that is true' is always treated separately as referring by itself to some bearer of truth, whether it be a sentence, proposition, or statement. On our account cross-referencing—without separate reference of 'that'—happens between the *whole* expression 'that is true' and its antecedent."[26] Another way to put this point is that where the classical account takes a subpart of the sentence as a referring term and takes '. . . is true' as a predicate that forms a sentence from that term by characterizing its referent, according to the prosentential theory the only expression standing in a referential relation is the whole sentence, which refers anaphorically to an antecedent. There are accordingly *two* innovations put forth concerning reference in sentences like 'that is true'. The sentence is seen as an anaphoric proform and 'that' is no longer seen as a referring term. "Reference can involve either (or both) anaphoric reference or independent reference, and since people have not seriously considered the former, the possibility that *the* relation between 'that is true' and its antecedent may be that of anaphoric reference has not occurred to them. In ignoring anaphoric reference philosophers have assumed that *the* reference involved in 'that is true' is, through 'that', like that between a pronoun (say 'she' used

independently) and its referent (say Mary). Once this picture dominates, the need for bearers of truth begins to be felt, and it is then but a small step to the claim that in using 'is true' we are characterizing those entities" (emphasis added).[27]

But why should one have to choose between, on the one hand, treating the whole expression 'that is true' as a prosentence anaphorically referring to a sentence tokening from which it inherits its content and, on the other hand, treating 'that' as a referring expression (in particular a sentence nominalization) that picks out the tokening on which the whole prosentence depends? Instead of seeing '. . . is true' as a syncategorematic fragment of a semantically atomic generic prosentence 'that is true', one can see it as a *prosentence-forming operator.* It applies to a term that is a sentence nominalization or that refers to or picks out a sentence tokening. It yields a prosentence that has that tokening as its anaphoric antecedent. To take such a line is not to fall back into a subject-predicate picture, for there is all the difference in the world between a prosentence-forming operator and the predicates that form ordinary sentences. Nor does it involve commitment to bearers of truth, apart from the tokenings that play the role of sentential antecedents, which no anaphoric account can do without.

IV. REFERENCE AND ANAPHORICALLY INDIRECT DESCRIPTIONS

1. From 'True' to 'Refers'

There is a further reason to prefer the account that treats '. . . is true' as a prosentence-forming operator as here recommended, rather than as a fragment of the single prosentence, 'that is true' (functioning almost always quantificationally), as the original theory has it. Conceived in the former way, the treatment of 'true' has an exact parallel in the treatment of 'refers'. 'Refers' can be understood as a pronoun-forming operator. Its basic employment is in the construction of what may be called *anaphorically indirect definite descriptions.* These are expressions such as 'the one Kissinger referred to [represented, described, talked about] as "almost a third-rate intellect"', understood as a pronoun whose anaphoric antecedent is some utterance by Kissinger. A full-fledged pronominal or anaphoric theory of 'refers' talk can be generated first by showing how other uses of 'refers' and its cognates can be paraphrased so that 'refers' appears only inside indirect descriptions and then by explaining the use of these descriptions as pronouns formed by applying the 'refers' operator to some antecedent-specifying locution.[28]

Treating 'true' as an operator that applies to a sentence nominalization and produces a prosentence anaphorically dependent upon the nominalized sentence token, and 'refers' as an operator that applies to an expression picking out a term tokening and produces a pronoun anaphorically dependent

upon it then permits a single theory form to explain the use of all legitimate semantic talk about truth *and* reference in purely anaphoric terms. Discussion of how anaphoric dependence is itself to be understood must be postponed until Chapter 7, when all the raw materials will have been assembled. In what follows, then, a construction is presented that is intended to account for the use of expressions containing 'refers' and its cognates in natural languages, in a fashion strictly analogous to that just rehearsed for 'true'. To do so, however, it will be necessary to look a little more closely at the phenomenon of pronominal anaphora than was done in introducing prosentences, even though the official account of the sort of commitment inheritance that underlies that phenomenon cannot yet be considered.

Reference comes in two flavors: word-*world*, or extralinguistic reference of the sort invoked when it is said that the phrase 'the author of *Dreams of a Spirit-Seer*' refers to a certain actual individual (namely Kant), and word-*word* intralinguistic or anaphoric reference instanced by pronouns such as 'he' in "Wittgenstein admired Frege's work, so he traveled to Jena to talk with the great man." Intralinguistic reference of this sort has not been of much interest to philosophers (as opposed to linguists), for it has seemed natural for semantic purposes to assimilate pronouns to bound variables, and so to expect to explain anaphoric reference as grammatically guaranteed *coref*erence. This coreference is in turn to be understood in terms of the primary, word-world sense of reference to the same extralinguistic item. It will be argued in Chapters 7 and 8 that this is a mistake. For one thing, anaphoric mechanisms are more fundamental than, and are presupposed by, deictic mechanisms—one cannot have a language with indexicals such as demonstratives but without expressions functioning anaphorically, though the converse is possible. For another, the explicitly representational locutions by means of which we grasp and express the distinction between what our thought and talk is *about* and what we think and say about it all depend on anaphoric mechanisms.

The point of this section is to show how an analysis in terms of anaphoric mechanisms can provide the resources for a purely *intra*linguistic account of the use of the English sentences by means of which philosophers make assertions about *extra*linguistic referential relations. More specifically, although we can and must distinguish between our words and what the words refer to or have as their referents, the truth of claims about what we are referring to by various utterances is not to be understood in terms of a *relation* of reference between expressions and the objects we use them to talk about. Following Sellars,[29] it will be argued that 'refers' not be semantically interpreted by or as a relation and, a fortiori, not a word-world relation. Instead, 'refers' will be explained as a complex anaphoric pronoun-forming operator, by analogy (in the category of terms) to the analysis of '. . . is true' offered above. To show this, the anaphoric roles that expressions can play are botanized. This leads to the specification of a new part of speech—*anaphori-*

cally indirect descriptions. Next, a formal test is offered for identifying expressions that play the anaphoric role of indirect descriptions, and 'refer' is explained as an operator that forms such descriptions. A paraphrase strategy is then offered by means of which reference claims ostensibly in other forms can be wrestled into forms in which 'refer' appears only inside indirect descriptions.

2. Anaphoric Chains

In his seminal article "Reference and Context," Charles Chastain suggests a novel approach to the understanding of singular term reference.[30] The basic concept he employs is that of an *anaphoric chain*, a notion best approached at this point in terms of examples. Consider the discourse:

> #A man in a brown suit approached me on the street yesterday and offered to buy my briefcase. When I declined to sell it, the man doubled his offer. Since he wanted the case so badly, I sold it to him.#

Two anaphoric chains are intertwined here, one corresponding to the buyer, and one to the briefcase:

> A man in a brown suit . . . the man . . . he . . . him

and

> my briefcase . . . it . . . the case . . . it.

The phenomenon may be indicated presystematically by saying that the reference of later elements in such chains (here 'it' and 'the man') is secured only by the relations these elements stand in to the singular terms that initiate the chains in which they appear. This is the word-word (token-to-ken)[31] relation of *anaphoric reference* or *anaphoric dependence.* The presence of an anaphoric chain in a discourse signals that not all its singular terms have reference independently. Rather, some elements are related to their referents only in a derivative manner, in virtue of their anaphoric links to other expressions.

Examining the kinds of expressions that can initiate and continue such chains enables Chastain to make two important, related observations. The first concerns the significance of indefinite descriptions. Since Russell's discussions early in the century, indefinite descriptions have been treated as though they were not singular referring expressions at all, but rather to be understood by means of a quantificational paraphrase. The presence of an indefinite description often does signal existential quantification rather than singular reference as the proper semantic construal, but Chastain points out that the role of indefinite descriptions in anaphoric chains indicates that these expressions can also have a purely referential function. As in the

example above, an indefinite description can initiate an anaphoric chain, which may then be continued by pronouns or definite descriptions. And it seems clear that, in the context in which it occurs above, 'a man' purports to refer to a unique individual, namely the man in the brown suit who approached me on the street yesterday and eventually purchased my brief-case.[32]

This observation leads to Chastain's second point, which is that the reason that apparently nonquantificational uses of indefinite descriptions have not been thought of as straightforwardly referential is that they do not behave enough like proper names, the paradigm of singular terms. Except under deviant circumstances, if a proper name is used somewhere in a discourse invoking a particular referent, then other tokens of that same type which appear elsewhere in the discourse will be coreferential with it, in a sense that can be explained in terms of intersubstitution.[33] In

> #Leibniz has been called a pluralist, and he has been called a monist, but no one has ever thought of that philosopher as a materialist.#

the sense is not altered if all the other elements of the anaphoric chain are replaced by the initiating expression to which they anaphorically refer and on which they anaphorically depend. An inelegant redundancy is the only cost of replacing 'he' and 'that philosopher' by 'Leibniz'. In the case of an anaphoric chain initiated by an indefinite description, on the other hand, such a substitution of terms alters the sense of the sentences in which the substituted terms appear.

Consider:

> #A Republican senator threatened to filibuster the Wilderness bill. The senator's staff persuaded him that this action was unwise, so he left the chamber.#

The anaphoric chain of interest here is

> A Republican senator . . . The senator . . . him . . . he.

The sense of the discourse is completely altered if the initiating expression is substituted for each of the terms that anaphorically depends on it:

> #A Republican senator threatened to filibuster the Wilderness bill. A Republican senator's staff persuaded a Republican senator that this action was unwise, so a Republican senator left the chamber.#

In this passage the indefinite descriptions, although all of the same lexical type, do not purport to corefer. Each initiates a distinct anaphoric chain, and these chains may or may not involve the same individual. To continue such a chain requires either the use of a pronoun, which always continues an

existing chain, or the use of a definite description, which can either initiate or continue a chain. The fact that a chain beginning 'a SORTAL . . .' cannot be continued by repeating the initiating phrase, but can be continued with a definite description of the form 'the SORTAL . . .', is called by linguists the requirement of a "definitization transformation." One may conclude from such special requirements either that indefinite descriptions cannot function as singular referring expressions or that not all singular referring expressions must behave like ideal proper names. Chastain's suggestion is that the second alternative is worth exploring.

3. Anaphoric Roles of Tokens

Singular term tokens can play various roles in anaphoric chains. Such a token may initiate an anaphoric chain, as 'A Republican senator' does in the first example above. Or it may continue an existing chain and so depend for its referent on an anaphoric antecedent, as 'the senator' does in that example. Besides being distinguished as anaphoric initiators and dependents, tokens can be sorted according to two distinctions regarding the term types they instantiate. Chastain's considerations concerning substitution show that in dealing with anaphoric chains, one cannot in general assume that cotypical term tokens are coreferential, even in the absence of overtly indexical elements. Those expressions that (according to substitution tests such as that involved in comparing the last two examples) do not vary in reference from token to token within the type may be called *invariant under intratype or cotypical substitution*. An example would be proper names, as conceived and idealized by the tradition. Expressions that are referentially variable from token to token within the type may be described as not cotypically intersubstitutable. Pronouns would be a paradigm.[34]

The third distinction it will be useful to make is that between lexically complex expressions and those that are lexically simple, though perhaps grammatically complex—that is, between phrases that are nouns and words that are nouns. Consider, for instance, two varieties of cotypically nonintersubstitutable anaphoric dependents: dependent uses of definite descriptions and personal pronouns. The lexically simple pronoun 'he' is limited, in the information it can give about its anaphoric antecedent and the chain of which it is a part, to a small number of dimensions such as gender and number, specified in advance by the grammar of the language. Lexically complex anaphoric dependents, in contrast, can use the full descriptive resources of the language to give anaphoric information. This open-endedness permits dependents such as 'the senator mentioned above, who opposed the Wilderness bill and was dissuaded by his staff from expressing his stand'. The same contrast of lexical complexity applies to anaphoric initiators, so the indefinite description 'a Republican senator' can be compared with the proper name 'Leibniz' in the examples above.

Deploying these three independent functional distinctions—between anaphoric initiating tokens and dependent tokens, between cotypically intersubstitutable and cotypically nonintersubstitutable types, and between lexically complex and lexically simple types—yields eight roles that tokens can be thought of as playing in anaphoric chains. So among anaphoric initiators that are invariant under cotypical intersubstitution there are those that are lexically simple, such as proper names like 'Leibniz' as used above, and those that are lexically complex, such as 'the first U.S. president'. Among the cotypically nonintersubstitutable anaphoric initiators there are again the lexically simple such as 'this', and lexically complex indefinite descriptions, such as 'a Republican senator' in the example above. Among the anaphoric dependents that are not cotypically intersubstitutable one can similarly distinguish lexically simple pronouns such as 'it' from lexically complex dependent uses of definite descriptions, such as 'the man' in the very first example. Finally, among the anaphoric dependents that are invariant under cotypical intersubstitution one can distinguish some uses of lexically simple proper names, as in

> #I met a man I'll call 'Binkley'. Binkley is a mechanic.#

from lexically complex dependents, which will be called *indirect definite descriptions*.

4. Anaphorically Indirect Descriptions

With the exception of the last category mentioned, this tripartite division just rearranges familiar facts about the linguistic behavior of standard kinds of singular terms. The categorization was presented, however, to introduce the notion of indirect definite descriptions—a kind of singular term whose existence has not generally been recognized. Indirect definite descriptions are accordingly characterized as lexically complex anaphoric dependents that are invariant under cotypical intersubstitution. From this specification it follows that expressions in this category are complex pronouns, as are ordinary anaphorically dependent definite descriptions. Unlike such descriptions, however, all cotypical tokens of expressions in this category are guaranteed to be coreferential with each other, since they all anaphorically depend upon and hence corefer with a single common antecedent token, and so with each other.

The idea is that an indirect definite description is a pronoun that actually contains a description specifying the term occurrence that is its anaphoric antecedent. Cotypical tokens of an indirect definite description type contain the same description and so (except in special cases) specify the same antecedent. One immediate expressive advantage of a language containing locutions of this sort would be that identities employing anaphorically indirect descriptions could be used to assert that two term *tokens* (or tokenings) were

coreferential, even if the tokens were of cotypically nonintersubstitutable types (such as indefinite descriptions or pronouns), for which, as Chastain showed, standard substitutional accounts of coreference fail, since they presuppose invariance under cotypical intersubstitution.

A useful picture of the functioning of these expressions (the picture that motivates calling them *indirect* descriptions) is offered by the indirect addressing function offered in most basic computer architectures. Ordinarily, the central processor uses addresses to pick out values, just as we use descriptions to pick out objects. But, in indirect mode, the CPU when given an address as an input does not return the value stored at that address as its output. Instead it treats that value as another address and returns the value stored in that second address as its output. Indirect descriptions are to be understood by analogy to this two-step process. First, a token is specified, perhaps by being described as to type and spatiotemporal location. But the token thus picked out is not the referent of the whole indirect description. For next, an indirectness operator is applied to that token specification to produce the indirect description, which only *anaphorically* refers to the specified token and so, as a whole, refers not to that token but rather to whatever that token (its anaphoric antecedent) refers to—just as with ordinary pronouns. The flexibility of the von Neumann computer architecture is in large part due to its capacity to treat the same expression both as datum (that is, as a value) and as instruction (the address at which a value can be found). Indirect descriptions exploit the analogous use/mention amphibiousness made possible by anaphora. The claim to be defended is that the expressive dividend that traditional *semantic* vocabulary pays in a language to which it is added consists in its performing this function.

To become entitled to claim that there actually are expressions in natural languages that should be understood as playing the anaphoric role just abstractly described, and to see what indirect descriptions have to do with specifically *semantic* vocabulary, it is necessary to look at some examples. Consider a discourse in which Joe says:

> #I should have known better than to let the mechanic Binkley work on my car. That airhead misadjusted the valves.#

Suppose that Jim, forgetting the name Joe used, later says:

> #For car repair, don't go to the mechanic Joe referred to as 'that airhead'.#

How should this latter remark be understood? In particular, how should one understand the singular term:

> (t) the mechanic Joe referred to as 'that airhead'?

Clearly this term refers to Binkley, Joe's hapless mechanic. But how is this

reference secured? The most obvious way to interpret such a singular term is as a straightforward definite description, by analogy to

> the mechanic who worked on Joe's car and misadjusted the valves.

In both cases some purportedly unique feature of Binkley is used to single him out—his relation to Joe either in being referred to by him in a certain way or in having abused his car in a certain way.

But the anaphoric category of complex cotypically nonintersubstitutable dependents and the brief discussion of anaphorically indirect definite descriptions suggest that an alternative analysis might be more illuminating. For the term (t) can be thought of as being an anaphoric dependent, having Joe's original token of 'that airhead' as its anaphoric antecedent. If tokens of the type of (t) are anaphorically dependent on the original token of 'that airhead', then they are coreferential with it and hence refer to Binkley the mechanic. On this account, (t) should be thought of as referring to Binkley in the way that a token of 'he' would, if Jim could arrange to ensure that the antecedent of that token of 'he' were Joe's tokening of 'that airhead'. Pronouns, as simply cotypically nonintersubstitutable anaphoric dependents, can take such antecedents if the antecedent and dependent tokens are sufficiently close to each other in time, space, or audience attention. But for distant antecedents, one may not simply rely on the meager resources grammar gives us to work backward from a simple dependent token such as 'he', which even with contextual supplementation can give us only so much information about its antecedent. Here, according to the current suggestion, is where indirect definite descriptions enter. For these locutions are grammatically complex, like ordinary definite descriptions, and enable the use of the full descriptive resources of the language to specify the antecedent token to which they are anaphorically linked.

In the example, the antecedent token is specified as that token whereby Joe referred to someone as 'that airhead', the token directly picked out by the phrase "Joe's utterance of 'that airhead'." Knowing what the individual was referred to *as* specifies the *type* of the antecedent tokening. Invoking Joe locates the particular token of that type that is in question. The presence of 'refer' marks the indirect-addressing feature, by which it is specified that the referent of the whole description is to be understood not to be the term token picked out as anaphoric antecedent but rather, as with simple pronouns, the referent of that antecedent token. Indirect definite descriptions such as (t) should be understood as complex pronouns (anaphoric dependents), and 'refers' and its cognates should be understood as complex anaphoric pronoun-forming operators.[35]

The 'refer' cognates consist of all the sorts of expressions that would normally be thought of as being used to assert semantic word-world rela-

tions. So the following examples ought to be understood according to the model of indirect or anaphoric descriptions: 'the philosopher John mentioned yesterday', 'the restaurant he talked about at the committee meeting', 'the difficulty discussed above', 'the person denoted by the second name on the list', 'the criminal described by the police in the morning paper', and 'the cluster of buildings Russell called "as like as monkeys can make" (to those at Cambridge)'.[36]

5. Iterability

Although there are important differences among these examples, all of them could be paraphrased so as explicitly to use some form of 'refer'. But even this rough characterization is of use only insofar as it is possible to say what is special about the functioning of 'refer' that would enable one, for instance, to tell whether some alien language possessed an expression playing an analogous role. Putting the question more generally, even if it turns out that one can properly account for the behavior of expressions like those in the examples according to the indirect-addressing model of anaphoric descriptions, how could one explain and justify enforcing such a radical distinction between the analyses of descriptions as apparently analogous to (t) as

> (u) the one Joe startled (insulted, deafened) by his remark about airheads?

Does not the most intuitive reading of (t) assimilate it to (u), treating both as ordinary definite descriptions of a man who in each case happens to be picked out by his relation to some utterance of Joe's? What difference between these cases *makes* the difference in virtue of which (t) should be treated as an indirect description, which essentially involves an anaphoric link, whereas (u) should be treated as an ordinary description, which uses a relation to an utterance to pick out an object? What is the crucial difference between being referred to by a certain token and being startled, insulted, or deafened by it?

The clearest manifestation of the difference in question concerns the *iteration* of pronoun-forming operators. Because the relation '. . . is an anaphoric dependent of . . .' is transitive, any operator that takes a term token and produces an expression that anaphorically depends on it should be iterable without change of resulting reference (assessed, as always, substitutionally). For tokens of the complex pronoun formed by applying the indirectness operator to (a token of) the result of applying that operator to an original initiating token should simply continue the anaphoric chain—dependents of dependents having the same original or ancestral antecedent.

Consider such iteration as applied to (t) and (u). If the description-forming

operators that produced these are iterated, with suitable variation of speakers, the results are:

> (t') the one John referred to as "the one Joe referred to as 'that airhead'"
> (u') the one John startled by his remark about the one Joe startled by his remark about airheads.

If descriptions formed in the appropriate way from 'refers' are anaphorically indirect descriptions, then, in virtue of the transparent iterability of anaphoric dependence, (t') ought to be coreferential (intersubstitutable) with (t) and, hence, with Joe's original tokening of 'that airhead'. And so they are, issues of speaker's reference aside.[37]

But though (t) and (t') are coreferential de jure, the superficially analogous (u) and (u') would be coreferential only by accident and under special circumstances.[38] These considerations can be formulated as the *iteration condition* (IC) below, which is a necessary condition for understanding an operator PF as a pronoun-forming operator. Subject to the conventions that a term designation surrounded by angle brackets forms a designation $\langle\text{term}\rangle$ of the *type* of that term, and that such a type designation surrounded by subscripted slashes forms a designation $/\langle\text{term}\rangle/_i$ of a particular token(ing) of that type, the iteration condition can be expressed:

(IC) $PF(/\langle\text{term}\rangle/_i) = PF(/\langle PF(/\langle\text{term}\rangle/_i)\rangle/_j)$,

where the identity sign marks the intersubstitutability of these expressions.[39] It is clear that nothing can be thought of as a pronoun-forming operator unless it meets this condition, which embodies the transitivity of anaphora. The strategy here will be to exploit such an *iterability* requirement (suitably qualified) as a sufficient condition for identifying operators that form indirect descriptions (which have been explained as lexically complex anaphoric dependents that are invariant under cotypical intersubstitution).

Enough weight will be placed on the strategy of transforming the iterability condition from a necessary into a sufficient condition for interpreting a syntactically relational expression as a complex pronoun-forming operator to make it worth stating precisely. Consider a construction that on the surface has the form:

> (v) the x [REF $(x, /\langle\text{term}\rangle/_i)$].

The overall expression appears to be a definite description that picks out a thing x by means of its relation REF to a token $/\langle\text{term}\rangle/_i$ of type $\langle\text{term}\rangle$. An example would be

> the man who was frightened when Bernadette uttered a token of type 'BOO'.

The iterability test says to consider terms of the form of (v) along with those of form

$$(v') \text{ the } y \text{ [REF } (y, /(\text{the } x \text{ [REF } (x, /(\text{term})/_i]])/_i]],$$

as in

> the man who was frightened when Bernadette uttered a token
> of type 'the man who was frightened when Bernadette uttered a
> token of type "BOO"',

as well as (t') and (u').

The claim is that the syntactically relational expression REF here should be understood *not* as standing for a relation (as in an ordinary definite description) but as an anaphoric operator forming indirect descriptions, if and only if the following three conditions are met:[40]

1. *If* (v) is a proper description (that is, in fact picks out one and only one object) and if (v') is a proper description, *then* they corefer (are intersubstitutable).
2. *If* the token by relation to which the individual in (v') is picked out were *not* of the same type as the expression (v), *then* (v) and (v') would not in general corefer.[41]
3. Accepting the identity statement by means of which the coreference of (v) and (v') is asserted is not accidental, in the sense that it is a condition of being taken to understand the expressions involved.

The first condition is required because expressions of type (v') need not always pick out unique objects, even when the expression of type (v) does. There need be no one who is the man Bernadette frightened by uttering a token of type 'the man who was frightened when Bernadette uttered a token of type "BOO"' (either because there is no such man or because there are too many), and similarly for genuinely anaphoric cases. Iterability is a relevant test in the (in general counterfactual) situations where the appropriate individuals exist. The second condition is required in order to rule out cases where the same individual is picked out no matter what token one looks at—the case where one and only one man is frightened, but he is frightened by *whatever* Bernadette says. The type of the intermediate antecedent of an anaphorically dependent expression is obviously essential to its having the reference that it has, so this condition represents a natural constraint. The third condition is required in order to rule out grammatically accidental coreference of (v) and (v'), as might happen in a psychologically homogeneous population with the relation 'is the first object one is reminded of on hearing the expression (term)'.[42] Together these three conditions ensure that any expression REF that satisfies them may appropriately be understood as forming expressions of type (v), which ought to be understood as really having the form of complex anaphoric dependents, that is, as indirect definite descrip-

tions like (t), rather than as ordinary definite descriptions like (u). Frege placed great theoretical weight on the intersubstitutability of the terms 't' and 'the Bedeutung (or referent) of "t"', and this same essential redundancy feature of referring lies at the center of the present account.[43]

6. Other Uses of 'Refers'

The account so far has described the anaphoric category of indirect descriptions as a form of pronoun, has offered a formal test discriminating operators that generate expression types of this category, and has pointed out that 'refers' as it appears in contexts such as

the one Joe referred to as 'that airhead'

can be understood as such a complex pronoun-forming operator. But there are other important uses of 'refers' and its cognates. The most fundamental of these are *Tarskian contexts,* such as

'Rabbits' refers to (denotes) rabbits;

denials of reference, such as

(The expression) 'the present king of France' does not refer (or refers to no one);

mere reference claims, such as

During his talk the speaker referred to Napoleon;

and *referential predications,* such as

The speaker talked about shadowy figures from the intelligence community.

The strategy is to approach these locutions in two stages. First, each such usage is paraphrased into a form in which the 'refers' cognate appears *only* inside an indirect description. Then that description is explained as functioning as a complex pronoun, according to the story already told. The present concern is thus with the paraphrase in terms of indirect descriptions.

The generalization of the token-based account of indirect descriptions required for Tarskian contexts is really a simplification to a special case. For statements about what a term refers to or denotes presuppose that the term type in question is invariant under cotypical intersubstitutions—that is, that all cotypical tokens corefer. So the Tarskian claim

The term 'Leibniz' denotes Leibniz

can be parsed as an identity involving an indirect description:[44]

The one denoted by the term 'Leibniz' is (=) Leibniz.

Given the presupposition of invariance under cotypical intersubstitution of the type corresponding to 'Leibniz' in the original claim, this indirect description is equivalent to

> the one denoted by any token of 'Leibniz',

which may be straightforwardly understood as a complex pronoun, anaphorically dependent on an antecedent that may be any token of the specified type (for instance the one that appears on the other side of the identity sign). So the felt triviality of such reference claims is explained. Of course, reference claims involving expressions that are invariant under cotypical intersubstitution need not be epistemically trivial if different languages or different term types are involved, as in

> In our world (the expression) 'the first postmaster general' refers to Benjamin Franklin (or: the inventor of bifocals),

which can be understood as using anaphoric relations to claim

> In our world, the one referred to as 'the first postmaster general' is (=) Benjamin Franklin,

which is not a trivial assertion.

If Tarskian truth conditions are set up using denotation claims for semantic categories besides terms, these can be accommodated as well by this scheme.

> 'Red' refers to (denotes) red things

is to be read as

> The ones referred to as (denoted by) 'red' are red things.

Also,

> 'Magnetic' applies to (has in its extension) magnetic things

is to be read as

> The ones 'magnetic' applies to are magnetic things.

It was pointed out above that the prosentential account of '. . . is true' underwrites the Tarski biconditionals—that is, claims of the form

> 'Snow is white' is true if and only if snow is white.

That all the biconditionals of this form be generated is, of course, the primary criterion of adequacy Tarski imposes on candidate conceptions of truth (or better, from the present point of view, conceptions of what is expressed by '. . . is true'). In his own formal theory of truth, he shows how these biconditionals can be recursively generated (in a first-order formal language) on the

basis of a finite number of primitive denotation relations expressed by statements of the form

> 'Benjamin Franklin' denotes Benjamin Franklin

and

> 'Red' denotes red things.

The account of 'refers' and 'denotes' as proform-forming operators underwrites all of these basis clauses of the Tarskian recursion in just the same way that the account of '. . . is true' underwrites the biconditionals that result. Furthermore, the basis clauses of the recursion do not have to be simply stipulated, one by one, in the way that Field finds theoretically objectionable.[45] Rather, they are generated in a principled way from an underlying anaphoric account of what reference claims express. That account deals gracefully with extensions of the language being considered and generalizes to deal with arbitrary languages. The anaphoric approach accordingly makes intelligible what McDowell has called "modest" Tarskian theories of truth. All that is presupposed is that one understands the various subsentential parts of speech involved (which is the subject of Chapter 6) and what anaphora is (which is the subject of Chapter 7). The present discussion proceeds subject to these promissory notes.

This account of Tarskian contexts in which 'refers' and its cognates appear respects the different modal status of

> The term 'Leibniz' denotes Leibniz,

which is only contingently true, and

> Leibniz is Leibniz,

which is necessarily true. For the possibility that the first claim is not true can be understood in terms of its paraphrase as the existence of a possible world w such that

> The one referred to as 'Leibniz' in w is not Leibniz,

that is, is not the one we refer to in our own world as 'Leibniz'. The explicit relativization of the indirect description to a possible world simply specifies which world its antecedent tokens are to be found in. The candidate antecedents of

> the one referred to as 'Leibniz' in w

are tokenings of the type 'Leibniz' that are uttered in w. The anaphoric approach accordingly has room for what has been thought of as the contingency of word-world semantic relations, although it is not based on such relations.

Coreference claims represent a simple variation on Tarskian contexts and

can be interpreted in much the same way. To say that the expressions $\langle type_1 \rangle$ and $\langle type_2 \rangle$ corefer is just to assert an identity between the corresponding indirect definite descriptions, that is, to say that the one referred to by (tokens of) $\langle type_1 \rangle$ *is* the one referred to by (tokens of) $\langle type_2 \rangle$. Asserting such an identity is licensing as (assertional) commitment-preserving the intersubstitution of expressions of those types. (The expressive role characteristic of identity locutions, making explicit substitution-inferential commitments, is discussed in Chapter 6.) In the case of cotypically nonintersubstitutable expressions such as demonstratives and pronouns, asserting an identity authorizes substitution of anaphoric dependents of the token on the left for anaphoric dependents of the token on the right of the identity, and vice versa. (The way in which anaphoric chains of such tokenings can play the same role in substitution inferences that classes of cotypical tokenings do for expressions that are suitably invariant is discussed in Chapter 7.)

It may be worth noticing that if attention is restricted to term types that are invariant under cotypical intersubstitution, the iteration test introduced above can be simplified correspondingly, by omission of token specifications. The necessary condition for operators PF to form complex pronouns that are invariant under cotypical intersubstitution is then:

$$PF(\langle PF(\langle type \rangle) \rangle) = PF(\langle type \rangle),$$

and the corresponding condition suggested as sufficient for REF to be an indirect description-forming operator is:

$$\text{the } y[\text{REF } (y, \langle \text{the } x[\text{REF } (x, \langle type \rangle)] \rangle)] = \text{the } x[\text{REF } (x, \langle type \rangle)].$$

Simple negations of statements of reference, as in

> (The expression) 'the shortest man in the room' does not refer to John,

raise no new issues, for the underlying identity that is being negated has already been explained. But claims that an expression does not refer to *anything* deserve special mention. The obvious way of extending to these cases the previous strategy of paraphrasing what look like assertions of reference relations as identities involving indirect descriptions is to quantify into the identity and read the result as a negative existential statement. That is, statements of the form

> $\langle type \rangle$ does not refer

are to be read as

> The one referred to as $\langle type \rangle$ does not exist,

where this last is to be understood in the same way as ordinary negative existentials, such as

> The present king of France does not exist.

(Just how existential and negative existential statements ought to be understood is discussed officially in Chapter 7.) That the indirect definite descriptions involved in denials of referentiality are anaphoric dependents makes no difference to the reading of the negative existentials, any more than it causes difficulty in understanding remarks like

> I would be comforted by the benevolence of a supreme being, except that such a being does not exist,

in which 'such a being' is an anaphoric dependent.
  Statements like

> During his talk the speaker referred to Napoleon

say that reference has taken place but give no information about what the referring tokens or types were. Such remarks may be understood as asserting that there is *some* term tokening t in the speaker's discourse such that the item referred to by t is Napoleon. Statements like

> The speaker talked about shadowy figures from the intelligence community

are similar, except that a predication rather than an identity is what is asserted of the items referred to or talked about. This sentence says that there were tokenings t, t' (and perhaps more) such that the items referred to as (or talked about by the use of) t and t' are (have the property of being) shadowy figures from the intelligence community. These are predications involving pronouns, intrinsically no more mysterious than sentences like

> They are confused.

Common nouns can be formed from indirect descriptions just as they can be from ordinary direct descriptions, and the present account extends straightforwardly to these expressions, as in

> All the animals the speaker mentioned tonight were quadrupeds.

This example indicates as well how *generalizations* about reference are to be approached anaphorically. Endorsement of this claim commits one to all the substitution instances of the form

> If t is an animal the speaker mentioned tonight, then t is a quadruped.

The antecedent of each such conditional is a referential predication, equivalent to

> There is a term token /s/ such that the speaker uttered /s/ tonight and the item referred to by /s/ is an animal and the one referred to by /s/ is (=) t,

a kind of claim that has already been given an interpretation. As long as one knows what to make of each of the substitution instances to which a universally quantified claim undertakes commitment, one knows what to make of the universal generalization itself.

A full discussion of such cases requires accounts of anaphora and quantification that are not yet on the table. (The substitutional significance of such quantificational claims is discussed more fully in Chapter 6, and the account of anaphora is to be found in Chapter 7.) The complications arise in part from the recognition that in the general case the term substituend t, which is repeated in the specification above of the form of each sentential substitution instance of the quantification, would need to be replaced by two (not necessarily cotypical) term tokens, one of which is anaphorically dependent on the other. The present point is that although an account is not yet being offered of quantification in general, it is clear from the example that no new difficulties are added by the presence of anaphorically indirect descriptions in the quantificational substitution instances, so that generalizations about reference can be understood if any sort of generalization can.

Consideration of generalizations about what is referred to by various expressions makes salient another issue, which can be dealt with only in passing here, namely the susceptibility of an anaphoric account of reference claims to the formulation of semantic paradoxes. In the presumably analogous case of truth, a naive substitutional understanding of quantification into truth claims commits one to interpreting paradoxical sentences such as the Liar. Of course generalization is not the only way in which such paradoxical expressions can arise, nor is the possibility of semantic paradox restricted to the category of sentences. It is possible to use 'refer' to formulate empirically paradoxical term tokens, such as

> (w) the square root of 2 that is the result of multiplying –1 by the one referred to by the term token marked 'w',

where 'square root of 2' is a sortal comprising the positive and negative square roots and 'one' is understood as a prosortal anaphorically dependent upon it. Interpreting such tokens as anaphorically indirect descriptions focuses attention on grounding conditions for anaphoric inheritance—a large and important topic. In "Inheritors and Paradox," Dorothy Grover elaborates an anaphoric approach to semantic paradoxes for the closely analogous anaphoric treatment of '. . . is true' discussed above.[46] Grover finds that the natural condition on anaphoric grounding yields an interpretation coinciding

in general with the sentences that Kripke assigns a semantic value to at the minimal fixed point, the interpretation he takes to provide the most natural model for the intuitive concept of truth. Her remarks can be applied to the present construction by means of the crucial analogy between the pronominal account of reference and the prosentential account of truth.

V. THE FUNCTION OF TRADITIONAL SEMANTIC VOCABULARY IS EXPRESSIVE, NOT EXPLANATORY

1. Summary

The conclusions of the discussion of truth in Sections II and III may be summed up as follows. The pragmatists' approach to truth introduces a bold phenomenalist strategy—to take as immediate explanatory target the practical proprieties of *taking*-true and to understand the concept of truth as consisting in the use that is made of a class of expressions, rather than starting with a *property* of truth and then seeing what it is to express a concept used to attribute that property. Their implementation of this strategy is flawed in its exclusive attention to taking-true as a variety of *force* or pragmatic significance—as a doing, specifically an asserting of something. For 'true' is used in other contexts, for instance, embedded in the antecedent of a conditional; the semantic content that it expresses is accordingly not exhausted by its freestanding assertional uses. Content-redundancy theories can incorporate the insights of these force-redundancy accounts, and in their most sophisticated (anaphoric) form they account for the wider variety of uses of 'true'. Indeed, starting with an analogous pronominal account of 'refers' and 'denotes', it is possible to generate Tarski-wise the truth equivalences that jointly express the content redundancy of '. . . is true'. The strategy of the classical pragmatists has been vindicated at least this far: It is possible to account for truth talk without invoking a property of truth that such talk must be understood as answering to.

The project pursued in Section IV was to make it plausible that the use of 'refers' and cognate locutions in natural languages can be understood by first paraphrasing contexts in which they occur into a form in which they appear only inside indirect descriptions and then understanding their role in those paraphrases as operators taking token (or type) specifications and forming from them lexically complex pronouns invariant under cotypical intersubstitution, whose anaphoric antecedents are the specified tokens (or tokens of the specified type). This account of 'refers' as a pro*noun*-forming operator is evidently parallel to the account of '. . . is true' as a pro*sentence*-forming operator. Each theory explains the use of a bit of traditional *semantic* vocabulary in terms of the formation of *anaphoric* proforms. Indirect descriptions formed from 'refer' both *mention* a term expression (in picking our anaphoric antecedents) and *use* that expression. The effect of applying an indirect

description-forming operator to a mentioned term is that of turning the mentioned occurrence into a used occurrence. Thought of in this way, 'refers' is an *anaphoric disquotation operator* in the same sense that 'true' is.[47]

It should be acknowledged that this anaphoric account of the use of traditional semantic vocabulary does not underwrite all of the idioms that have pressed those expressions into service. There is one sort of truth and reference talk that is not recoverable on a prosentential and pronominal rendering of 'true' and 'refers'. Talk in which the substantive 'truth' appears in a way not easily eliminable in favor of 'true' will receive no construal by such theories. Yet philosophers do say such things as "Truth is one, but beliefs are many" and "Truth is a property definable in the language of some eventual physics," which are outside the scope of the account of 'true' offered here.[48] Similarly, although accounts are offered of what someone referred to by an utterance and of what the reference of the utterance was, nothing is said about the *relation* of reference. The anaphoric approach does not say how to understand sentences such as "Reference is a physical, causal relation." The reason is clear enough. On the anaphoric account, although '. . . is true' has the surface syntactic form of a predicate, and '. . . refers to . . .' the surface syntactic form of a relational locution, the grammatical and semantic roles these expressions play are not those of predicative and relational locutions. Their grammar is quite different; they are operators forming anaphoric dependents—namely prosentences and anaphorically indirect descriptions.

Philosophers have misconstrued ordinary talk using 'true' and 'refers' on the basis of a mistaken grammatical analogy to predicates and relational expressions. On the basis of this mistaken analogy (though the mistake is understandable, given the surface forms), they have hypostatized a property of truth and a relation of reference as the semantic correlates of the apparently predicative and relational expressions. Competing theories of the nature of this odd semantic predicate and relation have then been forthcoming. Such a search is of a piece with the search for the objects corresponding to each expression that plays the surface syntactic role of a singular term—for instance quantificational expressions such as 'someone' and 'everyone', or the 'it' in "It is raining." A more careful look at the (substitution-inferential) use of these expressions shows that the initial analogy to singular terms is misleading; a more careful look at the (anaphoric) use of 'true' and 'refers' similarly shows that the initially tempting assimilation of them to expressions of properties and relations is misleading.

One who endorses the anaphoric account of what is expressed by 'true' and 'refers' must accordingly eschew the reifying move to a truth property and a reference relation. A line is accordingly implicitly drawn by this approach between ordinary truth and reference talk and various specifically philosophical extensions of it based on theoretical conclusions that have been drawn from a mistaken understanding of what such talk expresses.

Ordinary remarks about what is true and what is false and about what some expression refers to are perfectly in order as they stand; the anaphoric account explains how they should be understood. But truth and reference are philosophers' fictions, generated by grammatical misunderstandings. It is no defect in the anaphoric account not to generate readings of the fundamentally confused remarks that result. Taking a claim to be true must be understood in the first instance as adopting a normative attitude—that is, endorsing the claim and so acknowledging a commitment. This normative attitude is presupposed by the possibility of ascribing an objective property and is not to be explained in terms of it.

Chapter 8 (Section VI) discusses in deontic scorekeeping terms what it is to ascribe objective properties, or more generally to make objectively representational claims—claims subject to objective assessments of correctness, depending on how things are with what is represented by them, regardless of the attitudes or endorsements of anyone. The expressive power of 'true' ensures that where an objective property is ascribed to something, the resulting claim can correctly be said to be objectively true or false. Properly understood, however, no property of truth (objective or otherwise) is being invoked by such a remark. One who asserts "The claim that all integers are the sum of at most nineteen fourth powers is objectively true" ascribes an objective property to the integers, but not to the claim. Navigating in the idiom that distinguishes these (endorsing the first specification of the ascription and not the second) is a somewhat delicate matter. The discussion below of the concept of a *fact* may be helpful in acquiring this skill.

Parallel remarks may be made about a supposed word-world relation of reference. The present account distinguishes sharply between expressions and their referents—where that latter expression is understood as shorthand for 'what is referred to by those expressions'. For very different uses are associated with the expressions 'the expression "Leibniz"' and 'the referent of the expression "Leibniz"'. In particular, although doxastic commitments are preserved by intersubstitution of the latter with the singular term 'Leibniz', they are not preserved by intersubstitution of the former with 'Leibniz'. That is, Leibniz is (=) the one referred to by the expression 'Leibniz', but Leibniz is not the expression 'Leibniz'. These remarks will not officially be intelligible until the substitution-inferential role of singular terms has been explained (in Chapter 6), and the relation between anaphoric dependents and their antecedents has been explained in terms of inheritance of substitution-inferential role (in Chapter 7). For present purposes it suffices to say that treating tokens of the type ⟨the referent of the expression 'Leibniz'⟩ as anaphorically dependent on tokens of the type ⟨Leibniz⟩ commits one to the propriety of the intersubstitution that would be made explicit by endorsement of the assertible identity

Leibniz is (=) the referent of the expression 'Leibniz'

and does not commit one to the propriety of the intersubstitution that would
be made explicit by endorsement of the assertible identity

> Leibniz is (=) the expression 'Leibniz'.

So although anaphora is an *intra*linguistic (or word-word) relation, adopt-
ing an anaphoric account of 'refers' as a proform-forming operator does not
entail conflating linguistic items with extralinguistic items. No doubt, as
with any other two items in the causal order, there are many relations that
can correctly be said to obtain between a term tokening and what it refers
to. But the present considerations show that talk about referring and referents
provides no reason whatever to conclude that some one of these could be
singled out as *the* reference *relation*—that unique semantically significant
word-world relation in virtue of which the nonexpression is the referent of
the expression. Various word-world relations play important explanatory
roles in theoretical semantic projects, but to think of any one of these as what
is referred to as "the reference relation" is to be bewitched by surface syn-
tactic form.

In order to see what is and is not being claimed for this analysis of 'true'
and 'refers' and to see the significance of the replacement of an account of a
truth *property* and a reference *relation* by anaphoric accounts of what those
traditional semantic locutions express, it is helpful to redescribe and clarify
the explanatory role that the anaphoric analysis is supposed to play. Suppose
that from some language-in-use such as English one extrudes all the sen-
tences that contain specifically *semantic* vocabulary, such as 'true' and 're-
fers' and other words used in the same ways.[49] The anaphoric analysis
presented in this chapter permits the extension of an account of the use of
this *non*semantic fragment of the language to an account of the use of the
whole language. Put otherwise, it explains how to add the expressive power
provided by traditional semantic vocabulary to a set of linguistic practices
that does not have such locutions. In particular, once the account (in Chap-
ters 3 and 4) of discursive practices, the pragmatic significances they insti-
tute, and the semantic contents they confer has been extended (in Chapter
6) to incorporate the substitution-inferential commitments governing the use
of subsentential expressions, and (in Chapter 7) to the anaphoric inheritance
of such commitments, the considerations advanced in this chapter then
suffice to incorporate 'true' and 'refers' into the deontic scorekeeping model
of linguistic practice.

2. Semantic Deflationism

This is a *deflationary* account of the role of traditional semantic
vocabulary, paradigmatically 'true' and 'refers'. Three deflationary conse-
quences of the anaphoric approach to such vocabulary are particularly note-
worthy. First, it is denied that there is a *property* of truth or a *relation* of

reference. Second, it is denied that claims expressed using traditional semantic vocabulary make it possible for us to state specifically *semantic facts,* in the way that claims expressed using the vocabulary of physics, say, make it possible for us to state specifically *physical facts.* Third, it is denied that the notion of *truth conditions* can be appealed to in explaining (as opposed to expressing) the sort of *propositional contents* expressed by declarative sentences—and similarly that the notion of *association with a referent* can be appealed to in explaining the sort of semantic contribution the occurrence of a singular term makes to the contents of sentences in which it appears.

It will help in clarifying the status and significance of these denials to consider them in connection with a significant structural difficulty that Boghossian has diagnosed as afflicting many forms of deflationism.[50] He points out that corresponding to each of the three sorts of claims just rehearsed there is a danger of deflationism undercutting itself and lapsing into incoherence. The general worry is that the force of deflationist claims depends on the contrast between predicates (such as '. . . has a mass of more than ten grams') that do, and those (such as '. . . is true') that do not, correspond to properties—and declarative sentences (such as "Snow is white") that do, and those (such as "It is true that snow is white") that do not, state facts. But such contrasts seem to presuppose a robust correspondence theory of the contents of some predicates and claims—at least those the semantic deflationist finds unproblematic, paradigmatically those of natural science. Deflationary approaches to semantics seem to be saying: physical predicates correspond to physical properties and relations, but semantic predicates do not correspond to semantic properties and relations; physical claims have truth conditions and if true correspond to physical facts, but semantic claims do not have truth conditions and so cannot correspond to semantic facts. Yet at the same time, the deflationists want to deny that content can be explained in terms of truth conditions and correspondence to facts, properties, and objects.

So, it is claimed, deflationary approaches are conceptually unstable. Consistently following out the rejection of robust correspondence theories of content requires treating using an expression as a predicate as all there is to expressing a property, and using a declarative sentence to make a true claim to be all there is to stating a fact. So on a deflationary construal, one is forbidden to deny that the predicate '. . . is true' denotes a property, or that the claim "It is true that snow is white" states a semantic fact. Yet it is the essence of deflationism to deny these claims. So the very intuitions that deflationism seeks to develop and defend surreptitiously presuppose exactly the sort of robust truth-conditional correspondence theory of content they are concerned to reject.

These considerations present a serious challenge to many ways of pursuing the program of semantic deflationism, but the approach that has been presented here is immune to them. The first argument depends on treating

'. . . is true' as a predicate. If it is, then since that expression is used to make claims and state facts, it must, on deflationary accounts, be taken to express a property. But the essence of the anaphoric versions of semantic deflationism is precisely to take issue with this grammatical presupposition. According to the account endorsed here, '. . . is true' expresses a prosentence-forming operator. Its syntax and grammar are quite distinct from those of predicates, to which it bears only the sort of surface similarity that quantificational expressions bear to genuine singular terms. In particular, proform-forming operators such as those formed using 'true' and 'refers' are syntactically distinguished from superficially similar predicates and relational expressions by the iteration condition. The part of speech '. . . is true' is assimilated to by these theories does not have a directly denotational semantics but inherits its significance anaphorically, by an entirely distinct mechanism.

So when it is claimed here that '. . . is true' does not express a property, this means that it is not even of the right grammatical form to do so—anymore than 'no one' is of the right form to pick out an individual, although there are some features of its use that could mislead one on this point. Furthermore, this claim is not made ad hoc, to avoid the sort of theoretical circularity Boghossian points out, but is motivated by ground-level considerations having to do with the use of 'true' and 'refers' that, it is claimed, cannot otherwise adequately be represented. Thus from this point of view, the argument to the effect that: "the denial that a given predicate refers to or expresses a property only makes sense on a robust construal of predicate reference . . . But if this is correct, the denial . . . that the truth predicate refers to a property must itself be understood as framed in terms of a robust notion of reference"[51] depends upon the ultimately incorrect presupposition that truth is properly rendered as a predicate. Given this, the second claim does not follow from the first. It can be granted that denying of a predicate that it expresses a property presupposes a robust conception, without being committed thereby to anything about how '. . . is true' ought to be understood.

3. Facts Are True Claims

What, then, about the sentential level, at which claims are made and facts stated? Is the semantic "nonfactualism" (Boghossian's term) of the deflationist incoherent? In assessing this claim, it is important to distinguish two different ways in which one might adopt a nonfactualist attitude toward what is claimed by "It is true that snow is white." According to the usage endorsed here, facts are just true claims.[52] That is, phenomenalistically, to call something a fact is just to take it to be true. 'Claims' here has the semantic sense of what is claim*ed*, rather than the pragmatic sense of the claim*ing* of it—a matter of content, not of force or deontic attitude. Thus to

say that facts are just true claims does not commit one to treating the facts as somehow dependent on our claimings; it does not, for instance, have the consequence that had there never been any claimers, there would have been no facts. (There are no possible situations in which there would have been no facts. A situation or set of circumstances just is one sort of set of facts.)

This notion of facts as true claims is meant to contrast with a view of them as what *makes* claims true—at least where that latter formula is conceived as potentially of *explanatory* use, rather than as providing an *expressive* equivalent. To say that it is the *fact* that p that *makes* it *true* that p (or that it is true that p *because* of the fact that p) is to provide an *explanation* only in the misleading sense in which what *makes* it the case that the Greeks defeated the Persians at Plataea is that the Persians were defeated by the Greeks at Plataea—the sense in which the Greeks defeated the Persians at Plataea *because* the Persians were defeated by the Greeks at Plataea.

In a certain sense, facts are what make claim*ings* true. But claim*ings* are true at all only in a derivative sense: We say "What you say is true," not "Your saying of it is true." Your saying can be speaking tru*ly*, but that is just saying something (making a claim) that is true. Talk of facts as what *makes* claims true is confused if it is thought of as relating two distinct things—a true claim and the fact in virtue of which it is true—in such a way that the former might be *explained* by appeal to the latter. Rather, "The claim that p is true" and "It is a fact that p" are two equivalent ways of saying the same thing—expressing the *same* content, and so (if the claim they both express is true) stating the *same* fact.

Truth claims can be true, so some of them state facts. (For the reasons rehearsed above, this does not entail that there is a property of truth.) So far this does not sound "nonfactualist." Yet it is denied that there are any specifically *semantic* facts. "It is true that snow is white" expresses just the same fact that "Snow is white" expresses. The former uses anaphoric mechanisms to do so; its expression of that fact can thus involve presuppositions that are not involved in the latter expression (as becomes obvious if one considers other forms of truth claim), but these need not be considered part of the fact that is stated, any more than the differences involved in interpreting "Snow is white" and "La neige est blanche" or, under the right circumstances, "John is confused" and "He is confused" mean that these do not express the same fact.

Physical claims do state specifically physical facts because they essentially employ specifically physical predicates and so invoke specifically physical properties and relations. Semantic claims do not state specifically semantic facts because their use of specifically semantic vocabulary does not invoke specifically semantic properties and relations. This sounds like it ought to be called "nonfactualism" about truth talk. The distinction that must be kept in mind is that between claiming that "Snow is white is true"

states a fact (which deflationists had better not deny, for the reasons Boghossian points out) and claiming that it states a special *kind* of fact, namely a semantic fact. The 'deflating' part of deflationism can consist in its denial of this latter claim. Mastering the vocabulary of physics (or for that matter, etiquette) gives us expressive access to a range of facts we cannot otherwise express. Mastering semantic vocabulary just gives us a new way (useful for other reasons, having to do with communication and generalization) of getting at a range of nonsemantic facts we already had access to. This is just the point of redundancy deflationism. "It is true that snow is white" is a semantic *expression* of a nonsemantic fact.

In this connection it is worth making a third point about Boghossian's arguments against semantic deflationism. He thinks that the moral we should draw is that "we really cannot make sense of the suggestion that our thoughts and utterances do not possess robust truth conditions."[53] If contents must be explained as truth conditions, then an argument that such truth conditions cannot be understood in a deflationary way amounts to an argument that they, and so contents, must be construed robustly. But one could equally well conclude that one ought not to explain propositional contentfulness in terms of truth conditions. As Dummett argued long ago, anyone who holds to a deflationary theory of truth is precluded from explaining propositional contents in terms of truth conditions. For redundancy theories of 'true' presuppose the contentfulness of the nonsemantic sentences on which semantic claims are redundant, in order to explain how 'true' ought to be used. It would be circular to presuppose such contents in an account of truth, if the contents are themselves to be construed in terms of what then must be an antecedently intelligible notion of what it is to be true. If one can make robust antecedent sense of truth, then one can appeal to it to explain contents without circularity (a big 'if'). If, however, one is a deflationist about truth, then it is necessary to look elsewhere for the basic concepts one appeals to in explaining contentfulness. That is the line that has been pursued in this work.

One can say of anything that has a propositional content that it has truth conditions. According to the relaxed deflationary view, this characterization is just a harmless compliment paid to things whose contents can be expressed in declarative sentences or by the corresponding 'that' clauses (picked out in turn by their special role in the practice of making assertions). But it is one thing to say that whatever is contentful will, in consequence, have truth conditions. It is quite another to think that one could use the possession of truth conditions as part of an *explanation* of propositional contentfulness. As Dummett recognized, this latter strategy is forbidden to deflationists, on grounds of circularity. Deflationists ought to acknowledge the general possibility of *expressing* semantic content truth-conditionally, while denying the possibility of *explaining* semantic content truth-conditionally.

4. Correspondence, Constraint, and Representation

From the perspective provided by this way of talking, then, correspondence theories of truth are unsatisfactory because they are unenlightening, rather than because they are false. True claims do correspond to facts, and understanding claims does require grasp of what the facts must be for those claims to be true. For when the 'ing'/'ed' ambiguity is resolved, these theses take one of two forms. If 'claim' is understood as what is claimed, true claimable contents just *are* facts; the relation of 'correspondence' is just that of identity. For that reason, grasp of such contents can be identified with grasp of what the facts must be for them to be true. But the basic question is what one must be able to *do* in practice in order to count as grasping or understanding an assertible (hence propositional) content. Paraphrases in terms of corresponding facts serving as truth conditions provide no independent explanatory grip on the issue—only an alternate way to express it.

If 'claim' is understood as the act of claiming, rather than the content claimed, however, true claim(ing)s can be said to correspond to facts in a stronger sense. They *express* those facts; they are the acts of making explicit, in virtue of whose significance as acknowledgments of inferentially articulated commitments anything at all can be understood as a claimable—and hence, if true, as a fact. In exactly the same sense, false claims express their claimable contents. In either case, scorekeeping mastery of the significance of claimings depends on one's grasp of the claimed contents. This is, trivially, grasp of what (claimable contents) must be true, what the facts must be, if the claiming is to be a true-claiming—a claiming of a true claimable content.

The important thing to get clear about is what it is for an act of claiming to express a claimable content, that is, the activity of making something explicit. Once the expressive role of 'true' and 'fact' is properly understood, it becomes apparent that their use presupposes a notion of propositional content (hence of propositionally contentful acts and states); so what such traditional semantic vocabulary expresses is not in principle available to *explain* the nature of propositional contentfulness. Their parasitic expressive role precludes their playing a fundamental semantic explanatory role.[54] By contrast (as Chapter 3 shows) it is possible to explain the practical significance of acts of claiming, and so to approach the propositional contents they express, without appealing to notions of truth conditions or fact. The use of expressions such as 'true' and 'fact' can then (as this chapter shows) be explained in terms of these same social practices of giving and asking for reasons.

On neither of these construals of claims is there room for a robust correspondence between facts and claims. What the facts are does not depend on what claimings we actually effect. But the worry may remain that a semantic idiom that identifies facts with true claims (via the identification of taking to be a fact with taking to be true, that is, with acknowledging a doxastic

commitment) must inevitably 'lose the world'—trading its solidity for a froth of words. A threatening idealism of linguistic practice seems to be implicit in such an identification.

But this is a misplaced concern. What must not be lost is an appreciation of the way in which our discursive practice is empirically and practically *constrained.* It is not up to us which claims are true (that is, what the facts are). It is in a sense up to us which noises and marks express which claims, and hence, in a more attenuated sense, which express true claims. But empirical and practical constraint on our arbitrary whim is a pervasive feature of our discursive practice. Words form a distinct and largely independent realm within the world—in the sense not only that the nonlinguistic facts could be largely what they are even if the specifically linguistic facts (thought of as a class of facts about words) were quite different, but also in the sense that the words—as noises, marks, and so on—could be largely what they are, even if the nonlinguistic facts were quite different. But discursive *practices* as here conceived do not stand apart from the rest of the world in this way. The nonlinguistic facts could be largely what they are, even if our discursive practices were quite different (or absent entirely), for what claims are true does not depend on anyone's claiming of them. But our discursive practices could not be what they are if the nonlinguistic facts were different.

For those practices are not things, like words conceived as marks and noises, that are specifiable independently of the objects they deal with and the facts they make it possible to express. Discursive practices essentially involve to-ing and fro-ing with environing objects in perception and action. The conceptual proprieties implicit in those practices incorporate both empirical and practical dimensions. All our concepts are what they are in part because of their inferential links to others that have noninferential circumstances or consequences of application—concepts, that is, whose proper use is not specifiable apart from consideration of the facts and objects that responsively bring about or are brought about by their application. The normative structure of authority and responsibility exhibited by assessments and attributions of reliability in perception and action is causally conditioned.

This sort of causal contribution to the norms implicit in discursive practice means that even though it is the practices of a linguistic community that make their words express the concepts they do, the members of the community may be understood to have undertaken commitments by using those words that outrun their capacity to recognize those commitments. Earthlings and twin-earthlings may apply the same phonetic and orthographic sign design 'water' to samples of clear, tasteless, odorless, thirst-quenching liquids on their respective planets (and may make corresponding inferential moves with them to and from other sign-designs) and still be understood (by us, who are describing the case) to be applying different *concepts* thereby, if the noninferential circumstances of appropriate application of their concept involves the presence of XYZ rather than H_2O. This can be so even if neither

earthlings nor twin-earthlings can be trained reliably to discriminate XYZ from H_2O perceptually.

Practitioners are not in general omniscient about the commitments implicit in their own concepts. For the interpreter who is making sense of their practices—and who *is* able (not necessarily perceptually, but conceptually) to distinguish H_2O and XYZ—can understand transported earthlings as *mis*taking for water the XYZ they look at, as *in*appropriately applying the concept they express with their word 'water' to that unearthly stuff.[55] As with assessments of reliability, truth, and knowledge generally, the 'externalist' element in attributions of commitments implicit in conceptual contents reflects the social difference in perspective between the scorekeeper and those whose normative statuses are at issue. One can (according to an interpreter or scorekeeper) have bound oneself by one's practice, in part because of the things one was actually dealing with, in such a way that using a particular word is correct in one circumstance and incorrect in another—even when the individual so bound cannot tell the situations apart.

Discursive practices incorporate actual things. They are solid—as one might say, *corporeal:* they involve actual bodies, including both our own and the others (animate and inanimate) we have practical and empirical dealings with. They must not be thought of as hollow, waiting to be filled up by things; they are not thin and abstract, but as concrete as the practice of driving nails with a hammer. (They are our means of access to what is abstract—among other things—not its product.) According to such a construal of practices, it is wrong to contrast discursive practice with a world of facts and things *outside* it, modeled on the contrast between words and the things they refer to. It is wrong to think of facts and the objects they involve as constraining linguistic practice from the outside—not because they do not constrain it but because of the mistaken picture of facts and objects as outside it. What determinate practices a community has depends on what the facts are and on what objects they are actually practically involved with, to begin with, through perception and action. The way the world is, constrains proprieties of inferential, doxastic, and practical commitment in a straightforward way from *within* those practices.[56] So if I perceive a liquid as tasting sour, infer that it is an acid, infer further that it will therefore turn litmus paper red, and, intending to match a red pigment sample, accordingly dip litmus paper in the liquid, I may nonetheless subsequently acquire perceptually a commitment to the result being a blue, rather than a red, piece of paper, and hence an acknowledgment of my practical failure. In this way I can find myself with incompatible commitments (which need to be sorted out if I am to remain entitled to any of my commitments in the vicinity). The possibility of incompatible commitments arising from the cycle of perception, inference, action, and perception reflects the way the normative structure of perception and action incorporates elements of the causal order.

As a result, empirical and practical constraints get built into what commitments (including inferential commitments) one can sustain entitlement to.

Thus a demotion of semantic categories of correspondence relative to those of expression does not involve 'loss of the world' in the sense that our discursive practice is then conceived as unconstrained by how things actually are. It does involve giving up the picture of how things are as contrasting with what we can say and think. Facts are (the contents of) true claims and thoughts. As Wittgenstein says: "When we say, and *mean*, that such-and-such is the case, we—and our meaning—do not stop anywhere short of the fact; but we mean: this-is-so."[57] What is lost is only the bifurcation that makes knowledge seem to require the bridging of a *gap* that opens up between sayable and thinkable contents—thought of as existing self-contained on their side of the epistemic crevasse—and the worldly facts, existing on their side.[58] What the picture of facts as true claims loses is only "the little rift within the lute, / that by and by will make the music mute, / and ever widening slowly silence all."[59]

The world is everything that is the case, a constellation of facts. But as the author of these words hastened to point out, those facts are structured and interconnected by the objects they are facts about; they are articulated by the properties and relations the obtaining of which is what we state when we state a fact (claim when we make a claim). To make a claim is to say that things are thus and so—that is, to talk about objects, and to say how they are propertied and related. Propositional contents (and hence facts) cannot be properly understood without understanding their representational dimension—what it means for them to be *about* objects and their properties and relations. The next three chapters accordingly show how to move from an account of the *expression* of facts to an account of the *representation* of objects and properties.

6

Substitution:
What Are Singular Terms,
and Why Are There Any?

I start out from judgments and their contents, and not from concepts . . . I only allow the formation of concepts to proceed from judgments. If, that is, you imagine the 2 in the content of judgment $2^4 = 16$ to be replaceable by something else, by (-2) or by 3 say, which may be indicated by putting an x in place of the 2: $x^4 = 16$, the content of possible judgment is thus split into a constant and a variable part. The former, regarded in its own right but holding a place open for the latter, gives the concept '4th root of 16' or 'the individual 2 falls under the concept "4th root of 16" or "belongs to the class of 4th roots of 16"'. But we may also just as well say '4 is a logarithm of 16 to the base 2'. Here 4 is being treated as replaceable and so we get the concept 'logarithm of 16 to the base 2': $2^x = 16$ . . .

And so, instead of putting a judgment together out of an individual as subject and an already previously formed concept as predicate, we do the opposite and arrive at a concept by splitting up the content of a possible judgment.

FREGE, "Boole's Logical Calculus and the *Begriffsschrift*"

I. MULTIVALUED LOGIC AND MATERIAL INFERENCE

1. *Three Challenges for Inferential Approaches to Semantics*

The theoretical structure being explored here is animated by commitments both to a *deontic pragmatics* and to an *inferential semantics*. The first means that the states to be investigated, the original bearers of intentional contents, are to be understood normatively—more particularly as species of *commitments* and *entitlements*. The phenomenalist account of deontic statuses such as commitment, in terms of scorekeeping with the socially complementary deontic attitudes of *attributing* and *undertaking*, is offered as a way to begin filling in such an approach to pragmatics. The second theoretical commitment means that the *contents* that determine, in context, the deontic *significance* of adopting or altering a deontic status, or of performing a contentful act, are to be understood as broadly *inferential* roles. The content must, in context, fix the circumstances in which one would be entitled to adopt or undertake a commitment with that content and must fix the appropriate consequences of undertaking such a commitment. Employing an expression with that content then involves endorsing

the inferential commitment from those circumstances of entitlement to those consequences of commitment. The description of the game of asserting, of the inferentially articulated practices that confer assertible, that is, propositional, contents on states, acts, and utterances in virtue of their roles in that game, and the account of logical vocabulary as distinguished by its expressive task of making explicit as assertible contents precisely the inferential commitments that determine those roles, are offered as a way to begin filling in such an approach to semantics. This chapter continues the inquiry into inferential notions of semantic content.

There are three topics that would seem to pose special explanatory difficulties for attempts to understand semantic content in terms of proprieties of inference. First, the functional involvements that could plausibly be taken to be responsible for the conferral of such contents relate conceptually contentful deontic states not only to each other but also to the nondiscursive environment. Perception and action, as entries to and exits from the discursive realm, are governed by practical proprieties every bit as important as, and irreducible to, those governing purely inferential moves within that realm. So significant have the entries and exits seemed that each has been taken, by some empiricists and by some pragmatists in turn, to be the sole source of content for intentional states—to the exclusion not only of inferential articulation but of each other. How can a broadly *inferential* approach incorporate the aspects of semantic content conferred by the *non*inferential aspects of such entries and exits?

Second, the notion of content as inferential role seems naturally adapted to account only for *propositional* content, for it is only commitments with contents of this category that can play the role of premise and conclusion in inferences. But the sentences that express propositions typically have significant parts that are not sentences, which do not express propositions, and so which cannot serve as inferential premises and conclusions. Yet these subsentential expressions certainly ought to be said to be contentful, in virtue of what Dummett calls the "contribution" they make to the propositional contents expressed by sentences in which they occur. How can a broadly inferential approach to semantic content be extended from the grammatical category of sentences, the only sort of expression directly involved in inference, to various subsentential categories such as singular terms and predicates? For in the absence of contents corresponding to these categories, it would not be possible to understand important sorts of inferences, paradigmatically those codified explicitly by the use of identity and quantificational logical locutions.

Third, when the semantic theorist seeks to express conceptual contents explicitly, and so to reason about them—for instance when a question has arisen concerning how a certain remark should be interpreted—the semantic vocabulary employed includes, not only the logical locutions that have been construed as making *inferential* relations explicit, but also *representational*

locutions that should be understood as making referential relations explicit. Such locutions make it possible to say what someone is talking about, what is being referred to, what a belief is of or about, or what would make it true. How can a broadly inferential approach to semantic content account for the representational features of content that are expressed explicitly by means of such locutions?

The first of these prima facie difficulties has already been addressed (see Chapter 4). Although entry and exit moves are not themselves inferential moves, neither the noninferential acknowledgments of doxastic commitments that proximally terminate perceptual entries (as distinct from mere differential responses) nor the acknowledgments of practical commitments that noninferentially initiate actions (as distinct from other performances) can be understood apart from the role they play in the game of giving and asking for reasons, most directly as premises for cognitive reasonings, and as conclusions of practical reasonings, respectively. Perceptual reports are to be distinguished from mere reliable differential responses generally by their liability to demands for justification and their utility in providing justifications for other claims. Actions are to be distinguished from behavioral performances generally by their responsibility to assessment and deliberation concerning the inferentially articulated responsibilities they incur and discharge. So not only do perceiving that a content is true and acting so as to make it true involve endorsement of the inferential propriety of the move from the circumstances in which one is entitled to produce such a performance to the consequences one becomes committed to thereby, but those circumstances (of action) or consequences (of perception) themselves are inferentially significant.

The third of the cited challenges to an inferential approach to semantic content concerns its explanatory adequacy to the phenomena that make representational approaches to semantic content attractive and unavoidable. This is the most important and difficult issue. The general strategy for responding to it that is pursued here is to attempt to explain, in terms of the inferentially articulated social scorekeeping practices that institute discursive deontic statuses, what is expressed by the central sorts of representational semantic locutions. Where this can be done, the result is an account of what the theorist is *saying* when making claims about what represents what. Chapter 5 began this discussion by explaining the use of 'true' and 'refers' or 'denotes' (and so one crucial sense of 'represents') in terms of *anaphoric* links between expression tokenings. Chapter 8 completes the official treatment of representational locutions by specifying in discursive scorekeeping terms what it is to use locutions to make propositional-attitude ascriptions *de re*. This is the trope that makes it possible to specify what we are talking or thinking *of* or *about*, what objects our beliefs are directed at. This is the essential use in virtue of which expressions are properly interpreted as expressing attributions of ofness or aboutness in the intentional or semantic sense.

Both the discussion of what it is for a belief or claim to be of or about an object, or to be true of an object, and the discussion of 'refers', however, require that the inferentialist account of conceptual content be extended to subsentential expressions, paradigmatically singular terms and predicates. So the treatment of what is expressed by the central, explicitly representational locutions requires that the second challenge to the inferentialist order of semantic explanation be addressed. This should come as no surprise. For although some semantic thinkers (Davidson and Stalnaker are recent examples) conceive representational relations as obtaining in the first place between propositionally contentful intentional states and facts or states of affairs, they are in a distinct minority. Most representationalists have not taken the pragmatic priority of the propositional to entail a corresponding priority in the semantic order of explanation of conceptual contents. The more common position holds that the notion of representation is to be understood, to begin with, in terms of the representation of *objects*, particular things, and their properties and relations. According to this way of thinking, the basic representational bonds—in terms of which, for instance, the capacity for propositional representation, the capacity to represent possible states of affairs, is to be accounted for—are taken to be those linking represented objects to object-representings and represented properties to property-representings. If something like this turned out to be correct, adequate explanations of the function of attributions of representational purport and success could not be conducted entirely at the level of propositional contents.

There is an interaction between one's choice of semantic primitives (inference or representation) and one's choice of grammatical categorial primitives (sentences, or terms and predicates). The interaction is motivational rather than strictly conceptual, though—it is not that commitment to one semantic order of explanation entails commitment to a particular categorial order of explanation, or vice versa. Leibniz, who may serve as a paradigm for pre-Kantian inferentialists generally, begins his account with concepts standing in essentially inferential relations of inclusion to one another. Propositional contents are reached only by suitably combining these independently contentful items. So semantic inferentialism can coexist with a bottom-up categorial strategy.[1]

Conversely, semantic representationalism is compatible with a top-down categorial strategy, which takes the fundamental sort of content to be propositional. Representing states of affairs, purporting to represent facts, need not be thought of as semantically decomposable. If the propositions represented are thought of, for instance, as sets of possible worlds, there would seem to be no necessity to continue by explaining the capacity to represent these things in terms of more primitive capacities to represent objects or properties.[2] Talk of objects and object-representings and properties and property-representings would then proceed in terms of role in propositions and proposition-representings (as it does for Kant).

The bottom-up categorial strategy is obliged to explain propositional se-

mantic contents at some point, however, for these are the contents expressed by sentences, the only expressions with which, as Wittgenstein says, one can make a move in the language-game. Failure to ascend to an account of such contents would disqualify a theory as a semantic theory, for it would sever that theory from any account of the use of linguistic expressions or the significance of beliefs. It is precisely the role it plays in explaining the proprieties of the use of linguistic expressions or the possession of intentional states such as belief that qualifies something associated with those expressions as a semantic content. Dually, the top-down categorial strategy is obliged to explain the subpropositional contents expressed by subsentential expressions such as singular terms and predicates. Failure to descend to an account of such contents would doom a theory to explanatory inadequacy, for it would then be able to make no sense of the connection between saying something (expressing a proposition) and talking about something (characterizing an object). While these two ought not to be identified at the outset, the latter phenomenon is too central to our understanding of what we are doing when we talk and think simply to be ignored. Unless it accounts for the possibility of representing particular objects, a semantic theory will not address the concerns that many have taken to define its topic. The relation between these categorial strategies may be compared in this regard to that between inferentialist and representationalist commitments to fundamental semantic concepts. Each reductive order of explanation must account for the notions treated as primitive by the other, or independent accounts must be offered of each sort of primitive, together with a theory that specifies how they collaborate.

2. Freestanding and Ingredient Contents

The conclusion is that any account of the representational character of propositionally contentful states, acts, and utterances is obliged to offer a reading of singular reference (the representation of particular objects) and of property-representation. For the link between belief and particular objects is a sort of paradigm of representational directedness. Offering such a reading requires looking at subpropositional contents and the way in which one expression can occur as a semantically significant component in another. The only sort of contentful expressions that have been officially discussed so far are sentences (and a very special sort of sentential operator). So it will be well to begin by considering the concept of sentential embedding in general: how the content of one sentence can contribute to the content of a compound sentence in which it is embedded as a semantically significant component. Starting with the special case in which the only subsentential components considered are themselves sentences has the advantage that this grammatical category can already be specified and understood in terms of another aspect of its use, its directly inferential significance in expressing assertional commitments. Thus a sense can be given to the question, What is the relation

between the sort of content relevant to this fundamental assertional and inferential use of sentential expressions and their derivative use as (in general) unasserted components of assertible sentential expressions? With a grip on this relation it will be possible to move on to consider the contents of expressions whose *only* use is as unasserted components of sentences, paradigmatically singular terms and predicates.

The primary job of a concept of semantic content, it has been emphasized, is to account for the pragmatic significance of the states, performances, and expressions that are understood as exhibiting such contents. The more specific theoretical commitments that have been forwarded so far are intended to fill in notions of content and significance that can satisfy this basic principle. These subordinate endorsements include the practical and normative understanding of those significances in terms of deontic states, the social-phenomenalist understanding of those deontic states in terms of discursive scorekeeping by adoption of socially perspectival deontic attitudes, the idea that the sort of practice or use to begin with is linguistic, the idea that linguistic practice is distinguished by its government of assertional performances, the idea that assertional uses are essentially inferentially articulated, and the idea that inferential involvements correspond to propositional contents.

Frege builds a basic structure of semantics and pragmatics into his system from the beginning, distinguishing accounts of the significance of judging, under the heading "theories of force," from accounts of the contents judged, under the heading "theories of content." As part of his specification of the task of the theory of content, Frege recognizes that expressions can be contentful not only in the sense that a certain force can be attached to their utterance but also in the sense that their occurrence expresses something about the content, in the first sense, of sentences in which they appear.[3] As Dummett puts the distinction:

> In speaking of sentences themselves there are two different ways in which we may regard them; and these may give rise to two distinct notions of [content]. On the one hand, we may think of sentences as complete utterances by means of which, when a specific kind of force is attached, a linguistic act may be effected: in this connection, we require that notion of [content] in terms of which the particular kind of force may be explained. On the other hand, sentences may also occur as constituent parts of other sentences, and, in this connection, may have a semantic role in helping to determine the [content] of the whole sentence: so here we shall be concerned with whatever notion of [content] is required to explain how the [content] of a complex sentence is determined from that of its components. There is no a priori reason why the two notions of [content] should coincide.[4]

It is this second notion, and its relation to the first, that is the current topic. The technical terms Dummett introduces to capture the two dimensions

of sentential content that Frege discerned are "freestanding sense" and "ingredient sense." Each of these indicates an explanatory role that the notion of content as truth conditions has been thought to play: settling, in context, what the assertor of a freestanding (unembedded) sentence with that content thereby becomes committed to, and settling, in sentential context, the freestanding content of a compound sentence in which it is an (embedded) ingredient. Understanding these relations is particularly important from the point of view of a strategy, such as the present one, that seeks to work backward from notions of commitment and inference to notions such as truth conditions and representation. How should the notion of ingredient content be understood, and what does it have to do with talk about truth?

To begin with, it may be pointed out that in the passage above, where the bracketed word 'content' has been inserted, Dummett writes "truth-value." He is discussing Frege, and in the semantics of Frege's extensional logic, the concept of truth-value plays both sorts of role. Truth is what matters for the force of assertions of freestanding sentences. For it is what is preserved by good inferences, in particular the inferences that are good in virtue of their logical form—the ones Frege is codifying. Furthermore, possession, by a logically compound sentence, of the property preserved by logically good inferences is determined by the truth-values of its component sentences. When the same formal apparatus is maintained as much as possible—consistent with letting different notions play these two roles—the result is classical multivalued logic.

3. Multivalued Logic

The standard way of presenting these semantic ideas is as part of a bottom-up compositional definition of logical connectives, and of the validity of compound sentences formed by their use. The semantics is provided by a generalization of truth tables, defined not over Frege's two (truth) values but over many, perhaps an infinite number. Corresponding to each n-ary syntactic compounding device is a function mapping n-tuples of values assigned to component sentences onto the value assigned to the compound sentence in which they are components. These functions are most easily visualized in the form of the familiar sort of table:

!	[*1]	[2]	[3]
[*1]	1	2	3
[2]	1	3	3
[3]	2	3	1

According to this table, for instance, an interpretation that assigns p the value [2] and q the value [3] must assign $p!q$ the value [3].

Since the original role played by the notion of truth-value is being bifur-

cated, it is best not to beg questions by continuing to employ it for one or the other of these notions. The values in the set {[1], [2], [3]} may be called *multivalues.* One or more of the multivalues is distinguished or *designated* (indicated by the '*' attached to the multivalue [1] in the table). A compound formula is valid in virtue of its form in case it is assigned a designated value no matter what multivalues are assigned to its component sentences.

Designatedness here indicates whatever the force-relevant notion is, for instance truth or, more generally, what is preserved by good inferences (which might in another context be some sort of commitment or entitlement). According to such a scheme, an interpretation assigns each sentence two sorts of value: as designated or not, and as having a certain multivalue.[5] The designatedness value includes everything that matters for the pragmatic significance of the freestanding uses of the sentence (to which assertional force can be attached) as far as it is represented by this formal apparatus. Differences between sentences that are assigned the same designatedness value (in the example, designated or not designated) are significant at all only insofar as they affect the designatedness of compounds containing them. The two undesignated multivalues in the example differ in that substituting one for the other changes not only the multivalues but the designatedness of some compounds containing them.

The standard, bottom-up direction of explanation exploits this apparatus to move from an antecedent set of multivalues that can be associated with sentences, and from functions antecedently associated with compounding devices, via a notion of designatedness, to attributions of formal validity. The same apparatus, however, can be exploited in the service of the converse, top-down direction of explanation. Then the move is from antecedently understood attributions of material designatedness to assignments of multivalues to sentences and of functions to compounding devices. The essential principle is that if two sentences have the same multivalue, then substituting one for the other never changes the designatedness of any compound sentence in which they can appear as components. This is what is meant by saying that the multivalues express the contribution a sentence makes to the designatedness value of compounds containing it.

Since any sentence can be regarded as a degenerate compound containing itself, it follows from this principle that two sentences with the same multivalue must have the same designatedness value. This is what justifies the usual procedure—embodied both in the standard tabular way of setting out semantic definitions of connectives in multivalued logic and in the definitions that generalize it to semantic matrices—of treating multivalues, rather than sentences, as what take designatedness values. Turning the basic principle around, two sentences can be treated as having the same multivalue just in case substituting one for the other never changes the designatedness value of a compound sentence in which one appears as a component. In this way, sentences are assimilated into co-multivalue classes—and so taken as

sharing their ingredient contents—accordingly as their intersubstitution as components of compound sentences preserves the designatedness values of those compounds. Lindenbaum can be understood as employing an extreme form of this strategy in his mechanism for constructing, from the set of theorems of a logic meeting certain general conditions, a matrix of multivalues and compounding functions defined over them that would validate just those theorems, by identifying multivalues with equivalence classes of logically interderivable sentences, and designatedness with theoremhood. There is no guarantee that this procedure will not end up with an infinite number of small equivalence classes (as it does in the standard Lindenbaum algebra for the propositional calculus).

How finely the ingredient contents are individuated by this substitutional test depends on the expressive power of the language, specifically on what sentential embedding contexts and embedded sentences are discerned in it. Strictly speaking, substitutional assimilation according to multivalues or ingredient contents is always relativized to a class of embedded sentence occurrences, and so to a class of sentential embedding contexts. It is not implausible that in natural languages, for any two lexically distinct sentences, there is some context in which substitution of one for the other can affect the assertional designatedness of the compound sentence resulting from such substitution. For instance, 'S now thinks (or wishes) that p' is quite discriminating. This fact need not rob the substitutional form of analysis of its usefulness, for relative to various restricted classes of contexts, important assimilations are brought about nonetheless. Indeed, the partial ordering on sentential contexts that is brought about by looking at proper-inclusion relations among the multivalue equivalence classes they generate can contain interesting information about the semantic relations between those compounding devices.

Two embedding contexts can generate the same multivalues (in case they sort possible embedded sentences into just the same equivalence classes), or one can cut finer than another. Suppose, though, that every sentential embedding context that is discerned yields a different way of carving up the embedded sentences into equivalence classes, in a crazy-quilt of overlapping classes exhibiting no substantial identities or inclusions. In that case there would seem to be no theoretical advantage to discerning the semantically significant occurrence of one sentence in another. Occurrence of a sentential expression as a lexical part or syntactical subunit of another sentence is neither necessary nor sufficient to make it appropriate to discern the semantically significant occurrence of one sentence in another. The theorist may discern such occurrences where there are no lexical sentences, as with an embedded expression such as 'Kant's claim about Aufklärung and responsibility', or may deny them where there are, as Quine would do with direct quotation of sentences uttered. Discerning subsentential structure is enlightening only insofar as the assimilation of embedded sentences shows how the

capacity to use the embedded sentences, together with the capacity to use some of the embedding sentences, could generalize to a capacity to use the compound sentences with arbitrary embedded components.

It is worth considering a somewhat different, but closely related, point that Dummett makes, in connection with an approach to linguistic theorizing he associates with Wittgenstein:

> One way in which these passages from Wittgenstein may be taken is as rejecting the whole idea that there is any one key idea in the theory of meaning: the meaning of each sentence is to be explained by a direct characterization of all the different features of its use; there is no uniform means of deriving all the other features from any one of them. Such an account would have no use for any distinction between sense and force: while it could admit some rough classification of sentences, or particular utterances of sentences, according to the kinds of linguistic act effected by means of them, it could cheerfully regard the totality of such types of linguistic act as unsurveyable—as Wittgenstein does— and would not need to invoke the classification of linguistic acts in its accounts of the meanings of particular sentences . . . The difficulty with such a theory is to see how it could do justice to the way in which the meanings of sentences are determined by the meanings of the words which compose them. The great strength of a theory which admits something as the key concept for the theory of meaning—at least a theory which is as developed as that of Frege—is that it displays a plausible pattern for the determination of the meaning of a sentence by the meanings of the constituent words . . . If nothing is to be taken to be a key concept, then we are once more without any conception of what the meaning of a word, as opposed to that of a sentence, is taken to be.[6]

The idea is that understanding a word need consist only in understanding the contribution it makes to the sense or content of sentences containing it. From there, the speech-act theory is to explain how that content contributes to the force or significance of various sorts of performances involving it. Otherwise, understanding the word requires mastery of the contribution it makes to all of the different acts that can be performed by means of it.

What underlies the analogy

$$\text{force : sense}$$
$$\text{(or pragmatic significance : semantic content)}$$
$$::$$
$$\text{sentential content : subsentential content}$$

is the thought that each variety of significance a performance can have sorts sentences that may be uttered with that force or significance into content-equivalence classes accordingly as intersubstitution preserves it. Unless dif-

ferent kinds of force or significance sort sentences into content-equivalence classes in the same way, no theoretical advance is made by discerning contents in addition to significances. Just so with assimilation of subsentential components preserving the content (or significance) of sentential compounds.

It was said above that although Frege is the first to take seriously the requirement that some aspect of semantic content determines the contribution a contentful expression makes to compounds in which it occurs as a component, nonetheless in the semantics for his logic he employs one notion, truth-value, to play both freestanding and ingredient roles. He there codifies inferences that depend only on the logical form of the sentences involved, and not on their material or nonlogical content. For these purposes he finds that it is possible to treat truth- (or commitment-) preservation not only as necessary for goodness of inference but also as sufficient. Thus for the inferences codified by his classical conditional, not only is designatedness preserved by good inferences, but any inference that preserves designatedness is a good one. For this compounding device, the two-valued conditional, sameness of designatedness value (which does duty here for freestanding content) is sufficient for sameness of multivalue (which does duty here for ingredient content). In this sentential context, the force-relevant content determines the role of sentences as components as well.

Designatedness-functional contexts such as this may be said to embed *homogeneously* with respect to designatedness values, since those values are all that matter in determining the contribution made by an embedded sentence to the designatedness value of the whole (that is, they can serve as multivalues). This term is used to mark off one of the several distinct senses sometimes attached to the expression *extensional.* Whether or not a sentential context is homogeneous in this sense concerns the relation between designatedness values and multivalues. It is quite independent of any specific conceptions of what plays the role of the immediately pragmatically relevant freestanding content. That role could be played by an antecedent concept of truth values or (looking ahead) by a concept of truth conditions.

Truth is preserved by good inferences of a certain important class. That class can be thought of as corresponding to deductive inferences, provided the notion is broadened beyond the concern with formally good inferences that is traditionally tied up with the notion of deduction. (Here the principle that a good inference never leads from premises that are true to a conclusion that is not true is being thought of as only a necessary condition on the goodness of inferences.) The inferences in question are just the commitment-preserving ones, ('committive inferences', for short). The pragmatic force of freestanding utterances of the expressions that can take truth-values of the sort preserved by good inferences (that is, sentential expressions) is assertional commitment, overtly acknowledged by, and so appropriately attributed to, the utterer. If the inference in question is a good one of this sort,

then to be committed to the premises is to be committed to the conclusion. No further understanding of the notion of truth-value in its role as designatedness value is required in order to proceed to the assimilation of sentences into multivalue equivalence classes.

Start with any set of concomitantly attributed or undertaken commitments to claims, some of which are expressed by sentences that are sententially compound (in that substitution for sentences that occur as their components makes sense). Then assign to two sentences the same componential value or multivalue (relative to that set of commitments and that compounding vocabulary) just in case substituting one for the other never turns a sentence expressing a claim in the set of concomitant commitments into a sentence expressing a claim that is not in that set. It follows that if two sentences are componentially equivalent, then they have the same designatedness value—the commitments in question include either both of them or neither.

To recap: The sort of content that has been considered here previously is the broadly inferential content that determines the correct uses of freestanding sentential utterances, paradigmatically the significance of asserting them. Content understood in this way can be associated only with expressions whose freestanding use has a pragmatic significance. It is not available as an interpretation of the contribution made by the occurrence of essentially subsentential expressions, such as singular terms and predicates. Following Frege and Dummett, a further sort of sentential content, 'ingredient' content, has been discerned, corresponding to the role that sentences can play as components of compound sentences. Although freestanding content may play the role of ingredient content (in homogeneous contexts), in general the latter is not reducible to the former. Ingredient contents are a sort that can coherently be attributed to expressions functioning only as components of assertible sentences, although so far only the contents to be associated with sentential sentence components have been considered. In the usual synthetic use of contents as multivalues, to define logical connectives, one begins with contents of this sort and determines designatedness values and, eventually, formal validity by their means. But the same apparatus can be exploited analytically, to move down from a notion of formal validity (as Lindenbaum does) to the assimilation of sentences according to their componential roles or, as has just been seen, from a notion of material designatedness (for example as assertional commitment) to multivalue equivalence classes. The mechanism whereby a simple notion of ingredient content (multivalue) is extracted from a simple notion of freestanding content (designatedness value) is purely substitutional. Two expressions are assimilated as making the same contribution to compound sentences in which they occur relative to some property of freestanding sentences just in case substituting one for the other never changes an embedding sentence from one that has the property to one that does not. As concern shifts to material, rather than formal, issues,

validity ceases to be the key notion, and designatedness comes to the fore, as the topmost property with respect to which substitutional invariances are assessed. The route to the notion of semantically significant occurrences of subsentential expressions, then, goes through the notion of substitution.

II. SUBSTITUTION, SENTENTIAL EMBEDDING, AND SEMANTIC ROLES

1. *Substitution and Subsentential Content*

Frege's notion of *substitution* is the key to appreciating the characteristic theoretical role played by concepts of semantic content. This point begins to emerge when it is noticed that in the story just told, the relations between designatedness and multivalue, on the one hand, and between validity and designatedness, on the other, are of the same general sort. The fundamental pragmatic status that a notion of content is to help keep track of is that of assertional or doxastic commitment. As the previous chapter argued, this is a notion sufficiently intimately tied to that of truth claim that whatever sort of content ends up accounting for the pragmatic significance associated with that status for that reason has credentials as explicating one important dimension of truth talk.

Given the general understanding of the relation between material and formal proprieties of practice that has been urged earlier, the concept of the formal logical validity of claims should be treated as derived from that of material assertional commitment. The means of derivation are straightforwardly *substitutional:* A (logically) valid claim is one, first, that is designated (to which one does or ought to undertake or attribute commitment) and one, second, that cannot be turned into an undesignated claim by any substitutions restricted to a special class of vocabulary. In the case being considered, the nonlogical vocabulary consists just in the component sentences, from which the compound sentence is conceived as resulting upon the application of a logical sentential connective. If a sentential context is not valid, in that not all substitutions preserve designatedness, then it may be substitutionally homogeneous (designatedness-functional), provided that substitution within codesignatedness classes preserves designatedness. If not, then intersubstitution within multivalue classes, substitutionally heterogeneous with respect to designatedness, by definition will preserve designatedness. Valid claims are just those special sentential contexts with respect to which the multivalue substitutional equivalence class assimilates all sentences. Ordinary claims—which are not valid with respect to substitution for components generally, nor with respect to codesignated components—are valid with respect to substitution for component sentences by sentences sharing a multivalue. The same substitutional structure is responsible for moving up from material assertional commitments to assertional validity and moving down

from material assertional commitments to multivalues or ingredient contents.

If logically valid sentential contexts are just those that assimilate all sentences into one single multivalue equivalence class, what is their special interest? They are of interest because the way those valid contexts are compounded out of other, nonvalid ones has much to teach about those nonvalid contexts, which include the basic sentential connectives. Any metatheory that identifies a logic with the set of its theorems is committed to understanding the semantics of logical expressions only insofar as it is expressed by the capacity of those expressions to enter into assertionally valid combinations. As was pointed out in Chapter 2, Dummett correctly argues that this is an unduly restrictive view of the subject matter of logic.[7] He would identify that subject matter by reference not to the theorems characterizing a logic but to its derivability relation. For classical Boolean logic these two notions are equivalent, but in general the theorems need not settle the derivability relation. He shows how the apparatus of multivalues can be applied in the definition of valid inferences, and not just in the definition of valid claims. He is concerned with defining validity from antecedent sets of multivalues, that is with the synthetic rather than the analytic use of this substitutional machinery. And since his topic is logical validity, he is concerned only with formal, and not with material, inferences. But the point he makes carries over to the analysis of material inferences and the derivation of a notion of material ingredient content, where its real significance becomes apparent.

A move from material assertional commitment as designatedness to material inferential commitment as designatedness corresponds to the move Dummett recommends from formal assertional validity to formal inferential validity as the notion with respect to which substitutional equivalence is assessed. The suggestion is to look at inferential commitments and correctnesses of inference instead of, or as well as, looking at assertional commitments and correctnesses of claims. A condition on the individuation of sentential contents as inferential roles can be generated from the notion of goodness of inferences by considering two sequential applications of the methodology of substitution that generates multivalues from the designatedness of compound sentences. Extending Frege's usage, two claims can be said to have the same inferential content just in case substitution of a token of the one type for a token of the other never turns a good inference into one that is not good, no matter whether the sentence appears as a premise or as part of the conclusion of the inference.

This principle does not depend on the existence in the language in question of sentential operators producing compound sentences in which other sentences are embedded. An inference here can be thought of as a pair of sets: of premise claims and of conclusion claims. Inferences can be treated as themselves a sort of compound in which sentences can appear as embedded

components, and such inferences can be classified as 'designated' (good) or not. If attention is restricted to inferences involving only freestanding occurrences of sentences, the equivalence classes of claims defined by preservation of goodness of inference on intersubstitution within the class may be called 'freestanding inferential contents'. Inferential contents so defined are generated just the way multivalues are, except that instead of looking at the designatedness of compounds such as conditionals as what must be preserved by substitution, one looks at the goodness of inferences. They are the products of the first application to the analysis of inferences of the substitutional methodology suggested by multivalues. The result is just what Frege defined as "begriffliche Inhalt" (conceptual content) at the beginning of the *Begriffsschrift*, as discussed in Chapter 2.

Component conceptual contents can then be defined by a second application of the substitutional methodology that introduces multivalues—this time to a case where what must be preserved on substitution is not the designatedness of compound sentences but their inferential contents, which were constructed substitutionally by the first application of the analogy with multivalues. The result is to put into play two concepts of broadly inferential content: contents as the inferential potentials of freestanding sentential utterance (including both their employment as premises and as conclusions of inferences), on the one hand, and contents as the contribution a sentence makes to the inferential content of compound sentences in which it appears as a component, on the other. This latter sort of content, which may be called the 'component content' of a sentence, arises from considering substitution within compounds, rather than within inferences. Putting the two definitions together, it follows that two sentences have the same component content if and only if substitution of one for the other as embedded components of any compound sentence never turns from good to not good an inference in which the compound sentence appears freestanding. Assimilating sentences accordingly as their intersubstitution in inferences preserves the material goodness of inferences yields freestanding content equivalence classes, and assimilating them accordingly as their intersubstitution in sentential compounds preserves freestanding content, yields component or ingredient content-equivalence classes. On the side of assertional commitments and proprieties, beginning with material designatedness of compound claims yields one level of further substitutional assimilation, namely multivalues. On the side of inferential commitments and proprieties, the substitutional machinery can be applied twice—once to yield a notion of freestanding inferential content, and once again to yield a notion of component inferential content. (It is irrelevant for this contrast that in either case the top-level material notion, whether assertional or inferential, can also be used substitutionally to define notions of formal validity.)

The first step in generating the inferential hierarchy of substitutional levels of content was made by noticing that Frege's substitutional definition

of equivalence of conceptual contents from material goodnesses of inference is analogous in structure to the definition of equivalence of multivalues from truth-as-designatedness, that is, material goodnesses of claims. The main adjustment required for this analogy is that inferences must be treated as a kind of context in which sentences can appear embedded, as premises and conclusions, and which as a whole is assessable according to its correctness. The second application of the substitutional machinery is more closely analogous to the assertional designatedness-multivalue paradigm in that only substitution within compound sentences in which other sentences occur as components is envisaged. It is less closely analogous to the paradigm than the first step in that what is preserved as the test of assimilation (corresponding to multivalue equivalence in the assertional paradigm) is in the inferential case not an on/off property, designated or not designated, but possession of a certain freestanding inferential role or value—of which there are many, perhaps infinitely many. It is by way of preparation for this point that the initial account above of the relation between designatedness and multivalues speaks of intersubstitution within multivalue classes as preserving designatedness *value*, even where this just means preserving designatedness.[8] The substitutional conceptual machinery as such is indifferent as to whether what is preserved is membership in a single class (the designated ones), alike for all compound sentences whose components are being varied, or membership in whichever element of some partition the compound sentence whose components are being varied belongs to. That there is no technical difference does not mean, however, that there is no difference in the explanatory value of applying the technical machinery.

Dummett robustly acknowledges the requirement that a notion of semantic content qualifies as such only by its relevance to the pragmatic significance of acts, for which asserting serves as a prototype. He is concerned to argue that this requirement means that the substitutionally topmost level of interpretation, the level of designatedness, must be two-valued or on/off, since what must ultimately be settled is whether an assertion is or is not correct (assertible).[9] As a general point, this seems dubious—perhaps a normative pragmatics need not be founded on the application of the dichotomy correct/incorrect to performances such as assertions. The discursive scorekeeping account offered here is substantially more complex, as not only are commitment and entitlement distinguished, but track is kept of which are undertaken and which attributed—all articulating various ways in which a claim or an inference can count as correct. Again, thinking of pragmatic status as what must be preserved upon intersubstitution of sentences sharing a semantic content may be too narrow a formal paradigm—perhaps semantic contents can determine the correctness of material inferences without having for that reason to be conceived as preserved by good inferences.

Whether or not this on/off requirement could be shown to apply to semantic interpretation generally, there is a sense in which it is satisfied by the

inferential hierarchy that is here laid alongside the assertional one that Dummett considers. The topmost notion there is the goodness of inference, which can be thought of as a yes/no, correct/incorrect, two-valued affair: of designatedness rather than 'designatedness values'. This is because, though the purely substitutional machinery does not require it, the topmost notion in these hierarchies is a pragmatic one, as Dummett urges; furthermore, the sort of pragmatics being pursued here is one of deontic status and social attitude (commitment and entitlement, attributing and undertaking), and these are conceived as either characterizing an individual or not. The three-leveled inferential hierarchy shows, though, that the assessments of correctness that generate this two-valuedness at the top need not be directed in the first instance at sentences. When one starts with inferences, sentences are assimilated into many inferential role-equivalence classes, not simply into those that are designated as correct and those that are not.

2. Two Concepts of Extensionality

The assertional interpretive hierarchy—of designatedness as substantive assertional commitment and multivalues as equivalence classes of component sentences intersubstitutable saving the designatedness of compounds—gave rise to a natural notion of extensionality for sentential contexts as consisting in componential homogeneity. In this sense a context is extensional if the multivalue equivalence relation need cut no finer than the codesignatedness classes. How does the componential notion of extensionality apply to the inferential interpretive hierarchy? Since sentences are not inferences, the relation between designated inferences and the inferential roles of freestanding sentences cannot be homogeneous. It cannot strictly be that all one needs to know about freestanding sentences in order to assimilate them in such a way that intersubstitution within the resulting classes will preserve goodness of inference is whether or not they are good inferences, for they are not inferences at all. It will be necessary to look elsewhere for an analog of this sort of extensionality at the top level of the inferential hierarchy. At the lower level, when what is at issue is the relation between the inferential roles of sentences and their componential roles, however, a notion of extensionality as homogeneity does apply, since it is sentences in both cases that are assigned such roles.

A sentential context in which sentences can appear embedded as components is extensional in the sense of being componentially homogeneous just in case substitution of one claim for another with the same inferential role never alters the inferential role of the compound sentence containing them. It is a criterion of adequacy on semantically explicitating vocabulary—which has been picked out here as deserving to be called specifically *logical* vocabulary in virtue of playing that expressive role—that it generate embedding contexts that are inferential-role-functional or homogeneous in this sense.

Thus the inferential role of a conditional claim is to be settled by the inferential roles of its antecedent and consequent. Not all vocabulary is like this. In some of its uses, for instance, the inferential role played by the important, pragmatically explicitating expression 'S claims that p' in some speaker's mouth depends not on the inferential role played by p for that speaker but on the role it is taken to play for S.[10]

Componential homogeneity is a concept that has application only within a substitutional hierarchy. There sentences are assimilated at the lower level, as associated with one sort of semantic interpretant, in case substituting one for another does not alter the assimilation at a higher level. That assimilation corresponds to association of another sort of semantic interpretant with sentences that depend in some way upon the sentences substituted for. A different sort of reducibility can be conceived in terms of the relations between the two-leveled hierarchy of assertional semantic interpretation and the three-leveled hierarchy of inferential semantic interpretation. A particularly strong bond between the assertional and the inferential orders would be forged if the designatedness of sentences determined the designatedness of inferences involving those sentences. The topmost assertional level of interpretation (committed/not committed) would then assimilate sentences into inferential-role equivalence classes, and assertional multivalues would coincide with inferential component contents. Commitment to the goodness of inferences would be preserved by substitutions for premises and conclusions, provided those substitutions preserve assertional commitment.[11]

Notice that the concept of multivalue is not equivalent to that of component content unless never turning a designated claim into one that is not designated is sufficient for never turning a good inference into one that is not good. Since in any case preserving designatedness is a necessary condition of a good inference, sameness of component content will guarantee sameness of multivalue, but not in general vice versa. Multivalues capture the contribution that component sentences make to only the designatedness-functional inferences involving the compounds they are embedded in. Commitment to the goodness of an inference in this sense of goodness is what is expressed by the assertion of a classical two-valued truth functional, so-called material conditional. In connection with assessments of the formal correctness of certain kinds of logical inferences, treating preservation of assertional commitment as sufficient as well as necessary is not an entirely useless strategy, as Frege shows. But the principle it embodies is simply false if applied to genuinely material inferences, whose correctnesses constitute the possession of material content by the assertible sentences that appear as their premises and conclusions.

That inferential commitments should be determined (in this substitutional sense) by assertional commitments regarding their premises and conclusions—which is an interhierarchy rather than an intrahierarchy relation—is, however, another sense that has sometimes been associated

with the notion of extensionality. Dummett, in the chapter in which he discusses the distinction between truth-value as designatedness and truth-value as ingredient or multivalue, considers the tension in Frege between two conceptions of Bedeutung. In one sense this technical term is used just to mean something like 'semantic role', defined by substitutional assimilations. In another sense the relation between an expression and its interpreting Bedeutung is understood to be modeled on that between a name and its bearer. At this stage in the present exposition, names and bearers are not among the concepts it is officially permissible to pretend to understand. Sentences, and their pragmatic and semantic correlates—that is, assertional commitments and inferentially articulated propositional contents—are all that are onboard so far. But Bedeutung notoriously embraces not only particular objects as referred to by singular terms but also truth-values, taken by sentences. The analog at the categorial level of sentences, to the tension at the categorial level of terms between a substitutional notion of Bedeutung as semantic role and a representational notion of it modeled on the name/named relation, is the 'tension' between notions of content derived from the inferential hierarchy and those derived from the assertional.

For substitution within assertional codesignatedness classes to preserve inferential designatedness is only one way in which the assertional interpretants might determine the inferential ones. Another possibility is that two sentences might have the same freestanding inferential role in case they have the same multivalue or assertional component content. Whether or not this is so depends on the expressive resources of the language, on what sort of sentence-forming locutions it makes available. Where these resources include conditionals, since these codify inferential commitments as explicit assertional commitments, the assertional designatedness of conditionals will vary with substitution of antecedent for antecedent and consequent for consequent, unless the substituends share their freestanding inferential roles. Different kinds of conditionals may codify different classes of inferences, each of which defines a correlative substitutional notion of inferential role. Where the inferences corresponding to that role are expressible by conditionals in the language, assertional multivalues must cut as fine as freestanding inferential contents.

Of course, the assertional multivalues may partition the sentences into even smaller classes, as they will if the language permits compound sentential contexts interpretable as having the form 'S believes that if . . . then q'. Whether or not the expressive resources of the language suffice to establish a general determination of freestanding inferential contents by assertional multivalues depends not only on what conditionals exist but on how they behave. If only finitary conjunction is available, for instance, assertional multivalues adequately represent freestanding inferential contents in general only if the language is compact. Again, if infinitary conjunction is available, it must be able to form 'enough' conjunctions, and so conditionals. In any

case, it is clear that the proper order of explanation runs from inferential role to assertional codifying locutions (such as conditionals), to a multivalue defined substitutionally with respect to assertions formed by the use of those locutions, not the other way around.

That is, one does not start with an intrasentential notion of assertional multivalue and then use that to define conditionals, and those to define inferences. Freestanding inferential contents must be defined as part of the same conceptual package as assertional designatedness. Inferential commitment must be considered along with assertional commitment. For it is its inferential role that determines what asserting a sentence commits and entitles one to, and what could commit or entitle one to it. Apart from such inferential involvements, an assertional commitment would be without content. The assertional hierarchy of interpretation should not be conceived as independent of and antecedent to the inferential one. It may be noticed, furthermore, that even if sameness of assertional multivalue ensured sameness of freestanding inferential content, it would not follow that it ensured sameness of inferential component content, unless sameness of freestanding inferential role were sufficient for sameness of inferential component content, that is, unless the inferential hierarchy were componentially homogeneous.

If for these reasons the material-inferential interpretive hierarchy should not be seen as derivative from the material-assertional interpretive hierarchy, what about the other way around? Within these substitutional hierarchies there is a definite sense to the claim that one sort of content 'determines' another. Assertional multivalues determine assertional designatedness in the sense that two sentences cannot have the same assertional multivalue and different designatedness values. It is this sense in which freestanding inferential contents determine inferential designatedness, and inferential component contents determine the freestanding ones. One cannot in the same sense ask whether inferential designatedness determines assertional designatedness, since there is no one sort of thing that can take both values. One can ask whether it is possible for two interlocutors to undertake or have attributed to them just the same inferential commitments but different assertional ones. Apart from commitment to conditionals, it would seem that this possibility ought to be allowed. Two scientists or two politicians might agree entirely about what would be true if certain conditions obtained but nevertheless have quite different beliefs about what conditions do in fact obtain. Though it is true that in classical two-valued logic, fixing the truth-values of all the conditionals (corresponding to inferential commitments, of a sort) determines the truth-values of all of the atomic propositions, this property is an embarrassment and provides further good reason to deny that the classic horseshoe means "if . . . then . . ."[12] Hypothetical commitments ought not to settle categorical commitments. Rather, inferential commitments determine assertional commitments only taken together with other assertional

commitments. They cannot do the job all on their own. When purely formal inferences are at issue, there is a purely formal sort of assertion (namely of theoremhood) for which such determination can be envisioned: given a formal derivability relation M, one can consider $\{A: \Phi MA\}$, that is, the set of claims derivable from the empty set of premises. Of course the same set can be defined with respect to a material-inferential relation, but the resulting set of claims is not the only one compatible with the inferential commitments that generate it.

3. Compositionality and Decompositionality

The primary lesson that should be drawn from this discussion is that there is an intimate relationship between the notion of *semantic content* and the concept of *substitution*. That concept is one of Frege's grand themes, exploited everywhere in the official definitions of logical and semantic concepts. He is methodologically quite self-conscious about the importance of his substitutional approach—it is the basis for his technical concept of function, which in his most metaphysical writings he takes as an explicit topic for philosophical inquiry. Frege's earliest semantic and logical work introduces the concept of conceptual content in terms of substitutional behavior with respect to a kind of pragmatic significance: two claims have the same content if substituting one for the other never turns a good inference into a bad one. Goodness of inference is a pragmatic matter—in Fregean terms, a matter of force; in this paradigmatic case, it is a matter of the force of reasons. In the terms being recommended here, it is a matter of normative force, of deontic status, and so of social practice and attitude. However the pragmatic end is conceived, the route from pragmatics to semantics is that of assimilating expressions according to invariance (of pragmatic significance of some sort) under substitution. This same substitutional path that leads from inference to sentential conceptual content leads as well from the possession of freestanding inferential content by compound sentences to the possession of component-inferential content by embedded ingredient sentences and, as will appear in the rest of this chapter, from sentential content to the content of subsentential expressions such as singular terms and predicates. The substitutional way of working out a top-down categorial explanatory strategy is already implicit in the substitutional form taken by Frege's inferentialist approach to propositional content.

This *de*compositional methodology is what lies behind what is often called Frege's 'principle of compositionality'. According to that principle, the semantic interpretant associated with a compound sentence such as a conditional should be a function of the semantic interpretants associated with its semantically significant components. In spite of the way it is usually interpreted or exploited, this principle by itself is neutral between bottom-up and top-down categorial explanatory strategies. As Frege's own substitutional

understanding of functions indicates, the principle operates as a constraint on what it is for one expression to count as a semantically significant component of another, regardless of whether the compound is conceived as built up in the first place by operations on antecedently specifiable components or, conversely, the components are conceived as substitutionally precipitated out of antecedently specifiable compounds.

Typically, discussions of the compositional constraint are framed, not in terms of the generic notion of a semantic interpretant, as above, but in terms of the specific notions of sense and reference that Frege introduced in 1891. Standard sketches of the explanatory roles characteristic of those two semantic conceptions center on the following leading ideas:

1. The referent of a sentence is its truth-value, what must be preserved by good inferences.
2. The sense of a sentence is the thought it expresses, what is grasped by someone who understands it.
3. The referents associated with compound expressions are functions of the referents of their components.
4. The senses associated with compound expressions are functions of the senses of their components.
5. The sense of a sentence, together with how things actually are, fixes its referent, that is, its truth-value.

Great care is required in specifying just what commitments one is attributing to Frege under headings such as these, but the point to be made here concerns only gross structure. Again, if subsentential expressions were currently at issue, further doctrines would need to be included, most notably that the referent of a proper name is the object that sentences it occurs in are about, or on which their truth depends.

At the crude level of description expressed by these five dicta, it may be noted that there is a natural mapping of the substitutionally generated assertional and inferential semantic hierarchies onto the Fregean scheme of sense and reference. Corresponding to (1), assertional designatedness plays the role of truth-value as what must be preserved by good inferences. Corresponding to (2), what one must grasp in order to understand a sentence is conceptual content, the begriffliche Inhalt of the *Begriffsschrift*. Two claims are defined as having the same conceptual content in case substituting one for the other never turns a good inference into a bad one. If substitutions within compound sentences such as conditionals, appearing as premises or conclusions, are included (as the subsequent practice of the *Begriffsschrift* in fact does), then the restricted compositional principle (4) will obtain as well, for the ingredient inferential contents of component sentences determine the ingredient inferential contents of the compounds they occur in quite generally. If the relevant substitutions are restricted to sentences playing freestanding roles as premise or conclusion, then what must be grasped is only freestand-

ing inferential content. In that case the restricted compositional principle (4) will obtain only for substitutionally homogeneous contexts of embedding, that is, where sharing an inferential role is sufficient for sharing a component inferential content. In any case, the corresponding restricted compositional principle on the side of reference, (3), will obtain only where assertional codesignatedness is sufficient for multivalue assimilation, that is, in assertionally homogeneous sentential contexts.

Frege is the first to investigate systematically the concept of the special sort of semantic value or content that an expression can have in virtue of its contribution to the semantic value or content of compound expressions in which it appears as a significant component. The latter notion arises in the context of a substitutional understanding of semantic contents in their relation to pragmatic significances generally. The distinction between freestanding and ingredient semantic contents is blurred, however, by the fact that both on the side of reference (which in one of its explanatory functions has been identified here with the hierarchy of assertional contents) and on the side of sense (which in one of its explanatory functions has been identified here with the hierarchy of inferential contents), Frege seeks to have *one* notion from each hierarchy play both sorts of role. (Attention is still restricted for the time being to sentential interpretants, so this means truthvalues, on the side of reference, and thoughts, on the side of sense.) A less confining semantic metatheory answering to the same insights expressed in the compositional requirements (3) and (4) above is achieved if two different sorts of content—freestanding and ingredient—are distinguished, both within the (de)compositional hierarchy generated by the semantic interpretants immediately relevant to assertional force, and within that generated by the semantic interpretants immediately relevant to the inferential content that is asserted.

What about (5), the structural principle that sense determines reference? Does the inferential role (freestanding or ingredient) of a sentence determine its assertional designatedness value? In one sense, clearly not, insofar as inferential agreement among interlocutors need not entail assertional agreement. In another sense, though, it does. Since the inferential role of the sentence determines what an assertor is committed to by it, one interlocutor cannot licitly assign to two sentences the same inferential role and different assertional designatedness values. The most interesting sense of this question, however, is neither of these. It is rather whether inferential contents determine assertional commitments in a way analogous to that in which truth conditions are understood to determine truth-values in the familiar conceptions of sentential semantics that take their cue from the later Frege's willingness to specify the senses expressed by sentences in terms of their truth conditions. Characteristic of these conceptions is an understanding of truth conditions as supplying one sort of sentential content: meaning as intension, as defining a function, which given a world or set of facts, deter-

mines the truth-value of that sentence (relative to that world or those facts). Model theory and formal intensional semantics provide representationalist, bottom-up implementations of such an understanding.

The role being played by truth-values in this sort of story is that of assertional designatedness. In the rival scheme being developed here, that role is played by the notion of assertional commitment. The notion of free-standing inferential content, which is derived substitutionally from that of inferential commitment (or designatedness of inferences), is, like that of content as truth conditions, intended to specify the content of sentences in the sense of what it is that can be assertionally designated or true. The present question is whether and in what sense inferential contents can serve the function that truth conditions serve, of determining, together with the facts, truth-values-as-designatedness-values. Can the contents substitutionally extracted from the pragmatic status of inferential commitment be construed as related to those substitutionally extracted from the pragmatic status of assertional commitment as intension to extension?

There is a sense in which inferential contents, together with the facts, determine what is true (assertionally designated). Thus component contents may be seen as corresponding on the inferentially intensional side to multi-values on the inferentially extensional, that is assertional, side. Component contents, which determine inferential contents, can thus be thought of as expressing the contribution sentences make to the truth conditions-as-intensions of compounds containing them. Further consideration of the issue must be postponed until Chapter 8, because it cannot be pursued while continuing to suppress the additional level of analysis at which the deontic status of assertional commitment (designation) is resolved into social attitudes, which are explicitly relativized as to attributor and attributee.

When loose talk of deontic status is replaced by careful talk about deontic social attitudes, the essential clues will be seen to be that facts are just true claims, and that taking-true is just asserting. For each interlocutor, the inferential contents associated with anyone's sentences, together with the facts, determine which of those sentences express truths. The facts consulted in each case are the claims the attributing interlocutor takes to be true (that is, endorses) or acknowledges assertional commitment to. If inferential commitments, and so inferential contents, were uniform across a community, the determined truths and the determining facts would in every case coincide (though unless assertional commitments were also uniform, they would vary from attributor to attributor). But this is an extraordinary and degenerate sort of case, one that bears the same relation to the fundamental situation of communication involving the practice of truth assessment that a community in which assertional commitments are universally shared bears to the fundamental situation of communication involving the practice of assertion. The full story of how the inferential content attached to an expression by an interlocutor affects the proprieties of truth attributions by others, and espe-

cially the role played by (what are treated as) the facts, of how not only sentential but subsentential sense determines reference, is presented in Chapter 8 as the centerpiece of an account of the significance of representational idioms for semantic theories.

4. Summary

Dummett explicitly distinguishes the explanatory role played by freestanding and ingredient contents and recognizes Frege's accomplishment in conceiving of the latter. He also explains the relevance of multivalued logic to these two notions of semantic content by pointing out that the two notions of truth-value in play in such logics—what have been called here 'designatedness' and 'multivalue'—ought to be understood just as versions of freestanding and ingredient content, respectively. These insights of Frege and Dummett have been applied and extended here by conjoining them with three further theoretical orientations. First, the substitutional apparatus that induces the distinction between levels of content is applied *analytically*, or in a top-down categorial direction, rather than *synthetically*, or in the bottom-up categorial direction of explanation that has dominated logic and semantics since Frege. Second, where standard treatments focus exclusively on the pragmatic goodness of *asserting*, to generate a top-level notion of truth-value as assertional designatedness, attention has been drawn here as well to the pragmatic goodness of *inferring*, to assign inferential roles to sentences, on the basis of which ingredient contents can then be defined substitutionally. Dummett recommends a move like this in understanding multivalued logics, under the heading of shifting from concern with logical validity to concern with logical derivability, from formally good claims to formally good inferences.

Finally, and perhaps most important, the substitutional mechanism that relates designatedness to multivalues is applied to *contents* rather than to *forms*. It provides a general semantic structure answering to material commitments and endorsements and the proprieties they induce, not merely to formal ones. On the assertional side, assertional commitments generally are considered at the top level, not just formally validated theorems. On the inferential side, material-inferential commitments are considered at the top level, not just formally valid inferences. Furthermore, this shift from the formal to the material is extended down the substitutional hierarchy—not just from pragmatic significance to freestanding sentential semantic content, but from such content to the ingredient contents that matter for the behavior of sentences as components in compound sentences. So the sentential compounding devices that can be considered are extended from purely formal vocabulary such as the conditional to any materially contentful sentential context in which other sentences can appear embedded as components (embedded as components in the sense that they can be substituted for). The

account that emerges adds another piece to the puzzle concerning the relations between logic and semantics, a piece that dovetails with the semantically expressive view of the distinctive task of logical vocabulary.

The problem with which this discussion began is generated by the fact that concepts of *semantic content* must, to deserve that appellation, play a role in determining the *pragmatic significance* of producing an utterance or adopting a state that exhibits such content. However, direct assertional and inferential significance attaches only to sentences, and furthermore only to sentences appearing freestanding, that is, as asserted or as premise or conclusion of an inference. Introducing the notion of substitution provides a model of a sort of indirectly assertional or inferential significance that the subsentential occurrence of an expression can have. This sort of content can be associated with sentences occurring as significant components in other sentences, rather than freestanding in an assertional or inferential way. Once this sort of ingredient content has been introduced into one's semantic theory, however, it becomes available to be associated also with expressions that (unlike sentences) can occur *only* as parts of assertible sentences. So the notion of *substitution* and substitutional content—which have been used here to cash out the notion of the 'contribution' the occurrence of a sentential component makes to the freestanding content of compound sentences it appears in—makes available a route into the assignment of broadly assertional and inferential contents to expressions of *sub*sentential grammatical categories, such as singular terms and predicates, to which the concept of freestanding content does not apply. The rest of this chapter is devoted to considering the sort of semantic interpretant that can be substitutionally associated with such expressions.

A bonus arising out of this line of thought is that, while remaining entirely at the level of sentences, and while eschewing any appeal to notions of reference or representation (what Dummett calls "the semantic model of the name/bearer relation"), it has been shown how to make sense of the notion of extensionality of an embedding context. Dummett is wrong to say: "If the notion of reference were introduced in the first place simply as that of the semantic role of expressions of different kinds, without an appeal to the name/bearer relation as prototype, then, at the outset, we should have no inclination to distinguish intensional from extensional contexts, or to treat the former separately; on the contrary, there would be a natural presumption in favour of a uniform semantic treatment for all contexts."[13] In fact, two different (though related) senses have been specified in which a context may be called 'extensional'. One has to do with substitutional homogeneity, that is, the sufficiency of assimilations of sentences according to their freestanding role to serve as assimilations of sentences according to their ingredient role. The other has to do with the sufficiency of concepts of content extracted from the *assertional* substitutional hierarchy of freestanding and ingredient contents to do duty as concepts of content in the *inferential* substitutional

hierarchy of freestanding and ingredient contents. It is in this sense that an inference—thought of as a context in which sentences can appear as premises and conclusions—(and hence also the conditionals that makes that sort of inference propositionally explicit) can be called 'extensional'. The clues provided by these ways of conceiving extensionality will be exploited later in discussing the relations between *inferential* and *referential* semantics.

III. SUBSENTENTIAL EXPRESSIONS

1. Singular Terms

What conditions on the use of an expression are necessary and sufficient for it to be functioning as, or playing the role of, a singular term? What sort of expressive impoverishment is a language condemned to by not having anything playing that sort of role? The answers to these questions may seem straightforward, at least in the large. Singular terms are linguistic expressions that refer to, denote, or designate particular objects.[14] The point of having something playing this role in linguistic practice is to make it possible to talk about particular objects, which, together with their properties and relations, make up the world in which the practice is conducted.

The first of these claims may be accepted without accepting the order of explanation presupposed by the transition from the first claim to the second. To begin with, it may be questioned whether the concept *particular object* can be made intelligible without appeal to the concept *singular term*. Frege, for instance, implicitly denies this when in the *Grundlagen* he explains the ontological category of particular objects, to which he is concerned to argue numbers belong, in effect as comprising whatever can be referred to by using singular terms, to which linguistic category he argues numerals belong. Again, it may sensibly be doubted whether the concept of singular reference is itself sufficiently clear to serve as an unexplained explainer. If it is not, then the responses offered above provide not so much answers to the questions they address as recipes for turning a suitable theory of reference into such answers. Insofar as one is sanguine about the prospects for such a theory, this of course is no bad thing. But it is important to be clear about what such a theory must account for in order to be serviceable in this explanatory context.

It is not enough, for instance, to explain only *successful* reference. For put somewhat more carefully, the first answer forwarded above must be that singular terms are expressions that, in Quine's useful phrase, "*purport* to refer to just one object."[15] The qualification expressed in this slogan by the use of 'purport' has two different functions: to acknowledge the notorious possibility that a name or a definite description may *fail* in its referential bid, as 'the most rapidly converging sequence' (or 'the square root of 2', as opposed to 'the positive square root of 2') does, and to exclude *accidentally* singular

expressions such as 'natural satellite of the earth', which succeed at unique signification, though they do not profess it. Ruling out possibilities of failure would require either omniscience on the part of those speaking the language or unacceptable restrictions on the formation of definite descriptions from predicates.

Quine is suspicious of the full-blooded notions of representational purport implicit in intentional idioms, and the echoes in his phrase are a reminder of his desire to explain much of what they might be thought to explain by appeal to more austere linguistic analogs. For singular referential purport, in the sense he appeals to, need not be an intentional affair. As Quine is quick to point out, "Such talk of purport is only a picturesque way of alluding to the distinctive grammatical roles that singular . . . terms play in sentences." The real task is to specify this role. Explanatory ground is gained by appeal to the principle Quine states only in the presence of such an account. That story, however, would offer a direct answer to the question, What is a singular term? one that does not appeal to (but on the contrary can itself be used via Quine's principle to help explain) the dark and pregnant notion of referential or representational purport. It is such an account that the remainder of this chapter aims to provide.

A further reservation concerning the line of thought about singular terms just considered has to do with the part the concept of singular terms is envisaged as playing in an account of the use of a language as a whole. Of particular importance is the relative explanatory priority of the category of singular terms with respect to the category of sentences. Semantic theories typically do not treat expressions of all grammatical categories equally. It is not just that different sorts of semantic interpretants are assigned to sentences, say, than to singular terms. In addition, some of those assignments of interpretants are considered basic, while others are derived from them. These latter are expression kinds whose semantic interpretation proceeds by appeal to the semantic interpretation of other sorts of expressions. A familiar example is broadly Tarskian compositional theories, which appeal to primitive assignments (perhaps relative to an index, such as a context or a model) of particular objects to atomic singular terms and of sets of those objects to atomic predicates in order to generate interpretations for sentences compounded out of them (and, along the way, to compound terms and predicates).

A contrasting direction of explanation is exhibited by broadly Fregean functional-categorial grammars and their corresponding semantics. These are less restrictive, both syntactically and semantically, than the Tarskian ones, in that any categories can be chosen as basic, not just terms and predicates, and any sort of interpretants can be associated with items of those categories, not just objects and sets of objects. A general mechanism is provided whereby (an infinite number of) further grammatical categories can be derived from the basic ones, and categorially appropriate interpretants supplied for expressions of those categories.

Suppose, as is usual, that singular terms (T) and sentences (S) are chosen for the basic categories. Then the derived category of (single-place) predicates $(T \to S)$ is understood to consist of expressions that combine with a term to yield a sentence, as 'writes' combines with 'Frege' to give 'Frege writes'. Adverbs, such as 'carefully', are expressions that combine with predicates to produce further predicates. They are $((T \to S) \to (T \to S))$s, taking 'writes' into 'writes carefully', for instance.[16] Exactly corresponding to this infinite syntactic hierarchy of derived categories is a semantic one, which turns arbitrary assignments of kinds of semantic interpretant to expressions of the basic categories into assignments of kinds of interpretant to expressions of the derived ones.

The general rule is that the semantic interpretant of an item of derived category $(X \to Y)$ is a function taking arguments of the kind semantically associated with the category X into values of the kind semantically associated with the category Y. So if singular terms were associated with objects, and sentences with sets of possible worlds, then predicates would be assigned functions from objects to sets of possible worlds, and adverbs would be assigned functions from functions of that kind to functions of that kind. The functional mechanism is completely indifferent as to the interpretation of the primitive categories—singular terms could as well be associated with recognizability conditions, and sentences with assertibility conditions. It would then be settled automatically that quantifiers, as $((T \to S) \to S)$s, must be semantically interpreted by functions taking functions from recognizability to assertibility conditions into assertibility conditions.[17]

2. Subsentential Expressions and Projecting the Use of Novel Sentences

In these two schemes for deriving the interpretation of some categories from the interpretation of more basic ones, sentences appear either as a derived semantic category or as a basic category on a par with singular terms. But it has been argued (under the heading of the pragmatic, and therefore semantic, priority of the propositional) that sentences are more special than this—that expressions of other categories count as having semantic content at all only insofar as they contribute to the content of sentences. The pre-Kantian tradition took it for granted that the proper order of semantic explanation begins with a doctrine of concepts or terms, divided into singular and general, whose meaningfulness can be grasped independently of and prior to the meaningfulness of judgments. Appealing to this basic level of interpretation, a doctrine of judgments then explains the combination of concepts into judgments, and how the correctness of the resulting judgments depends on what is combined and how. Appealing to this derived interpretation of judgments, a doctrine of consequences finally explains the combination of judgments into inferences and how the correctness of inferences depends on what is combined and how. Kant rejects this. One of his

cardinal innovations is the claim that the fundamental unit of awareness or cognition, the minimum graspable, is the judgment. For him, interpretations of something as classified or classifier make sense only as remarks about its role in judgment. In the *Grundlagen* Frege follows this Kantian line in insisting that "only in the context of a proposition [*Satz*] does a name have any meaning."[18] Frege takes this position because it is only to the utterance of sentences that pragmatic force attaches, and the explanatory purpose of associating semantic content with expressions is to provide a systematic account of such force.

That is, a further presupposition of the direction of explanation embodied in the answers forwarded above is that by saying what some expression represents (or purports to represent), one has thereby said how it ought to be used. As those candidate answers acknowledge, the category of singular terms whose nature and utility is being inquired into here is a *semantic* category. Associating objects (concrete or abstract) with expressions amounts to *semantic* interpretation of the expressions only if that association figures in the right way in accounts of how it would be correct to *use* them. Semantic properties and relations of expressions are distinguished from other sorts by the role they play in explaining the circumstances under which it is correct to use those expressions to perform various speech acts, and the appropriate consequences of so using them. Syntactic theory is concerned only to formulate rules determining what expressions are well formed, that is, can appropriately be used to perform standard speech acts. So it may group 'something' and 'everyone' in with 'the longest sentence in the A edition of Kant's first *Kritik*' and 'Aristotle'. Semantic theory ought nonetheless to distinguish the first two expressions from genuine singular terms, in virtue of the very different sorts of contribution their occurrence makes to the pragmatic significance of an utterance in a particular context.

Since semantics must in this way answer to pragmatics, the category of sentences has a certain kind of explanatory priority over subsentential categories of expression, such as singular terms and predicates. For sentences are the kind of expression whose freestanding utterance (that is, whose utterance unembedded in the utterance of some larger expression containing it) has the pragmatic significance of performing a speech act. Declarative sentences are those whose utterance typically has the significance of an assertion, of making a claim. Accordingly, there is available a sort of answer to the question,

What are sentences, and why are there any?

that is not available for any subsentential expression—namely, sentences are expressions whose unembedded utterance performs a speech act such as making a claim, asking a question, or giving a command. Without expressions of this category there can be no speech acts of any kind, and hence no specifically linguistic practice.

From this point of view it is not obvious why there should be subsenten-

tial expressions at all. For they cannot have the same sort of fundamental pragmatic role to play that sentences do. As a result, they not only cannot have the same *sort* of semantic content that sentences do, they cannot even have semantic content in the same *sense* that sentences do. Sentences have pragmatic priority, in that they are the category of expressions whose use constitutes linguistic practice. Accordingly, it is sentences whose proper deployment must be determined, in context, by associating semantic interpretants with them.

From this perspective, it is necessary to ask a question more general than that of the subtitle of this chapter:

> What are subsentential expressions, and why are there any?

Given the pragmatic priority of sentences, why should other semantically significant categories be discerned at all? Sentences are assigned semantic contents as part of an explanation of what one is doing in asserting them, what one claims, what belief one avows thereby. But the utterance of an essentially subsentential expression, such as a singular term, is not the performance of this sort of speech act. It does not by itself make a move in the language game, does not alter the score of commitments and attitudes that it is appropriate for an audience to attribute to the speaker. Accordingly, such expressions cannot have semantic contents in the same sense in which sentences can. They can be taken to be semantically contentful only in a derivative sense, insofar as their occurrence as components in sentences contributes to the contents (in the basic, practice-relevant sense) of those sentences.

So if, with Davidson, one takes the semantic interpretation of linguistic expressions to be an aspect of the intentional interpretation of behavior—assigning truth conditions to sentences according to the beliefs they express, and assigning truth conditions to beliefs and desires so as to make possible the explanation and prediction of behavior as largely rational for one who has beliefs and desires with those contents—then one ought to follow him as well in taking the *only* constraint on an assignment of denotations to subsentential expressions to be that it makes the truth conditions come out right. That is, one ought not to take there to be some independent notion of primitive denotation for such expressions that constrains or even determines the assignment of truth conditions.[19] The Tarskian technical apparatus is indifferent to whether it is exploited philosophically in the compositional, bottom-up direction Tarski originally envisaged or in the decompositional, top-down direction Davidson recommends. If one starts with the interpretation of subsentential expressions, then the primacy of the category of sentences in the linguistic practice one aims ultimately to account for provides sufficient motivation for moving up, compositionally, to generate truth conditions. What needs explanation is not this move but the concept of primitive denotation that provides its starting point. By contrast, if, because of their

pragmatic priority, one begins rather with the semantic interpretation of sentences, what is the motivation for decomposing them so as to interpret subsentential expressions as well? Why recognize the semantically significant occurrence of expressions of any category other than sentences?

Frege begins one of his later essays with this response: "It is astonishing what language can do. With a few syllables it can express an incalculable number of thoughts, so that even a thought grasped by a human being for the very first time can be put into a form of words which will be understood by someone to whom the thought is entirely new. This would be impossible, were we not able to distinguish parts in the thought corresponding to parts of a sentence."[20] The ability to produce and understand an indefinite number of novel sentences is a striking and essential feature of linguistic practice. As Chomsky has since emphasized, such creativity is the rule rather than the exception. Almost every sentence uttered by an adult native speaker is being uttered for the first time—not just the first time for that speaker, but the first time in human history. This high proportion of sentential novelty appears in surveys of empirically recorded discourses and becomes evident on statistical grounds when one compares the number of sentences of, say, thirty or fewer words, with the number there has been time for English speakers to have uttered, even if we never did anything else.[21] "Please pass the salt" may get a lot of play, but it is exceptionally unlikely that a sentence chosen at random from this book, for instance, would ever have been inscribed or otherwise uttered elsewhere.

The point is often made that individual speakers in training are exposed to correct uses only of a relatively small finite number of sentences and must on that basis somehow acquire practical mastery, responsive and productive, of proprieties of practice governing an indefinitely larger number.[22] The need to explain the possibility of projecting proper uses for many sentences from those for a few is, however, not just a constraint on accounts of language learning by individuals. For what is of interest is not just how the trick (of acquiring practical linguistic competence) might be done, but equally what the trick consists in, what counts as doing it. As just remarked, the whole linguistic community, by the most diachronically inclusive standards of community membership, has only produced (as correct) or responded to (as correct) a set of sentences that is small relative to the set of sentences one who attributes to them a language is thereby obliged to take it they have somehow determined the correct uses for. The idea that there is a difference between correct and incorrect uses of sentences no one has yet used involves some sort of projection.

There are a number of ways in which the use of a smaller number of sentences might determine the use of a larger number. If the alternative populations of sentence uses one is seeking to project to are sufficiently constrained, a small sample may suffice to determine which population is being sampled—for instance if no two of the candidate populations have any

subsets larger than n members in common, then a sample of that size will suffice to distinguish them. On the syntactic side, Chomsky follows this strategy in proposing that the reason grammar is learnable is that the final grammars available to the human learner are so severely constrained that the 'evidence' provided to such a learner in the form of sample grammatical sentences will in this way pick one out. On the semantic side, if there were relatively few in some sense possible constellations of correct usage for indefinitely many sentences, then specification or practical mastery of the correct use of a relative handful of them might well determine (in theory or in practice) the use of the rest—as ethologists have taught us that the most elaborate canned behavioral routines can be triggered by the occurrence of a few ordinary events.

Still, before resorting to the heroic postulation of the sort of structure that could make projection comprehensible even in the absence of subsentential structure (where sentences are individuated the way numerals are, for instance), it would seem the better part of valor to follow Frege in taking seriously the fact that the sentences we are familiar with do, after all, have parts. A two-stage compositional strategy for the explanation of projection would take it that what is settled by proprieties of use governing the smaller, sample set of sentences, which is projected, is the correct use of the subsentential components into which they can be analyzed or decomposed. The correct use of these components is then to be understood as determining the correct use also of further combinations of them into novel sentences.[23] The linguistic community determines the correct use of some sentences, and thereby of the words they involve, and so determines the correct use of the rest of the sentences that can be expressed using those words. By learning to use a relatively small initial sample of sentences, the individual learns to use the words they involve and thereby can learn to use all the sentences that can be formed out of those words by recombining them.

The need to project a distinction between proper and improper use for novel sentences provides the broad outlines of an answer to the question, What are subsentential expressions for? or Why are there any subsentential expressions? But what *are* subsentential expressions, functionally? According to the two-stage explanatory scheme, there are two sorts of constraints on the correct use of subsentential expressions, corresponding to their decompositional and compositional roles respectively. Their correct use must be determined by the correct use of relatively small subset of the sentences in which they can appear as components, and their correct use must collectively determine the correct use of all the sentences in which they can appear as components. In the passage cited above, Frege points out that this semantic projection of what he calls "thoughts" depends upon the possibility of syntactically analyzing the sentential expressions of those thoughts into elements, which can then be recombined to form novel sentences, expressing novel thoughts.

The key to the solution Frege endorses is the notion of *substitution*. For the first, or decompositional stage, sentences are to be analyzed into subsentential components by being assimilated as substitutional variants of one another—that is, related by being substitutionally accessible one from another. Regarding two sentences as substitutional variants of one another is discerning in them applications of the same function, in Frege's sense. In the second, or recompositional stage, novel sentences (and their interpretations) are to be generated as applications of familiar functions to familiar substitutable expressions. Familiar sorts of substitutional variation of familiar classes of sentences result in a host of unfamiliar sentences. This substitutional clue to the nature of subsentential expressions and their interpretation is pursued in what follows.

IV. WHAT ARE SINGULAR TERMS?

1. Syntax: Substitution-Structural Roles

What are singular terms? The question has been posed from the point of view of someone who understands (or is prepared to pretend to understand) already what it is to use an expression as a sentence but who admits to puzzlement concerning the distinctive contribution made by the occurrence of singular terms in such sentences. One way to get into this situation (that pursued in Chapter 3) is to begin with a pragmatics, an account of the significance of some fundamental kinds of speech act. A line can then be drawn around the *linguistic* by insisting that for the acts in question to qualify as *speech* acts, the fundamental kinds must include *asserting*. A general pragmatic theory then specifies for each speech act the circumstances in which, according to the practices of the linguistic community, one counts as entitled or obliged to perform it, and what difference that performance makes to what various interlocutors (the performers included) are thereby entitled or obliged to do. Assertional performances (and thereby specifically linguistic practices) are in turn picked out by *inferential* articulation— namely by the way in which the pragmatic circumstances and consequences of acts of asserting depend upon the inferential relations of ground and consequent among sentences. The category of sentences is then defined as comprising the expressions whose (freestanding or unembedded) utterance standardly has the significance of performing a speech act of one of the fundamental kinds. A pair of sentences[24] may be said to have the same pragmatic potential if across the whole variety of possible contexts their utterances would be speech acts with the same pragmatic significance (Fregean force).

Frege's notion of *substitution* can then be employed again to define subsentential categories of linguistic expression. Two subsentential expressions belong to the same syntactic or grammatical category just in case no well-

formed sentence (expression that can be used to perform one of the fundamental kinds of speech act) in which the one occurs can be turned into something that is not a sentence merely by substituting the other for it. Two subsentential expressions of the same grammatical category share a semantic content just in case substituting one for the other preserves the pragmatic potential of the sentences in which they occur. (Where sentences can occur embedded in other sentences, of course they too can be assigned semantic contents as well as pragmatic significances and potentials.) Then the intersubstitution of cocontentful subsentential expressions can be required to preserve the semantic contents of the sentences (and other expressions) they occur in, according to the structure laid out in the first two sections of this chapter. In this way, the notion of substitution allows both syntactic and semantic equivalence relations among expressions to be defined, beginning only with an account of force or pragmatic significance. The relations differ only in the substitutional invariants: expressions assimilated accordingly as well-formedness is preserved by intersubstitution share a syntactic category; those assimilated accordingly as pragmatic potential is preserved share a semantic content.

There are three sorts of roles that expression kinds can play with respect to this substitutional machinery. An expression can be substituted *for*, replacing or replaced by another expression, as a component of a compound expression. An expression can be substituted *in*, as compound expressions in which component expressions (which can be substituted for) occur. Finally, there is the substitutional frame or remainder: what is common to two substituted-in expressions that are substitutional variants of each other (corresponding to different substituted-for expressions). '$q \rightarrow r$' results from '$p \rightarrow r$', by substituting 'p' for 'q'. (In this example the expressions that are substituted in, '$p \rightarrow r$' and '$q \rightarrow r$', are sentences, and so are the expressions substituted for, 'p' and 'q'.)[25] The substitutional frame that is common to the two substitutional variants may be indicated by '$\alpha \rightarrow r$', in which 'α' marks a place where an appropriate substituted-for expression would appear.

Being substituted in, substituted for, or a substitutional frame are the *substitution-structural roles* that (sets of) expressions can play. The relation *being a substitutional variant of* obtains between substituted-in expressions, which must accordingly already have been discerned. Substitutional variation is indexed by pairs of expressions that are (substituting and) substituted for, which accordingly also must be antecedently distinguishable.[26] Substitution frames, by contrast, are not raw materials of the substitution process; they are its products. To discern the occurrence of a substitution frame—for instance '$\alpha \rightarrow r$' in '$p \rightarrow r$'—is to conceive of '$p \rightarrow r$' as paired with the set of all of its substitutional variants, such as '$q \rightarrow r$'. These are available only after a substitution relation has been instituted. For this reason, being substituted for and being substituted in may be said to be *basic* substitution-structural roles, while being a substitution frame is a (substitutionally) *derived* substitution-structural role.

Frege was the first to use distinctions such as these to characterize the roles of singular terms and predicates. Frege's idea is that predicates are the substitutional sentence frames formed when singular terms are substituted for in sentences.[27] One of the features that Strawson finds important in distinguishing singular terms from predicates follows immediately from this characterization in terms of substitution-structural role: namely that predicates do, and singular terms do not, have argument places and fixed adicities.[28] But it is clear that playing the substitution-structural roles of substituted for and frame with respect to substitutions in sentences is not by itself sufficient to permit the identification of expressions as singular terms and predicates, respectively. For, as in the schematic example of the previous paragraph, what is substituted for may be sentences, rather than singular terms, and the frames exhibited by substitutionally variant (sets of) sentences thereby become sentential connectives or operators, rather than predicates.[29] It will be seen, though, that the substitution-structural roles do provide important necessary conditions for being singular terms and predicates.

Why not think of predicates also as expressions that can be substituted for? If "Kant admired Rousseau" has "Rousseau admired Rousseau" as a substitutional variant when the category substituted for is terms, does it not also have "Kant was more punctual than Rousseau" as a substitutional variant when the category substituted for is predicates? Indeed, does not talk about predicates as a category of expression presuppose the possibility of such replacement of one predicate by another, given the substitutional definition of 'category' offered above? It does. Though either notion can be used to assimilate expressions accordingly as it preserves well-formedness of sentences, however, it is important to distinguish between *substituting* one expression (of a basic substitution-structural kind) for another and *replacing* one sentence frame with another. These differences are discussed in detail in Section V below. A few brief observations suffice here.

To begin with, it should not be forgotten that the frames on which the latter sort of replacement operates must themselves be understood as products of the former sort of substitution operation. Whatever items play the substitutionally derivative roles, for instance of sentence frames, can be counted as *expressions* only in an extended sense. They are more like *patterns* discernible in sentential expressions, or sets of such expressions, than like parts of them. A sentence frame is not a prior constituent of a sentence (though its occurrence may be marked orthographically that way) but a product of analyzing it, in particular by assimilating it to other sentences related to it as substitutional variants, when one or more of its actual constituents is substituted for. As a result, relative to such an analysis a sentence can exhibit many occurrences of expressions that can be substituted for, but only one frame resulting from such substitutions. A further difference, which is also a consequence of the substitutionally derivative status of sentence frames, is that replacing sentence frames, or more generally discerning substitutional variants in the second, wider sense, which involves replacement

of derived categories, requires matching argument places and keeping track of cross-referencing among them.[30] This has no analog in substitution for expressions of substitutionally basic categories. Thus although replacement of derivative expressions is sufficiently like substitution for basic expressions to define syntactic equivalence classes of expressions, they differ in ways that will later be seen to be important.

2. Semantics: Substitution-Inferential Significances

Raising the issue of the inferential significance of the occurrence in a sentence of some kind of subsentential expression shifts concern from the syntactic consequences of substitutional relations to their specifically semantic significance. Syntactic substitutional categories are defined by specifying which substitutions preserve sentencehood—where being a sentence is understood as having a pragmatic significance of its own, in that its freestanding utterance standardly counts as performing a basic speech act, paradigmatically making an assertion (overtly and explicitly acknowledging a doxastic commitment). Semantic substitutional contents can be defined by specifying which substitutions preserve the basic feature or features of sentences, in terms of which the pragmatic theory explains the proprieties of their use—namely the significance of the various speech acts they can be used to perform. This might be truth, justification, assertional commitment, or entitlement to such commitment (or whatever), as discussed in Sections I and II.[31]

Inferences that relate substitutionally variant substituted-in sentences as premise and conclusion, whether or not their goodness consists in the preservation of some semantically relevant whatsit, are called *substitution inferences.* An example is the inference from

> Benjamin Franklin invented bifocals

to

> The first postmaster general of the United States invented bifocals.

The premised sentence is substituted *in,* and a singular term is substituted *for,* to yield the conclusion. Because Benjamin Franklin was the first postmaster general of the United States, the inference from the premise to its substitutional variant is truth preserving: in the appropriate context, commitment to the premise involves commitment to the conclusion.

The substitution inference above *materially involves* the particular singular terms that occur (and are substituted for) in it. The particular predicate is not materially involved. For it is possible to replace that predicate with others without affecting the correctness (in this case, status-preservingness) of the inference. Thus if "α invented bifocals" is replaced by "α walked," the substitution inference from

Benjamin Franklin walked

to

The first postmaster general of the United States walked

will be correct under the same assumptions as the original. The predicates of concern are complex predicates, not simple ones. As such, they cannot be substituted for in a strict sense—they are substitution frames, playing a derived substitution-structural role. But a suitable notion of replacement of one sentence frame by another can be defined in terms of substitution for expressions playing basic substitution-structural roles.

The idea of replacing substitutional frames permits, for instance, substitution instances quantified over in "Anyone who admires someone admires himself," such as

Rousseau admires Montaigne, and Rousseau admires Rousseau

to appear as *frame-variants* of

Rousseau writes about Montaigne, and Rousseau writes about Rousseau,

when "α admires β, and α admires α" is replaced by "α writes about β, and α writes about α." The notion of substitution inference may be broadened to include inferences whose conclusion results from the premise upon replacement of a substitutional frame or pattern it exhibits by another. That is, the conclusions of inferences to be called 'substitution inferences' may be either frame-variants or strict substitutional variants of the premises (corresponding to basic and derived substitutional variation). Examples would be the inference from "Benjamin Franklin walked" to "Benjamin Franklin moved," and from "Benjamin Franklin walked" to "The inventor of bifocals walked," respectively.

The substitution inferences (in this broad sense) in which singular terms are materially involved differ in their formal structure from the substitution inferences in which predicates are materially involved. This difference provides another way of distinguishing the characteristic role of singular terms from that of other subsentential expressions, paradigmatically predicates. The point is noted by Strawson, who observes that predicates, but not singular terms, stand in "one-way inferential involvements." If the inference from "Benjamin Franklin walked" to "The inventor of bifocals walked" is a good one, then so is that from "The inventor of bifocals walked" to "Benjamin Franklin walked." Substitutions for singular terms yield reversible inferences. But it does not follow that the inference from "Benjamin Franklin moved" to "Benjamin Franklin walked" is good one, just because the inference from "Benjamin Franklin walked" to "Benjamin Franklin moved" is a good one. Replacements of predicates need not yield reversible inferences. Substitution inferences materially involving singular terms are de jure sym-

metric, while all predicates are materially involved in some asymmetric substitution inferences (though they may be involved in some symmetric ones as well).

One way to think about this difference is that where the goodness of a substitution inference is defined by its preservation of some semantically relevant whatsit, reflexivity and transitivity of those inferences is guaranteed by the nature of the preservation relation. The stuttering inference from p to p preserves any status that p might be accorded, while if the inference from p to q preserves that status, and that from q to r preserves it, then so must that from p to r. The symmetry of the relation, however, is assured neither by its status as an inferential relation nor by its holding accordingly as some status of the premise is preserved or transmitted to the conclusion.[32] Predicate substitution inferences may be asymmetric, while singular-term substitution inferences are always symmetric.

So singular terms are grouped into equivalence classes by the good substitution inferences in which they are materially involved, while predicates are grouped into reflexive, transitive, *a*symmetric structures or families. That is to say that some predicates are simply inferentially weaker than others, in the sense that everything that follows from the applicability of the weaker one follows also from the applicability of the stronger one, but not vice versa. The criteria or circumstances of appropriate application of '. . . walks' form a proper subset of those of '. . . moves'. Singular terms, by contrast, are not materially involved in substitution inferences whose conclusions are inferentially weaker than their premises.[33] To introduce a singular term into a language one must specify not only criteria of application but also criteria of identity, specifying which expressions are intersubstitutable with it.

Each member of such an inferential interchangeability equivalence class provides, symmetrically and indifferently, both sufficient conditions for the appropriate application and appropriate necessary consequences of application, for each of the other expressions in the class.[34] So, when the material substitution-inferential commitments that govern the use of singular terms are made explicit as the contents of assertional commitments, they take the form of identity claims. Identity locutions permit the expression of claims that have the significance of intersubstitution licenses. Weakening inferences, the one-way inferential involvements that collectively constitute the asymmetric substitutional significance of the occurrence of predicate expressions, are made assertionally explicit by the use of quantified conditionals. Thus "Benjamin Franklin is (=) the inventor of bifocals" ($t = t'$) and "Anything that walks, moves" $[(x)(Px \to Qx)]$.[35]

3. Simple Material Substitution-Inferential Commitments

The substitution inference from "The inventor of bifocals wrote about electricity" to "The first postmaster general of the United States wrote

about electricity" is a material inference. Since the inventor of bifocals is (=) the first postmaster general of the United States, it is a good inference (as is its converse). This last remark is worth unpacking a bit. Part of my associating the material content I do with the term 'the inventor of bifocals' consists in the commitment I undertake to the goodness of the substitution inferences that correspond to replacements of occurrences of that term by occurrences of 'the first postmaster general of the United States' (and vice versa). (This point is independent of the availability within the language of identity locutions permitting this substitution-inferential commitment to become explicit in the form of an assertional commitment.) That commitment has a general substitution-inferential significance, which is to say that the particular material inference endorsed above is correct as an instance of a general pattern. That same material-substitutional commitment regarding "the inventor of bifocals" and "the first postmaster general of the United States" governs also the propriety of the inference from "The inventor of bifocals was a printer" to "The first postmaster general of the United States was a printer," also that from "The inventor of bifocals spoke French" to "The first postmaster general of the United States spoke French," as well as a myriad of others. So one simple material substitution-inferential commitment regarding two expressions determines the correctness of a great many substitution inferences materially involving those expressions, across a great variety of substituted-in sentences and residual sentence frames, of which '. . . wrote about electricity' is only one example.[36]

Also, the substitution inferences to and from "The inventor of bifocals was a printer" are determined by all the simple material substitution-inferential commitments (SMSICs) that link the expression 'the inventor of bifocals' with another. Not all occurrences of those expressions, however, have their substitution-inferential significances determined in this way. For instance, that the inventor of bifocals is the first postmaster general of the United States does not settle the propriety of the substitution inference from

> The current postmaster general of the United States believes that the first postmaster general of the United States was a printer

to

> The current postmaster general of the United States believes that the inventor of bifocals was a printer.[37]

These observations motivate the discrimination of certain occurrences of an expression, in a syntactic sense of 'occurrence', as in addition semantically significant occurrences of it. A subsentential expression has a syntactic occurrence as a component of (is exhibited by) a sentence just in case it is replaceable by other expressions of its category (either in the original sense of being substituted for or in the secondhand sense appropriate to expressions

of substitutionally derived categories), saving sentencehood. (Syntactic categories are interreplaceability equivalence classes, since replacement is reversible and preservation of sentencehood symmetric.) For an occurrence of an expression in this syntactic sense to count also as having primary substitution-semantic occurrence in a sentence, the substitution inferences to and from that sentence, in which that expression is materially involved, must be governed (their proprieties determined) by the set of simple material substitution-inferential commitments that link that expression with another.[38]

How do SMSICs relating subsentential expressions settle the correctness of the substitution inferences in which the sentences exhibiting primary substitution-semantic occurrences of those expressions figure as premises and conclusions? According to a general pattern. A material substitution-inferential commitment regarding A and A' is a commitment to the effect that for any B such that AB is a sentence in which A has primary substitution-semantic occurrence, the inference from AB to $A'B$ is good. Likewise, a material substitution-inferential commitment regarding B and B' is a commitment to the effect that for any A such that AB is a sentence in which B has primary substitution-semantic occurrence, the inference from AB to AB' is good. Five points may be noted concerning this structure relating substitutional commitments to substitutional inferences.

First, *all* of the substitution inferences in which a sentence such as AB figures as premise or as conclusion are determined according to this pattern by all of the SMSICs dealing with expressions having primary substitution-semantic occurrences in AB (which might, but need not, be just A and B). Second, responsibility for those proprieties of substitution inferences to and from a sentence is apportioned between the various subsentential expressions having primary occurrences in it, with the SMSICs dealing with a particular expression responsible for the inferences in which that expression is materially involved. The content (determiner of material proprieties of inference) of each expression is represented by the set of SMSICs that relate it to other expressions. Only the collaboration of all of the SMSICs corresponding to subsentential expressions having primary occurrence in a sentence settles the correctness of the whole set of substitution inferences it appears in as premise or conclusion. Third, a consequence of this division of labor in the determination of the correctness of material inferences (assigning aspects of it to different sorts of expression) is that material-inferential roles are determined thereby for novel compounds of familiar components. So even if no one has ever encountered the sentence $A'B'$, the SMSICs cited above determine a commitment to the propriety of the inference from AB to $A'B'$. Other SMSICs already in place may in the same way license the inference from $A'B'$ to $A''B'$, and so on. Accumulating the content (what determines material proprieties of inference) to be associated with subsentential expressions in the form of substitutional commitments regarding pairs of expressions, then, permits the projection of material proprieties of substitu-

tion inference involving a potentially large set of novel sentences from the proprieties involving relatively few familiar ones. Fourth, on this model it is clear how to understand additions to or alterations of content. For when I discover or decide (what would be expressed explicitly in the claim) that the inventor of bifocals is the inventor of lightning rods and thereby undertake a new simple material substitution-inferential commitment, the substitution-inferential potentials both of sentences in which these expressions have primary occurrence and of others substitutionally linked to them are altered in determinate and predictable ways. Fifth, for the same reason, it is easy to understand what is involved in introducing new subsentential vocabulary, as expressing novel contents. Such vocabulary will make exactly the same sort of contribution to the strictly inferential contents of sentences that the old vocabulary does, as soon as its use has been tied to that of the old vocabulary by suitable SMSICs.

The criteria of adequacy responded to by these five points jointly constitute the *point* of discerning semantically significant subsentential structure, once the pragmatic, and so semantic, priority of sentences is acknowledged. Against the background of this sort of understanding of the semantically significant decomposition of sentences into their components, the formal difference between the material-substitutional commitments governing singular terms and those governing predicates becomes particularly striking. The SMSICs that determine the material-inferential significance of the occurrence of singular terms are symmetric: a commitment to the correctness of the inference that results from substituting A' for A is also a commitment to the correctness of the inference that results from substituting A for A'. The set of SMSICs that determine the material-inferential significance of the occurrence of any predicate, by contrast, include asymmetric ones. From this point of view, what is special about singular terms is that the simple material substitution-inferential commitments relating pairs of terms partition the set of terms into equivalence classes. This is what it is for it to be (particular) objects that singular terms purport-to-refer-to. An equivalence class of intersubstitutable terms stands for an object. It follows from the substitutional definition of the object-specifying equivalence classes of terms that it makes no sense to talk of languages in which there is just one singular term (*pace* 'the Absolute' as Bradley and Royce tried to use that expression), nor of objects that can in principle only be referred to in one way (by one term). The SMSICs that confer material-inferential content on predicates, by contrast, do not segregate those expressions into equivalence classes, and so do not confer a content that purports to pick out an object. The asymmetric structure conferred on the material contents of predicates is quite different.

There are, then, two fundamental sorts of substitution-inferential significance that the occurrence of expressions of various subsentential categories might have: symmetric and asymmetric. The claim so far is that it is a necessary condition for identifying some subsentential expression-kind as

predicates that expressions of that kind be materially involved in some asymmetric substitution inferences, while it is a necessary condition for identifying some subsentential expression-kind as singular terms that expressions of that kind be materially involved only in symmetric substitution inferences. These paired necessary semantic conditions distinguishing singular terms from predicates in terms of substitution-inferential significance (SIS) may be laid alongside the paired necessary syntactic conditions distinguishing singular terms from predicates in terms of substitution-structural role (SSR). The suggestion then is that these individually necessary conditions, symmetric SIS and substituted-for SSR, are jointly sufficient to characterize the use of a kind of expression that distinguishes it as playing the role of singular terms. In the rest of this work, the expression 'singular term' is used to signify expressions that play this dual syntactic and semantic substitutional role. It is to whatever expressions play this role that the argument is addressed.

V. WHY ARE THERE SINGULAR TERMS?

1. Four Alternative Subsentential Analyses

Here is an answer to the question, What are singular terms? They are expressions that are substituted for, and whose occurrence is symmetrically inferentially significant. The question, Why are there any singular terms? can now be put more sharply. Why should the expressions that are substituted for be restricted to symmetric inferential significance? What function does this arrangement serve? What expressive necessity enforces this particular combination of roles?

It is clear enough why the use of a substitutional scalpel to dissect sentential contents into subsentential components requires distinguishing expressions substituted for from substitutional frames. But why should any sort of subsentential expression have a symmetric SIS? And if some sort for some reason must, why should it be what is substituted for rather than the corresponding substitutional frames? Why should not both be symmetrically significant? The argument developed in the rest of this chapter is an attempt to answer these questions, and so in the specified sense to answer the question, Why are there singular terms?

What are the alternatives? They are structured by the previous pair of distinctions, between two sorts of substitution-structural syntactic role and between two sorts of substitution-inferential semantic significance. The possibilities are:

> (i) substituted for is symmetric;   substitutional frame is
>   symmetric
> (ii) substituted for is asymmetric;   substitutional frame is
>   symmetric

(iii) substituted for is asymmetric;   substitutional frame is
                                                    asymmetric
(iv) substituted for is symmetric;   substitutional frame is
                                                    asymmetric

The final arrangement (iv) is the one actualized in languages with singular terms. One way to ask why this combination of syntactic and semantic roles is favored is to ask what is wrong with the other ones. What rules out the combinations (i), (ii), and (iii)? What sort of consideration could? The strategy pursued here is to look at the constraints on the expressive power of a language that are imposed by each of those varieties of complex substitutional roles.

The first alternative is a good place to begin, for it is fairly easily eliminated from contention. The semantic point of discerning subsentential structure substitutionally is to codify an antecedent field of inferential proprieties concerning sentences by associating material contents with recombinable subsentential expressions so as to be able to derive those proprieties of inference and to project further ones, according to a general pattern of substitution-inferential significance of material-substitutional commitments. But the substituted-in sentences whose inferences are to be codified themselves stand in "one-way inferential involvements." The goodness of an inference may require that when the conclusion is substituted for the premise(s), some status (doxastic or assertional commitment, truth, and so on) is preserved. But the converse replacement need not preserve that status. Substitution inferences are not always reversible, saving correctness. Conclusions are often inferentially weaker than the premises from which they are inferred. A restriction to sentential contents conferrable by exclusively symmetrically valid material inferences is a restriction to sentential contents completely unrecognizable as such by us. But if both substituted-for expressions and the substitutional frames that are the patterns according to which they assimilate substituted-in sentences are significant only according to symmetric SMSICs, then asymmetric inferential relations involving substituted-in sentences can never be codified as substitution inferences materially involving subsentential expressions, and so licensed by the SMSICs regarding those expressions. Since the inferences to be codified include asymmetric ones, either the substituted-for expressions or the substitutional frames, or both, must be assigned asymmetric substitution-inferential significance.

The other two alternatives, (ii) and (iii), are alike in assigning the substituted-for expressions asymmetric substitution-inferential significance. If a good reason can be found for ruling out this combination of syntactic and semantic substitutional roles, then the employment of singular terms and their corresponding sentence frames will have been shown to be necessary. For if it can be shown that what is substituted for must have symmetric substitution-inferential significance, then since by the argument just offered the expressions playing some substitution-structural role must be asymmet-

ric (besides the substituted-in expressions), it will follow that the substitution frames must permit asymmetric substitution. Just this combination of roles has been put forward as characteristic of singular terms and predicates.

The first task addressed (in Section IV) was to answer the question, What are singular terms? The answer that has emerged is that they are expressions that on the syntactic side play the substitution-structural role of being substituted for, and on the semantic side have symmetric substitution-inferential significances. The second task is to answer the question, Why are there any singular terms? by presenting an explanation of why the inferential significance of the occurrence of expressions that are substituted for must be symmetric (and so segregate expressions materially into equivalence classes whose elements accordingly jointly purport to specify some one object). It takes the form of an argument that certain crucial sorts of expressive power would be lost in a language in which the significance of substituted-for expressions was permitted to be asymmetric.

2. The Argument

What is wrong with substituted-for expressions having asymmetric inferential significances? An asymmetric simple material substitution-inferential commitment linking substituted-for expressions a and b is a commitment to the goodness of all the inferences that are instances of a certain pattern. Where Pa is any sentence in which a has primary occurrence, the inference from Pa to Pb (the result of substituting b for a in Pa) is a good one, though perhaps its converse is not. The point of discerning primary occurrences of substituted-for expressions depends on these generalizations, for they provide the link that permits the projection of proprieties of substitution inference, based on associating particular substituted-fors with material contents in the form of determinate sets of simple substitution-inferential commitments relating their use to that of other substituted-fors. Whether the generalizations that animate asymmetrically significant substitutional commitments regarding substituted-fors make sense or not depends on the contents expressed by the sentences substituted in, and it is this fact that in the end turns out to mandate symmetric substitutional significances for what is substituted for. To see this, consider three ever more radical ways in which the generalizations associated with simple material substitution-inferential commitments might fail to obtain.

First, suppose there were sentences Qa, Qb, $Q'a$, $Q'b$, such that the inference from Qa to Qb is a good one, though the converse is not, and the inference from $Q'b$ to $Q'a$ is a good one, though the converse is not. Then although Qb results from Qa by substituting b for a, and correspondingly in the case of $Q'\alpha$, so that a and b have syntactic occurrences in these sentences, a and b cannot be admitted as having primary substitution-semantic occurrence in these contexts. For there is no simple material substitution-inferen-

tial commitment determining the substitution-inferential potential of those sentences. No symmetric SMSIC governs those inferences, since they are not reversible. But no asymmetric SMSIC governs them either, since such a generalization can at most license either the inference from Qa to Qb as a replacement of a by the inferentially weaker b, or the inference from $Q'b$ to $Q'a$ as a replacement of b by the inferentially weaker a, but not both. There need be nothing anomalous about such a situation. $Q\alpha$ and $Q'\alpha$ are inferentially complementary frames with respect to a and b. This precludes codifying substitution inferences involving those expressions in terms of SMSICs relating them, because the generalizations that would be appropriate to a and b with respect to $Q\alpha$ are different from and incompatible with those that would be appropriate with respect to $Q'\alpha$. The only cost of not being able to discern semantically primary occurrences of a and b in these contexts is that some of the good substitution inferences involving them are not captured by the material contents associated with a and b. This is no different than what happens in other cases where one ought not to discern primary substitution-semantic occurrence, for instance in "S believes that Qa."

For the next step, suppose that $Q\alpha$ and $Q'\alpha$ behave in this inferentially complementary fashion for every pair of substitutional variants in which they appear. That is, suppose there were a pair of predicates $Q\alpha$, $Q'\alpha$ such that for any substituted-for expressions x and y, if the inference from Qx to Qy is a good one, but not the converse, then the inference from $Q'x$ to $Q'y$ is not a good one. The presence of such a pair of predicates would block the possibility of the substitution-inferential generalization that would be required to give a substitutional commitment asymmetric significance, no matter what substituted-for expressions it involves. This situation would preclude discerning primary substitution-semantic occurrences of *any* substituted-for expression, in sentences exhibiting the substitution frames $Q\alpha$ and $Q'\alpha$. Again, the only cost is that certain proprieties of substitution inference, the ones that involve those frames, will not be projectable based only on the material contents associated with subsentential expressions, crystallized in sets of relational SMSICs.

Finally, in order to see how one might argue against admitting asymmetrically significant substituted-for expressions, strengthen the supposition yet again and consider what happens if there is a general recipe for producing, given any frame $Q\alpha$, a frame $Q'\alpha$ that is inferentially complementary to it, in the prior sense. That is, each $Q'\alpha$ is to be so constructed that whenever the inference from Qx to Qy is good, but not vice versa (intuitively, because y is inferentially weaker than x, the way 'mammal' is inferentially weaker than 'dog'), the inference from $Q'y$ to $Q'x$ is good, but not vice versa, for any substituted-for expressions x and y. Such a situation precludes discerning any primary substitution-semantic occurrences of any substituted-for expressions. For there are no syntactic occurrences of any substituted-for expressions whose substitution-inferential significance is correctly captured by an

asymmetric SMSIC (the symmetric ones are not currently at issue). For an asymmetric substitution-inferential commitment relating a to b governs inferential proprieties via the generalization that for any frame $P\alpha$, the inference from Pa to Pb is a good one, though not in general the converse.

Under the hypothesis being considered, no matter what particular instance $P\alpha$ is chosen, it is possible to construct or choose a complementary predicate $P'\alpha$ for which only the complementary pattern of substitution-inferential proprieties obtains. In the presence of a recipe for producing for arbitrary substitution frames other frames that are inferentially complementary to them, then, no proprieties of substitution inference can be captured by asymmetric SMSICs, and so no primary substitution-semantic occurrences of substituted-for expressions corresponding to them. The upshot of this line of thought, then, is that the existence of asymmetrically significant substituted-for expressions is incompatible with the presence in the language of expressive resources sufficient to produce, for arbitrary sentence frames, inferentially complementary ones. To explain why substituted-for subsentential expressions have symmetric substitution-inferential significances, which on the current understanding is to explain why there are singular terms, then, it will suffice to explain what sort of expressive impoverishment a language is condemned to if it eschews the locutions that would permit the general formation of inferentially complementary sentence frames.

When it has been seen that the particular constellation of syntactic and semantic roles characteristic of singular terms is necessitated by the presence in the language of vocabulary meeting this condition, it becomes urgent to see what locutions make possible the production of arbitrary inferentially complementary frames, and how dispensable the role they play in linguistic practice might be. What locutions have this power? Examples are not far to seek. The one to focus on is the *conditional*. Because conditionals make inferential commitments explicit as the contents of assertional commitments, inferentially weakening the antecedent of a conditional inferentially strengthens the conditional. "Wulf is a dog" is inferentially stronger than "Wulf is a mammal" because everything that is a consequence of the latter is a consequence of the former, but not conversely. But the conditional "If Wulf is a mammal, then Wulf is a vertebrate" is inferentially stronger than "If Wulf is a dog, then Wulf is a vertebrate." For instance, everything that combines with the first conditional to yield the conclusion that Wulf is a vertebrate also combines with the second to yield that conclusion, but not conversely. Again, endorsing all the inferences from sentences exhibiting the frame "α is a dog" to the corresponding "α is a mammal" does not involve commitment to the goodness of the inferences from sentences exhibiting the frame "If α is a dog, then α belongs to an anciently domesticated species" to those exhibiting the frame "If α is a mammal, then α belongs to an anciently domesticated species." Instances of the first conditional are true claims expressing correct inferences, while instances of its substitutional variant are

false conditionals expressing incorrect inferences. Quite generally, let Qa be a particular sentence in which the substituted-for expression a has primary occurrence, and Qb be a substitutional variant of it, and let r be some other sentence. Then $Qa \rightarrow r$ is a sentence in which a has primary occurrence, and the symbol $Q'\alpha$ may be introduced for the sentence frame associated with its occurrence, writing the conditional above as $Q'a$. If a is inferentially stronger than b, asymmetrically, then the inference from Qa to Qb is good, but not its converse (Wulf is a dog, so Wulf is a mammal).[39] But if that is so, then the inference from $Q'a$ to $Q'b$ cannot be good, for inferentially weakening the antecedent of a conditional inferentially strengthens the conditional.

This last formulation suggests another example. Inferentially weakening a claim inside a negation inferentially strengthens the compound negation. If the substitution inference from Qa to Qb is good but the converse not, then the substitution inference from $\sim Qa$ to $\sim Qb$ cannot be good. Embedding as a negated component, like embedding as the antecedent of a conditional, reverses inferential polarities. The conclusion is that any language containing a conditional or negation thereby has the expressive resources to formulate, given any sentence frame, a sentence frame that behaves inferentially in a complementary fashion, thereby ruling out the generalizations that would correspond to asymmetric simple material substitution-inferential commitments governing the expressions that are substituted for in producing such frames.

3. The Importance of Logical Sentential Operators

The conditional and negation are fundamental bits of logical vocabulary. Is it just a coincidence that it is logical sentence-compounding locutions that permit the systematic formation of inferentially inverting sentential contexts? The sentence q is inferentially weaker than the sentence p just in case everything that is a consequence of q is a consequence of p, but not vice versa (consequences are not preserved but pruned). It is an immediate consequence of this definition that inferentially weakening the premises of an inference can turn good inferences into bad ones. The defining job of the conditional is to codify inferences as claims (make it possible to express inferential commitments explicitly in the form of assertional commitments). It is essential to doing that job that embedded sentences that can play the role of premises and conclusions of inferences appear as components—antecedents and consequents—in the conditional. The contexts in which component sentences occur as antecedents accordingly must be inferentially inverting. Notice that this argument presupposes very little about the details of the use of the conditional involved. It is enough, for instance, if the conditional has the designated (semantic or pragmatic) status in case the inference it expresses preserves the designated status. As the defining job of the conditional is to codify inferences, that of negation is to codify incompatibilities. The negation of a claim is its inferentially minimal incompat-

ible—~p is what is entailed by everything materially incompatible with p.[40] These underlying incompatibilities induce a notion of inferential weakening: "Wulf is a dog" incompatibility-entails, and so is inferentially stronger than, "Wulf is a mammal" because everything incompatible with "Wulf is a mammal" is incompatible with "Wulf is a dog" but not vice versa (incompatibilities pruned, not preserved). It follows that incompatibility-inferentially weakening a negated claim incompatibility-inferentially strengthens the negation. "It is not the case that Wulf is a mammal" is incompatibility-inferentially stronger than "It is not the case that Wulf is a dog," just because "Wulf is a mammal" is incompatibility-inferentially weaker than "Wulf is a dog." Thus negation also enables the formation of arbitrary inferential complements. It was argued in Chapter 2 that what makes both conditionals and negation, so understood, specifically logical vocabulary is that the material inferences and material inference-inducing incompatibilities, of which they permit the assertionally explicit expression, play a central role in conferring material contents on prelogical sentences. It is a direct result of this defining semantically expressive function that they form semantically inverting contexts.

Since it is the availability of such contexts that rules out asymmetrically significant substituted-for expressions, it follows that a language can have either the expressive power that goes with logical vocabulary or asymmetrically substitution-inferentially significant substituted-for expressions, but not both. It is leaving room for the possibility of logical locutions that enforces the discrimination of singular terms (and as a consequence, of predicates) rather than some other sorts of subsentential expression. This conclusion would not be surprising if the logical vocabulary in question were that employed to make explicit the substitution-inferential relations in virtue of which singular terms and predicates can be assigned distinct material contents. For, as has already appeared, the symmetrically significant SMSICs associated with singular terms can be made assertionally explicit with the use of identity locutions, and the asymmetrically significant SMSICs associated with predicates can be made explicit with the use of quantificational locutions (together with sentential logical vocabulary). But like all logical locutions, the use of these presupposes, and so ought not be appealed to in trying to explain, the material contents of the nonlogical expressions that are explicitated. In short, the use of identity and quantificational locutions presupposes singular term and predicate use. So of course any language whose use is sufficient to confer on expressions the significance of such locutions must already have in play the symmetric and asymmetric SMSICs associated with nonlogical subsentential expressions, and the expressions whose use they governed will be singular terms and predicates respectively.

But these are not the logical locutions appealed to in the argument against asymmetrically significant substituted-for expressions. On the contrary, the only logical locutions required for that argument are those whose roles are

definable solely in terms of the behavior of *sentences*, before any sort of subsentential substitutional analysis has been undertaken. The argument does not depend on any particular features of the sentential contents that are available to begin with, determining the proprieties of material inference that provide the targets for substitutional codification in (implicit or explicit) SMSICs. All that matters is the availability of the expressive power of logical sentential connectives.

There is of course no absolute necessity that such vocabulary be available in a language. It would be a mistake to conclude from the true premise that something can be thought of as *propositionally contentful* only in virtue of its relation to proprieties of *inferential* practice, to the conclusion that such practice must be *logically* articulated. Such a move depends on the formalist error of assimilating all correctnesses of inference to logical correctness of inference, thereby denying the possibility of material, content-conferring inferential proprieties. Material proprieties of inference are antecedent to formal proprieties in the order of explanation, because to say that an inference is valid or good in virtue of its K form (for instance logical form) is just to say that it is a (materially) good inference, and it cannot be turned into one that is not good by replacing non-K (for instance nonlogical) vocabulary by (syntactically cocategorial) non-K vocabulary. There is nothing incoherent about a language or stage in the development of a language in which the only vocabulary in play is nonlogical.[41] But insofar as the material contents associated with substituted-for expressions are introduced and modified in units corresponding to asymmetrically significant SMSICs involving those expressions, the language containing them is not just in fact bereft of the expressive power of logical vocabulary, it is actually precluded from acquiring it (until and unless the offending subsentential semantic structure is reorganized). That is a reason not to have substituted-for expressions behave this way semantically, even in languages in which logical locutions have not yet been introduced.

Having to do without logical expressions would impoverish linguistic practice in fundamental ways. The use of any contentful sentence involves implicit commitment to the (material) correctness of the inference from the circumstances of appropriate application associated with that sentence to the consequences of such application. Introducing conditionals into a language permits these implicit, content-conferring, material-inferential commitments to be made explicit in the form of assertional commitments. This is important at the basic, purely sentential level of analysis for the same reason it becomes important later at the subsentential level, when identity and quantificational locutions can be introduced to make explicit the SMSICs that confer distinguishable material-inferential content on subsentential expressions. In each case, once made explicit in the form of claims, those content-conferring commitments are brought into the game of giving and asking for reasons. They become subject to explicit objection, for instance by

confrontation with materially incompatible assertions, and equally to explicit justification, for instance by citation of materially sufficient inferential grounds. The task of forming and nurturing the concepts we talk and think with is brought out of the dim twilight of what remains implicit in unquestioned practice into the daylight of what becomes explicit as controversial principle. Material contents, once made explicit, can be shaped collectively, as interlocutors in different situations, physically and doxastically, but in concert with their fellows, provide objections and evidence, claims and counterclaims, and explore possible consequences and ways of becoming entitled to assert them. Logic is the linguistic organ of semantic self-consciousness and self-control. The expressive resources provided by logical vocabulary make it possible to criticize, control, and improve our concepts. To give this up is to give up a lot.[42] Yet, it has been argued, it is a direct (if unobvious) consequence of leaving open the possibility of introducing such inferentially explicitating vocabulary that the subsentential expressions that are substituted for will be singular terms, and their corresponding sentence frames will be predicates, as judged by the symmetric and asymmetric forms of their respective substitution-inferential significances.[43]

VI. OBJECTIONS AND REPLIES

1. Are Singular Terms Symmetrically Substitution-Inferentially Significant Substituted-Fors?

Before pursuing further the significance of this result, it will be well to look a little more closely at the argument that has been offered for it. There are several sorts of objection it might elicit, which call for different sorts of clarification. The argument went like this. Singular terms and predicates were distinguished, as essentially subsentential expressions, by the coincidence of a particular syntactic substitutional role with a particular semantic substitutional role. Singular terms are substituted *for*, and the significance of an occurrence is determined by the content associated with it in the form of a set of *symmetric* simple material substitution inferential commitments linking that term to others. Predicates are substitutional sentence *frames*, and the significance of an occurrence is determined by the content associated with it, in the form of a set of *asymmetric* simple material substitution inferential commitments linking that predicate to others. So to ask why there are singular terms is to ask why these two sorts of substitutional roles coincide as they do.

The possibility that both what is substituted for and the resulting sentence frames might have symmetric substitutional significance is dispatched by the observation that the sentences displayed by this analysis as substitutional variants of one another stand in asymmetric substitution-inferential relations to one another. So the question, Why are there singular terms?

reduces to the question of why syntactically substituted-for expressions have symmetric rather than asymmetric substitution-inferential significances. For a substituted-for expression to have such a significance is for the contribution it makes to the substitution-inferential potential of sentences in which it occurs to be articulated as a set of SMSICs relating it to other expressions. An asymmetric SMSIC relating t to t' dictates that for any sentence Pt in which t has primary substitution-semantic occurrence, the inference from Pt to Pt' is a good one, though not conversely, where Pt' is a substitutional variant of Pt resulting from it by substituting t' for t. It was then pointed out that any language expressively rich enough to contain conditional or negating locutions—in short any language equipped with the conceptual resources of elementary sentential logic—precludes the existence of any such asymmetrically significant SMSICs governing substituted-for expressions. For those resources suffice to ensure that whenever, for some sentence frame $P\alpha$, the inference from Pt to Pt' is a good one but not conversely, there is another sentence frame $P'\alpha$ such that the inference from $P't'$ to $P't$ is a good one but not conversely. It follows that in such languages discerning asymmetrically significant substituted-for expressions would codify no inferences at all, permit no assignment of material content, in the form of a set of SMSICs, to any substituted-for expression. This way of combining syntactic and semantic substitutional roles is accordingly ruled out, and only that actually instantiated by singular terms and predicates remains. Thus the availability of the expressive capabilities of sentential logic dictates that subsentential substitutional analysis discern singular terms and predicates, rather than some other fundamental subsentential categories. It will be helpful to consider some different kinds of objection that might arise to this argument.

One potentially vulnerable premise is the definition of singular terms (and so of predicates) relied upon. The key final move in the argument, invoking the expressive power of logical locutions to form new sentences (and so new sentence frames) purports to show only why expressions that play the syntactic role of substituted-for must play the semantic role of being symmetrically substitution-inferentially significant. It might be denied that this result (for the moment treating the argument as successful) shows anything of importance about singular terms. For it might be denied that the syntactic and semantic substitutional roles that have been identified as characterizing this subsentential category of expression pick out its fundamental properties.

This complaint can be subdivided. It might be denied that the dual syntactic-semantic substitutional role attributed to singular terms is one they actually play—that is, that playing that dual role provides at least a necessary condition for expressions to be functioning as singular terms. This seems a difficult line to pursue. Singular terms are substituted for in sentences, generating derivative substitutional sentence frames. These two roles are clearly discriminable by the fact that the resulting frames have fixed possibilities of combination with substituted-for expressions (fixed adicities), always being

exhibited by sentences with the same number of substituted-fors,[44] while those latter expressions are subject to no analogous constraint. Of course, singular terms can also play the role of the expression that is substituted in. Thus there can also be expressions of the derivative category of singular term frames, such as "the father of α," which are substitutional patterns according to which singular terms can be assimilated.

This much is not special to substituted-for expressions. There is no reason "α wrote about ß" cannot be understood as a one-place sentence-frame frame, for instance that exhibited by sentence frames such as "α wrote about Rousseau" and "α wrote about Kant," as well as its playing the role of two-place sentence frame exhibited by any sentences exhibiting any of those one-place frames.[45] Sentences, in virtue of the direct pragmatic (paradigmatically assertional) significance of their use, are the original substituted-ins, the ones that get the whole substitution-analytic enterprise off the ground. Subsentential expressions are pragmatically significant, and so semantically contentful only in a substitutionally indirect sense (for instance, it is only via their substitutional relations with other expressions that they can be said to stand in inferential or incompatibility relations). Thus the fact that they can be substituted in, in no way qualifies the characterization of singular terms as substituted for. By contrast to the case of sentences, it is only because terms can play the latter syntactic substitutional role that they can play the former one.

Alternatively, one might argue that playing the semantic substitutional role of being governed by symmetrical SMSICs does not provide a necessary condition for a substituted-for expression to be a singular term. Here there are two possible cases. The claim might be that some terms are governed by *a*symmetrical SMSICs, or it might be rather that the use of some genuinely occurring singular terms is not governed by SMSICs at all.

For a first try at constructing singular terms that stand in one-way, or asymmetrical, inferential relations, one might consider definite descriptions of the following form: (the x)[Dx] and (the x)[$Dx\&Fx$]. This, however, will not provide a counterexample to the symmetry claim. To generate such a counterexample, it would have to be the case that in general either

(1) P((the x)[Dx]) entails P((the x)[$Dx\&Fx$]),
    and not always vice versa

or

(2) P((the x)[$Dx\&Fx$]) entails P((the x)[Dx]),
    and not always vice versa.

But (1) does not hold because an object may qualify as (the x)[Dx] without thereby qualifying as (the x)[$Dx\&Fx$], namely in those cases in which it is not also F. Description (2) does not hold because an object may qualify as (the x)[$Dx\&Fx$] without thereby qualifying as (the x)[Dx], namely in those cases

where, though only one thing is both D and F, more than one thing is D. These correspond to failures respectively of the descriptive and definiteness conditions on the applicability of definite descriptions.

2. Asymmetrically Substitution-Inferentially Related Singular Terms?

Seeing where this sort of prospective counterexample falls down suggests a much more serious prospect.[46] Consider the relation between pairs of terms such as 'Benjamin Franklin' and 'Benjamin Franklin, who was a printer'. Anyone who is committed to the claim:

Benjamin Franklin, who was a printer, invented lightning rods

is thereby committed to the corresponding claim:

Benjamin Franklin invented lightning rods,

but not everyone who is committed to the second claim is thereby committed to the first. In particular, anyone who has no opinion about whether Benjamin Franklin was a printer, or who denies that he was, can endorse the second without the first.

Such examples point out that there is a kind of term-forming operator that, given a term t, produces a compound term of the form t, who (or which) Φs, which systematically induces asymmetric substitution inferences relative to the first. The existence of such a recipe for producing terms from terms requires a genuine restriction in the scope of the claims made about terms so far. The example shows that there are some pairs of terms and some sentence frames such that substitution inferences relating those terms in those sentence frames are asymmetric. Acknowledging the existence of such cases, however, does not require relinquishing any claims on which the preceding argument relies.

As a preliminary, it may be noticed that the asymmetry pointed to here involves only relations to compound terms of a special form, and not terms generally—by contrast to predicates, *all* of which stand in one-way inferential involvements to other predicates that are *not* formed from them at all, never mind in a particular way. Again, the formation of terms of the form t, who Φs presupposes the existence of terms not of that form, which can appear in the ordinary predications of the form Φt, from which these special compound terms are formed (of course, t could itself be of the form t', who Ψs, but it cannot in principle be terms of this form all the way down). Thus the existence of terms that behave this way is parasitic on the existence of terms that do not.

But the point of central importance is that, as could be deduced already from the argument concerning recipes for generating logically compound predicates that reverse inferential polarity, it is *not* the case even for these

special pairs of terms that their substitution-inferential significance is governed by asymmetric SMSICs. To be so governed involves suitable substitutional variants standing in one-way inferential involvements for *all* predicates. And someone who is committed to the claim:

> It is *not* the case that Benjamin Franklin, who was a printer, invented lightning rods

is *not* thereby committed to the claim:

> It is *not* the case that Benjamin Franklin invented lightning rods.

Anyone who denies that Benjamin Franklin was a printer but who believes he did invent lightning rods can endorse the first claim but not the second. A similar point will hold for sentence frames formed from conditionals in which the terms in question appear in the antecedent. Thus pairs of terms of the form ⟨t, and t, who Φs,⟩ fail to be *systematically* asymmetrically substitution-inferentially significant in the way predicates are. So the symmetric and asymmetric *patterns* of substitution-inferential significance can still be appealed to in distinguishing the category of singular terms from that of predicates.

The restriction of asymmetric significance to some sentential contexts points up that these cases belong in a box with asymmetric inferences such as that from:

> The skinniest person in the room can't fit through the narrowest door

to

> The fattest person in the room can't fit through the narrowest door,

which are also not reversible.[47] These examples clearly turn on interactions between the predicates used to form definite descriptions and those involved in the sentence frame in which the description is embedded. Just by their nature, such asymmetries do not generalize across sentence frames generally and so have no systematic significance of the sort appealed to in the substitutional account of the difference between singular terms and predicates.[48]

3. Symmetrically Substitution-Inferentially Related Expressions That Are Not Singular Terms?

The second possible form of complaint about the necessity of the semantic condition of symmetrical substitutional significance stems from the observation of expressions that intuitively seem to be singular terms but that do not play a symmetric substitutional role. Consider occurrences of

expressions that function syntactically as singular terms, that is, that are intersubstitutable, saving sentential well-formedness, with substituted-fors that also have primary substitution-semantic occurrences.[49] The semantic significance of those occurrences could be of any of four different kinds. First, they might be primary substitution-semantic occurrences. In that case, which of the substitution inferences the expression is materially involved in are taken to be good or correct is governed by the SMSICs undertaken by the primary assessor of those inferences—typically, us (who are discussing them). These are the occurrences here treated as characteristic and fundamental.

Second, at the other end of the spectrum, there can be occurrences of what are syntactically singular terms that are semantically inert, for instance appearances in direct quotation, or that are unprojectable, such as the 'it' in "It is raining." The right thing to say about these seems to be that they are singular terms merely in a syntactic sense—that is, syntactically substitutable for full-blooded singular terms—but not functioning semantically as singular terms. The proposed substitutional definition of the dual functional role of singular terms explicitly leaves room for this sort of occurrence.

Third, in between primary substitution-semantic occurrences and semantically accidental occurrences, are what may be called secondary semantic occurrences. These have a systematic semantic significance for the propriety of inferences, and it is determined by a set of (symmetric) SMSICs. But they fail to qualify as primary substitution-semantic occurrences because the SMSICs relevant to the assessment of the propriety of inferences in which the expressions are materially involved are not the SMSICs associated with the one assessing those inferences. The central examples are occurrences in opaque positions in propositional attitude ascriptions. Thus my assessment of the propriety of the inference from "Carlyle believed (or acknowledged commitment to the claim) that Kant ascribed to each of us a duty to make ourselves perfect and others happy" to "Carlyle believed (or acknowledged commitment to the claim) that the author of *Dreams of a Spirit-Seer* ascribed to each of us a duty to make ourselves perfect and others happy" does not depend on whether, according to me, Kant is the author of *Dreams of a Spirit-Seer;* that is, it does not depend on the SMSICs that govern my primary uses of those expressions. But it is governed by some (symmetric) SMSICs— namely those I attribute to Carlyle, rather than those I undertake myself (that is, take to be true).

The definition of singular term in the full-blooded sense could be sharpened so as to acknowledge explicitly the possibility of secondary semantic occurrences. Since these are defined by government by (symmetric) SMSICs, such a clarification, qualification, or emendation is entirely in the spirit of the original proposal. The only kind of occurrence of expressions that syntactically qualify as singular terms that would be troublesome to the definition offered is if there should be a fourth sort of occurrence, intermediate between secondary semantic occurrences and semantically accidental ones.

These would be occurrences that were systematically significant for the goodness of inferences involving sentences they occur in, but not according to the pattern of government by SMSICs (symmetric or not). It is hard to know what to say about this possibility, in the absence of a specific candidate. The response to such a candidate ought to be to try to show that that role, like that of secondary semantic occurrences of singular terms, presupposes and is built on their primary role as governed by symmetric SMSICs. The possibility of a counterexample from this direction cannot simply be dismissed, though.

4. Does the Definition of Singular Terms Offer Sufficient Conditions?

A more promising line of attack might seem to be to focus on the sufficiency rather than the necessity of the characterization—to accept that singular terms do play the dual syntactic and semantic substitutional role attributed to them, but to deny that specifying these roles by itself suffices to pick out singular terms. Here the idea might be that there are other, genuinely central or essential roles of singular terms that are not derivable from those substitutional roles already mentioned, in that there could be expressions that satisfied the substitutional criteria but did not qualify as genuine singular terms through failure to play the genuinely critical roles. For instance, it might be noticed that the referential, rather than the inferential, role of singular terms has been ignored. This objection may be responded to in two ways—concessively and aggressively.

The concessive response would be to acknowledge that the characterization of singular terms omits features that may turn out to be essential, so that that characterization does not suffice to pick out singular terms. It may immediately be pointed out, however, that such an admission would not be very damaging to the argument presented. For that argument still goes through for an important class of expressions, call them 'ralugnis mrets', which contains singular terms as its most prominent proper subset. Even if this class should turn out to be substantially wider than that of singular terms, the question, Why are there any ralugnis mrets? would still be a fundamental one and would have to be addressed as the first part of an explanation of why there are singular terms. For that question would then assume the form "Why are there ralugnis mrets that also have the additional discriminating feature F (the hypothetical specific difference required to distinguish singular terms within the broader class of ralugnis mrets)?"

The aggressive response, though, is more interesting and potentially more satisfying. It is to attempt to derive the remaining features of singular-term usage from the dual substitutional characterization of them as symmetrically and substitution-inferentially significant substituted-for expressions. (Recall the suggestion earlier that having its substitution-inferential significance

determined exclusively by symmetric SMSICs can help provide an analysis of purporting to refer to or pick out one and only one object.) The explanatory strategy to which the present work (in particular, Part 2) is conceived as a contribution is to try to explain the referential and representational relations linguistic expressions stand in by appeal to the inferential, substitutional, and anaphoric relations they stand in. (These are three successive levels of analysis of material-sentential contents, each one of which presupposes the prior levels.) This is an ambitious undertaking, to be sure, and one this chapter does not make it possible fully to assess, since the substitutional relations discussed here provide only some of the necessary raw materials. Any difficulties of principle that arise in pursuing it will be forces pushing toward the concessive, rather than the aggressive, response to objections to the sufficiency of the characterization of singular terms.

But it would seem to be good Popperian methodology to adopt the strongest, most easily falsifiable, apparently sustainable commitment, to see where it falls down. In this case, concrete reason to be pessimistic about deriving the remaining representational properties would be concrete information about the missing feature F, which must be added to the substitutional characterization of singular terms to make it sufficient to pick them out within the wider class of ralugnis mrets. Detailed pursuit of the aggressive explanatory enterprise requires two further discussions: of anaphoric commitments that link unrepeatable tokenings and explain how anaphoric inheritance according to such links determines which SMSICs govern the use of those tokenings, and of how those links are made explicit in the form of *de re* ascriptions of propositional attitude. (*De re* ascriptions, such as "Carlyle believed *of* or *about* the author of *Dreams of a Spirit-Seer* that he ascribed to each of us a duty to make ourselves perfect and others happy" are the primary locutions with which we discuss representational relations. They permit us to talk about what people are talking about, and so provide a suitable explanatory target for an account of representational properties of linguistic expressions.)

5. Can the Same Argument Be Used to Show That Sentences or Predicates Must Have Symmetric Inferential Significances?

Another line of attack would be to object that the argument cannot be correct, for if it were, it would prove too much. For, it might be asked, why cannot one apply exactly the same considerations used to show that substituted-for *terms* must have symmetric inferential significances in order to show that *sentences* must have symmetric inferential significances? Sentences can be substituted for, as well as substituted in; they can appear as antecedents of conditionals or as negated; but their occurrences have an *a*symmetric inferential significance. Why does this not show that there must

be some defect in the argument purporting to show that there cannot be asymmetrically significant substituted-for expressions in a language containing logically expressive locutions? It does not because, unlike singular terms and predicates, sentences are not essentially subsentential expressions. Since sentences can occur freestanding, possessing pragmatic significance all on their own, their semantic content is not to be understood exclusively in terms of the contribution their occurrence makes to the inferential behavior of compound sentences in which they occur as significant components.

Sentences have a directly inferential content, since they can play the role of antecedent and consequent in inferences. Essentially subsentential expressions such as singular terms and predicates, by contrast, have an inferential content at all only in a substitutionally indirect sense. The inferentially articulated pragmatic significance of asserting a sentence is asymmetric already in the direct sense, since swapping premises for conclusions does not in general preserve the correctness of inferences. Considered just as expressions that are substituted for in other expressions—that is, just as embedded expressions—sentences need only be sorted into equivalence classes of various sorts, accordingly as they are symmetrically intersubstitutable preserving pragmatically important properties of the compound sentences in which they are embedded. This is exactly the procedure that was employed in Sections I and II, which investigated just such embedded occurrences of sentences and the sorts of subsentential sentential content they involve. Because the same sentences that stand in one-way inferential involvements as freestanding also are sorted into equivalence classes as substituted for, those equivalence classes can accordingly then be sorted into asymmetric families.

But this is entirely a consequence of their roles as premises and conclusions, and there can be no analog of this role for essentially subsentential expressions. As the discussion in the opening sections shows, the way to move from the sort of content associated with freestanding uses of sentences to that associated with embedded uses is broadly substitutional in nature. But it does not take the form of government by simple material, substitution-inferential commitments, which is all that is available for essentially subsententially occurring expressions. In Section II two sorts of sentential context of embedding for sentences were distinguished as inferentially 'extensional' in different senses, accordingly as intersubstitutabilities that preserve pragmatic statuses are determined by the freestanding assertional and inferential behavior of the component sentences, respectively. Both of these notions of extensional context depend essentially on how the expressions that occur in those contexts behave freestanding. Thus these notions do not apply to essentially subsentential expressions. Yet it is only by conflating them with the related but distinct notion of primary substitution-semantic occurrence (exclusive government by SMSICs), which applies to essentially subsentential expressions, that it could seem that the argument used to

forbid asymmetric significance of essentially subsentential substituted-for expressions would forbid it also to substituted-for sentences.

The objection was that the argument that essentially subsentential substituted-for expressions cannot have asymmetric substitution-inferential significance proves too much, for if good it would also show that sentences, which can also be substituted for, cannot have asymmetric substitution-inferential significance. That objection can be seen to fail because it ignores the further semantic resources available for expressions that are inferentially significant in a direct sense, as well as the substitutionally indirect one that the argument addresses. This immediately suggests a further objection that would not be vulnerable to this sort of response: to modify the previous objection by appealing to predicates, rather than sentences, to try to show that the argument for the symmetric significance of substituted-for expressions would prove too much. For predicates are essentially subsentential expressions, and the argument acknowledges that they do have asymmetric substitution-inferential significances, even in languages that have conditionals and negation in them. Although they are not expressions that are substituted for—indeed, as substitution frames, they are defined explicitly by their contrasting and complementary syntactic substitutional role—it must nonetheless be conceded by the argument that it makes sense to talk of predicates being replaced. For the semantic substitutional characterization of predicates as having the indirect inferential significance of their occurrence determined by families of asymmetric SMSICs presupposes such a notion. Why does not the argument from the expressive necessity of inferentially inverting sentential contexts to government of essentially subsentential substituted-for expressions by symmetric SMSICs go through equally well for essentially subsentential, merely replaceable, expressions—in particular, predicates? This is really another way of asking what role is played in the argument by the fact that the essentially subsentential expression involved is syntactically categorized as substituted for, and not as a substitution frame.

The argument was that for the simple material, substitution-inferential commitments associated with singular terms to govern not only occurrences in logically atomic sentences but also occurrences in embedded, inferentially inverting, logically articulated sentences, the significances associated with expressions by those commitments must be symmetric. Since asymmetric sentential substitution inferences are to be projected on the basis of the substitutional decomposition of sentences into compounds of essentially subsentential expressions, it follows that the other, derived, substitutional syntactic category—sentence frames—must be governed by SMSICs having asymmetric significance. It is now objected that there must be something amiss with this argument, for not only terms but predicates are recognizable when embedded in the antecedent of a conditional, or inside a negated sentence. So how is it possible for their occurrence to have asymmetric significance?

Here it is important to keep clearly in mind the distinction between substitutionally defined sentence frames and the predicate-letters or other expressions used as marks for them. Predicates, as substitution frames, are defined to begin with as equivalence classes of sentences, those that can be turned into one another by substitution for expressions of the substitutionally basic category.[50] '$P\alpha$', as a one-place sentence frame has been represented here, stands for such an equivalence class. Suppose it is the class of which "Carlyle wrote *Sartor Resartus*" and "Ruskin wrote *Sartor Resartus*" are members, because substitutional variants of the 'Carlyle' ↔ 'Ruskin' sort. But "If Carlyle wrote *Sartor Resartus*, then Carlyle had an ambiguous relationship with Hegel" is not obtainable as the result of substituting terms for terms in these sentences. It is not a member of the equivalence class denoted by '$P\alpha$'. In short, while the term 'Carlyle' has primary semantic occurrence in the conditional, as well as in its antecedent, the sentence frame "α wrote *Sartor Resartus*" has primary syntactic, as well as semantic, occurrence only in the antecedent (when freestanding), and not in the conditional as well. This is reflected in the fundamental difference between *substituting* for an expression and *replacing* it, as those words have been used here. What is replaced is always a frame associated with the whole sentence.[51] Thus although there is a sense in which a sentence can exhibit more than one sentence frame—for instance, "Carlyle wrote *Sartor Resartus*" exhibits not only "α wrote *Sartor Resartus*" but also "α wrote ß"—it is not in general possible to replace one of them, resulting in a sentence that still exhibits others of them. By contrast, a sentence may contain many occurrences of different substituted-for expressions, and substitution for one of them by and large results in sentences in which the rest still occur.[52] Thus frames have adicities, and substituted-fors do not.

It is a joint consequence of the requirement that primary semantic occurrences of either substituted-for expressions or of substitutional frames must have asymmetric significances, and the requirement that anything that has primary semantic occurrence both in freestanding sentences and embedded in inferentially inverting contexts such as negations and the antecedents of conditionals must be governed by symmetric SMSICs, that substituted-for and substitutional frames cannot both have primary semantic occurrence (the kind whose significance is governed by the SMSICs associated with the expression) in those embedded contexts. It is a consequence of the substitutional definition of sentence frames that they do not have primary occurrence in those contexts. This is not to say, of course, that no connection can be recognized between "X wrote *Sartor Resartus*" and "If Carlyle wrote *Sartor Resartus*, then Carlyle had an ambiguous relationship with Hegel" (if that were so, it would not be possible to state modus ponens without equivocation). On the contrary, the former is a frame exhibited by freestanding occurrences of the sentence that appears as antecedent in the latter conditional. Logical vocabulary produces inferentially inverting contexts because of its expressive role in making explicit inferential and incompatibility relations

among sentences. That same role guarantees that the contribution made to the use of a logical compound by the occurrence in it of component sentences is determined by the freestanding inferential role of those component sentences, which is what is being made explicit.[53] Thus the only contribution the predicate occurring in the antecedent need make to the projection of a semantic role for the conditional is to help project the semantic role of the antecedent as a freestanding, unembedded sentence. This much, however, is settled by discerning primary semantic occurrences of the predicate (replaceable ones whose significance is governed by SMSICs) only in the antecedent as a freestanding sentence. Thus in fact no semantic projectability is lost by refusing to discern in conditionals occurrences of the sentence frames exhibited by their antecedents.

6. Why Must It Be Possible to Substitute for Singular Terms in Logically Compound Sentences?

Offering this response to the objection naturally elicits a question as rejoinder: Why not treat occurrences of singular terms this way? That is, why not treat them as having their significance determined by the two-step process of examining first their contribution to the freestanding use of sentences and then the contribution of that freestanding use to the use of logically compound sentences in which they occur? The basic argument shows that both syntactic substitutional kinds of essentially subsentential expression cannot be taken to have primary semantic occurrence in conditionals; and it has been shown that, as defined, sentence frames do not. But what is the warrant for the asymmetry in treatment of the two syntactic substitutional sorts? Why see even terms as having primary semantic occurrence in conditionals, as well as in their antecedents? Asking this question makes it possible to highlight what is in some sense the heart of the difference between substituting for substitutionally basic subsentential expressions and replacing substitutionally derived sentence frames. There is an expressive reason for insisting on discerning primary semantic occurrences of singular terms in logically compound sentences, of which conditionals are paradigmatic, which has no analog for predicates.

To say that 'Carlyle' can play the role of substituted-for expression in "If Carlyle wrote *Sartor Resartus*, then Carlyle had an ambiguous relationship with Hegel" as well as in "Carlyle wrote *Sartor Resartus*" is to say that 'Ruskin', for instance, can be substituted for it in that context, to yield *conditional* substitutional variants, such as "If Ruskin wrote *Sartor Resartus*, then Ruskin had an ambiguous relationship with Hegel." These conditional substitutional variants define a conditional substitution frame: "If α wrote *Sartor Resartus*, then α had an ambiguous relationship with Hegel." The question accordingly becomes, What is the special virtue of discerning such frames? To yield the result in question, it must be a virtue that does not correspondingly attach to discerning the second-order conditional 'frames'

that would result from assimilating conditionals that were variants under replacement of first-order frames occurring in their antecedents—what might be symbolized in the example by "If $\Phi\alpha$, then α has an ambiguous relationship with Hegel."

The virtue in question is an expressive one, namely that conditional sentence frames must be discerned if the conditional locution that is used to make explicit material-inferential relations among sentences is to be able also to make explicit material substitution-inferential relations among sentence frames.[54] The material content associated with sentence frames, in virtue of which discerning expressions playing that essentially subsentential substitutional role contributes to the projectability of material-inferential contents for novel sentences, can be factored into simple material substitution-inferential commitments relating frames to other frames. When these content-conferring commitments remain implicit—that is, do not take the form of assertional commitments—they determine the significance of replacing one frame by another. When they are made explicit as the contents of logically articulated sentences, it is as quantified conditionals. It is a necessary condition of introducing quantificational locutions that one be able to discern the conditional sentence frames that are the quantificational substitution instances. In a language that has the logical expressive resources supplied by quantifiers, a typical predicate SMSIC might be made explicit as follows: (x)[if x is a whale, then x is a mammal]. Even in a relatively expressively impoverished language, one that still lacks the subsentential quantificational logical locution, if it has the sentential logical expressive resources supplied by the conditional, then the ability to project the semantic contents of novel material conditionals requires the ability to discriminate the substitutional variants that make up the conditional sentence frame "If α is a whale, then α is a mammal." That is, in a language containing a conditional, mastery of the material inferential content associated with the logically atomic frame "α is a whale" requires being able to assimilate conditional sentences into frame-equivalence classes accordingly as they are mutually accessible by substitution for expressions of the basic substitutional syntactic category. It is this mastery that may then be made explicit in the form of quantified conditionals.

The picture that emerges, then, is that substituted-fors and substitution for them (substitution mappings indexed by them) must be discerned in order to define frames in the first place, and again in defining replacement for them. Frames need to have material contents associated with them in order for singular terms to do so because projection requires the cooperation of both. Furthermore, logical frames are needed in order to make explicit the material contents associated with singular terms because identity locutions are such frames. *Conditional* (and negated)—that is, logically compound—frames are needed to make explicit the material substitution inferential and incompatibility contents of frames. Conditional sentence-frames—formed by assimilat-

ing conditionals according to accessibility relations defined by substituting for various basic essentially subsentential expressions—are required in order to codify explicitly the SMSICs that govern nonlogical first-order sentence frames.

There is no comparable necessity to be able to distinguish logically compound second-order frames by assimilating conditionals according to replacement of substitutionally derivative first-order frames in antecedents. Doing so is not required in order to codify and express explicitly the material contents either of singular terms or of first-order frames.[55] So essentially subsentential substituted-for expressions must be taken to have primary semantic occurrence in conditionals, as well as in the sentences embedded in those conditionals. For the corresponding conditional sentence frames articulate the substitutionally indirect, material-inferential contents governing the significance of primary semantic occurrences of nonlogical sentence frames. This reason does not analogously require that sentence frames be taken to have primary semantic occurrence in logically compound sentences.

Since it has been shown that discerning primary semantic occurrences of an expression both in the nonlogical sentences whose content is explicitated and in the inferentially inverting logical sentences that explicitate them requires that those occurrences be governed symmetrically by SMSICs, and that either frames or substituted-for expressions must be governed asymmetrically, it follows that frames and substituted-for expressions cannot both be taken to have primary semantic occurrence in logically compound sentences. Taken together, these arguments show why it is substituted-for expressions, and not the resulting substitutionally derivative sentence frames, that have the pattern of primary semantic occurrence across logically atomic and logically compound sentences that requires government by symmetric simple material substitution-inferential commitments. Since singular terms are essentially subsentential expressions that play the dual syntactic substitutional role of being substituted for, and the semantic substitutional role of having a symmetric substitution-inferential significance, this explains why there are singular terms.

7. Can the Substitutional Significance of the Occurrence of a Subsentential Expression Be Determined in Different Ways for Different Contexts?

One final objection should be considered. This stems from the thought that the inferential patterns associated with SMSICs are too rigid. Why should it be required that *all* of the primary occurrences of subsentential expressions have their significance determined by a SMSIC in the *same* way? If this requirement is relaxed, it seems that there is a way to evade the argument against the possibility of asymmetrically substitutional commitments governing expressions that are substituted for. Suppose, then, that the

terms a and b are linked, not by a symmetric relation of intersubstitutability that could be made explicit by an assertible substitution license in the form of an *identity* $a = b$ (defining an equivalence class), but by an asymmetric relation of *domination* that could be made explicit by an assertible substitution license in the form of an *inequality* $a \rangle b$. When this possibility was considered above, it was insisted that the significance of such a relation be that for *all* predicates or sentence frames $P\alpha$, if Pa then Pb, but not necessarily vice versa. The suggestion being considered now is that this requirement be relaxed.

Suppose that the logically atomic predicates are sorted into two classes, according to their *inferential polarity*. If $P\alpha$ has *positive* inferential polarity, then if Pa, then Pb, but not necessarily vice versa, just as before. By contrast, if $P\alpha$ has *negative* inferential polarity, then if Pb, then Pa, but not necessarily vice versa, under the same assumption that $a \rangle b$—that is, that a dominates b. It may turn out (but on the assumption being considered it need not) that all the logically atomic sentence frames have positive polarity. In any case, as the crucial step in the overall argument shows, there will be some logically compound sentence frames that have negative polarity. If $P\alpha$ has positive polarity, then $\sim P\alpha$ and $P\alpha \to r$ have negative polarity. To keep track, one might express all the logical compounds in disjunctive normal form and count the number of negations the term placeholder α is within the scope of. If it is odd, the polarity of its proximal logically atomic frame is reversed by the whole context; if even, that original polarity is retained.[56] Now it seems that it is possible to project substitution-inferential proprieties for logical compounds on the basis of an asymmetric relation of domination between expressions that are substituted for, by projecting the polarities of those compounds and applying the generalization that is appropriate for the polarity of the compound in each case. Although there is no *single* generalization specifying the significance of a certain asymmetric SMSIC of the sort demanded, there will be a *pair* of them, of just the sort demanded, one for each of the two polarities.[57]

If the suggestion that asymmetrically significant substitutional commitments governing the inferential significance of the occurrence of expressions that are substituted for can be accommodated in this way could be made to work, it would be devastating for the overall argument that has been offered here. It will not work, however, and something further can be learned from seeing why not. The general response is straightforward. Such a procedure will work as offered only if all predicates are one-place, including logically compound ones. If not, the polarity of a predicate can be different in different argument places. Thus $Pa \to Pb$ will have opposite polarities for the two argument places—one in the antecedent and the other in the consequent. The two-place predicate $P\alpha \to P\beta$ will not be sorted, then, into either polarity class by the procedure outlined above, and its inferential proprieties will accordingly not be determined by either pattern corresponding to the domination relation between a and b.

One might try to overcome this difficulty by assigning polarities not to *predicates* or *sentence frames* but to *argument places*. Then, assuming that the underlying frame $P\alpha$ has positive polarity, in $P\alpha \rightarrow P\beta$, the α position will have negative polarity, and the β position will have positive polarity. But this proposal will still not determine what should be said about the inferential relations between $Pa \rightarrow Pa$ and $Pb \rightarrow Pb$ in the case where $a \rangle b$: the first argument place has negative polarity and is being weakened by the substitution of b for a, so on that basis it should be the case that the overall claim is being strengthened inferentially by the substitution, while the second argument place has positive polarity and is being weakened by the substitution of b for a, so on that basis it should be the case that the overall claim is being weakened inferentially by the substitution.

Even if it is possible to fix up the proposal so as to deal with this difficulty, there is another that is decisive. One can take a sentence with two terms occurring in it at argument places of different polarities and form from it a *one-place* predicate or sentence frame: $P\alpha \rightarrow Q\alpha$, which can be represented as $R\alpha$. In this sentence frame, α has both positive and negative polarity. This is fatal to the scheme suggested for keeping track of polarities to permit projection of substitution-inferential proprieties in the face of asymmetrically significant substituted-for expressions.

Why not, then, exclude from the range of projection predicates that would require being assigned to each or neither of the polarities? Because these predicates are expressively essential. They are the ones that are required to codify the inferences involving predicates, for instance in the way that will eventually be made explicit by the use of quantifiers to bind conditional predicates ("Whatever walks, moves"). Thus the proposal cannot be carried through and poses no threat to the overall argument presented here.

VII. CONCLUSION

The title of this chapter asks a double question: What are singular terms, and why are there any? The strategy of the answer offered to the first query is to focus on substitution. The fundamental unit of language is the sentence, since it is by uttering freestanding sentences that speech acts are performed. Thus sentences are fundamental in the sense that it is coherent to interpret a community as using (its practices conferring content on) sentences but not subsentential expressions, while it is not coherent to interpret any community as using subsentential expressions but not sentences. But in fact there are good reasons why any community that uses sentences should also be expected to use subsentential expressions, indeed subsentential expressions of particular kinds.

The notion of substitution provides a route from the discrimination of the fundamental sentential expressions to the discrimination of essentially subsentential expressions. To carve up sentences substitutionally is to assimilate

them accordingly as occurrences of the same subsentential expressions are discerned in them. Such a decomposition is accomplished by a set of substitution transformations. The functional significance of discerning in a sentence an occurrence of one out of a set of expressions that can be substituted for is to treat the sentence as subject to a certain subclass of substitution transformations relating it to other, variant sentences. So the expressions that are substituting and substituted for can be used to index the transformations.[58] Two sentences are taken to exhibit the same substitutional sentence frame in case they are substitutional variants of one another—that is, are accessible one from the other by substitution transformations. These substitutional assimilations define two basic substitution-structural roles that essentially subsentential expression kinds could play. The first half of the answer to the first question, "What are singular terms?" is, then, that *syntactically,* singular terms play the substitution-structural role of being substituted *for,* while predicates play the substitution-structural role of sentence frames.

The second half of the answer to that question is that *semantically,* singular terms are distinguished by their *symmetric* substitution-inferential significance. Thus if a particular substitution transformation that corresponds to substituting one singular term for another preserves some semantically relevant sentential status (commitment, entitlement, truth, or whatever) when only primary occurrences are involved, no matter what the sentence frame, then the inverse transformation also preserves that status, regardless of frame. By contrast, every sentence frame is involved in weakening inferences where there is some other frame such that replacing primary occurrences of the first by the second always preserves the relevant sentential status, no matter what structure of substituted-for expressions is exhibited, while the converse replacement is not always status-preserving. Because the simple material substitution-inferential commitments that articulate the semantic content associated with singular terms are symmetric, their transitive closure partitions the set of singular terms into equivalence classes of intersubstitutable substituted-for expressions. It is in virtue of this defining character of their use that singular terms can be said to "purport to refer to just one object."

The full answer to the question, What are singular terms? is then that singular terms are substitutionally discriminated, essentially subsentential expressions that play a dual role. Syntactically they play the substitution-structural role of being substituted *for.* Semantically their primary occurrences have a *symmetric* substitution-inferential significance. Predicates, in contrast, are syntactically substitution-structural *frames,* and semantically their primary occurrences have an *asymmetric* substitution-inferential significance. This precise substitutional answer to the first question supplies a definite sense to the second one.

To ask why there are singular terms is to ask why expressions that are

substituted for (and so of the basic substitution-structural kind) should have their significance governed by symmetric commitments, while sentence frames (expressions of the derivative substitution-structural kind) should have their significance governed in addition by asymmetric commitments. The strategy pursued in answer to this question is to focus on the use of logical vocabulary to permit the explicit *expression*, as the content of sentences, of relations among sentences that are partly constitutive of their being contentful. To say that subsentential expressions are used by a community as substituted-fors and substitution-structural frames is to say that the contents conferred by the practices of the community on the sentences in which those expressions have primary occurrence are related systematically to one another in such a way that they can be exhibited as the products of contents associated with the subsentential expressions, according to a standard substitutional structure. The problem of why there are singular terms arises because that structure need not, for all that has just been said, assume the specific form that defines singular terms and predicates.

But suppose the condition is added that the sentences whose proper use must be codifiable in terms of the proper use of their subsentential components is to include (or be capable of being extended so as to include) not only logically atomic sentences but also sentences formed using the fundamental sentential logical vocabulary, paradigmatically conditionals and negation. This condition turns out to interact in intricate ways with the possibility of substitutional codification of sentential contents by subsentential ones— ways that when followed out can be seen to require just the combination of syntactic and semantic substitutional roles characteristic of singular terms and predicates. So the answer offered is that the existence of singular terms (and so of their complementary predicates) is the result of a dual expressive necessity. On the one hand, the material-inferential and material-incompatibility commitments regarding sentences must be implicitly substitutionally codifiable in terms of material-inferential and material-incompatibility commitments regarding the subsentential expressions that can be discerned within them or into which they can be analyzed, if the contents of novel sentences are to be projectable. On the other hand, those same commitments regarding sentences must be explicitly logically codifiable as the contents of assertional commitments, if the contents of nonlogical (as well as logical) sentences are to be available for public inspection, debate, and attempts at improvement. It is these two expressive demands, each intelligible entirely in terms of considerations arising already at the sentential level, that jointly give rise to the structure of symmetrically significant substituted-fors and asymmetrically significant substitution-structural sentence frames that defines the functional roles of singular terms and predicates.

The argument presented here may be called an *expressive deduction* of the necessity of basic subsentential structure taking the form of terms and predicates. A language must be taken to have expressions functioning as

singular terms if essentially subsentential structure is (substitutionally) discerned in it at all, and the language is expressively rich enough to contain fundamental sentential logical locutions, paradigmatically conditionals (which permit the assertionally explicit expression of material-inferential relations among nonlogical sentences) and negations (which permit the assertionally explicit expression of material-incompatibility relations among nonlogical sentences). The only way to combine the presence of logical vocabulary with a subsentential substitutional structure that does not take the term/predicate form is to preclude the formation of semantically significant logically compound sentence frames, by denying substituted-for expressions primary occurrence in logically compound sentences. The expressive cost of this restriction, however, is also substantial.

Unless sentence frames formed by substitutional assimilation of logically compound sentences are already available, it is not possible to introduce logical vocabulary (in the case of singular terms and predicates, identity and quantificational locutions) that will do for the commitments articulating the contents of subsentential expressions what the conditional and negation do for the commitments articulating the contents of sentences—namely express them explicitly in the form of assertional commitments. To put the point otherwise, the expressive power of sentential logical vocabulary derives in part from the interaction between a direct and a substitutionally indirect mode of making explicit the commitments that articulate sentential contents. The direct mode permits the formulation as the content of assertional commitments—of inferential commitments, for example. Without the expressive capacities provided here by conditionals, reasons could be demanded and debated for premises and conclusions, but not for the material inferences whose correctnesses are part and parcel of the content of those premises and conclusions. The substitutionally indirect way of making explicit the commitments that articulate sentential contents is by making explicit the commitments that articulate the contents associated with the subsentential expressions into which they can be analyzed, which commitments regarding subsentential expressions implicitly codify the same commitments regarding sentences that can also be made explicit directly. If a language has sentential logical vocabulary suitable to play both sorts of explicitating role, then its subsentential structure is obliged to take the specific form of singular terms and predicates.

Logical vocabulary has the expressive role of making *explicit*—in the form of logically compound, assertible sentential contents—the *implicit* material commitments in virtue of which logically atomic sentences have the contents that they do. Logic transforms semantic practices into principles. By providing the expressive tools permitting us to endorse in what we say what before we could endorse only in what we did, logic makes it possible for the development of the concepts by which we conceive our world and our plans (and so ourselves) to rise in part above the indistinct realm of mere tradition,

of evolution according to the results of the thoughtless jostling of the habitual and the fortuitous, and enter the comparatively well-lit discursive marketplace, where reasons are sought and proffered, and every endorsement is liable to being put on the scales and found wanting. The expressive deduction argues that subsentential structure takes the specific form of singular terms and predicates because only in that way can the full expressive benefits of substitutional subsentential analysis—codifying material correctnesses implicit in the use of sentences in material correctnesses implicit in the use of subsentential expressions—be combined with those afforded by the presence of full-blooded logical vocabulary of various sorts, performing its task of making explicit in claims what is implicit in the practical application of concepts.

In other words, languages have singular terms rather than some other kind of expression so that logic can help us talk and think in those languages about what we are doing, and why, when we talk and think in those languages. The full play of expressive power of even purely sentential logical vocabulary turns out to be incompatible with every sort of substitutional subsentential analysis except that in which essentially subsentential expressions playing the substitution-structural role of being substituted for have symmetric, substitution-inferential significances, and those playing the substitution-structural role of sentence frames have asymmetric, substitution-inferential significances. For to play its inference-explicitating role, the conditional, for instance, must form compound sentences whose antecedent substitution-position is inferentially inverting. Only symmetrically significant expressions can be substituted for, and so form sentence frames, in such a context. That is why in languages with conditionals, subsentential structure takes the form of singular terms and predicates.

In the opening paragraph of Section III it was pointed out that the principle that singular terms are used to talk about particular objects can be exploited according to two complementary directions of explanation. One might try to give an account of what particulars are, without using the concept *singular term*, and then proceed to define what it is to use an expression as a singular term by appeal to their relations to particulars. Or one might try to give an account of what singular terms are, without using the concept *particular*, and then proceed to define what it is for something to be a particular by appeal to their relations to expressions used as terms. (It should of course be admitted that in either case the talking about relation will require substantial explanation, though that explanation may have to look quite different depending on which explanatory strategy it is conceived as abetting.) The answer presented here to the question, What are singular terms? does not appeal to the concept of objects. So it provides just the sort of account required by the first stage of the second (Kant-Frege) strategy for explaining the concept of objects.

It is not the business of this chapter to pursue the later stages of that

direction of explanation, nor, therefore, to argue for its ultimate viability. But it is worth pointing out here that in the context of this order of explanation, to explain why there are singular terms is in an important sense to explain why there are objects—not why there is something (to talk about) rather than nothing (at all), but rather why what we talk about comes structured as propertied and related objects. "The limits of language (of that language which alone I understand) means the limits of my world."[59] To ask the question, Why are there singular terms? is one way of asking the question, Why are there objects? How odd that the answer to both should turn out to be: because it is so important to have something that means what *conditionals* mean!

Appendix I: From Substitutional Derivation of Categories to Functional Derivation of Categories

In functional-categorial grammars of the sort Lewis discusses in "General Semantics," one starts with some basic categories and defines derived categories by functions whose arguments and values are drawn from the basic categories, and from the derived categories already defined. Thus where singular terms (T) and sentences (S) are the basic categories, predicates, $(T \rightarrow S)$, are defined syntactically as functions taking (ordered sets of) terms as arguments and yielding sentences as values. Semantically, they are interpreted by functions taking (ordered sets of) whatever sort of semantic interpretant is assigned to terms as arguments and yielding as values whatever sort of semantic interpretant is assigned to sentences. Items of the basic categories play the roles of arguments and values of the functions associated with items of the derived categories. Items of derived categories of course play the roles of functions, but they also can serve as arguments and values of other functions.

These three roles correspond to substitution-structural roles. To play the role of value of a function is to be an expression that is substituted in. To play the role of argument of a function is to be an expression that is substituted for. To play the functional-categorially derivative role of a function is to be a substitutional frame—that is, a substitutionally derivative role. The triadic division of substitution-structural roles is accordingly orthogonal to that of functional-categorially basic and derived categories, in that derived categories, for instance, play all three substitution-structural roles. The two sorts of structure are nevertheless intimately related. The substitutional syntactic structure is a way of thinking about and constructing the functional-categorial syntactic structure (and thereby the corresponding semantic one).

The substitutional construction corresponding to a functional-categorial

hierarchy standardly generated by using sentences and terms as basic categories begins with sentences and substitutional transformations relating them. The sentences act as expressions to be substituted in, corresponding to the values of predicate functions. As to the expressions that are substituted for, they correspond to the indices of the substitution transformations. If they are, like the sentences substituted in, antecedently distinguishable, then the facts concerning which transformations remove occurrences of which terms, and which produce occurrences of other terms, can be used to index the substitution transformations, assigning to each a set of pairs of substituted and substituting terms. If only the transformations and not the term occurrences are given, then the sentences substituted in can be indexed by the substitution transformations that apply nontrivially to them, in order to determine what terms occur in them. Each transformation is then assigned a pair of sets of sentences—those it applies nontrivially to, and those it nontrivially results in. (A nontrivial substitution transformation is one that results in some sentence different from that to which it is applied. Intuitively, a transformation will apply nontrivially to a sentence only if one or more of the expressions it substitutes for occurs in the sentence.) In either case, predicates as sentence frames are defined as equivalence classes of substitutionally variant sentences.

Depending upon how the substitution transformations are conceived, it may take some special effort to see to it that a proper equivalence relation is defined from these substitutional accessibility relations. Thus if substitution for t needs to replace all occurrences of t, then it need not be the case that wherever substituting t' for t in s yields s', that substituting t for t' in s' yields s. This failure of immediate symmetry is evident if one substitutes 'Hegel' for 'Kant' in "Hegel wrote about Kant"—the converse substitution will not recover the original sentence from "Hegel wrote about Hegel." Similar phenomena afflict transitivity. These may be resolved either by defining substitutional accessibility in terms of the symmetric and transitive closures of these basic substitutional relations or by permitting partial substitutions in the base relation, at the cost of making each transformation one-many instead of one-one. For present purposes, it does not matter which route is adopted.

Defining sentence frames as equivalence classes of substituted-in expressions in this way suffices to determine their role as functions that apply to sets of substituted-in expressions. Applying the function to such an argument is just selecting some of the substitutionally variant sentences contained in the equivalence class, depending upon which substitution transformations apply nontrivially to it. This is another way of saying that selecting the right substitutionally variant instance depends on what substituted-for expressions occur in it. It may be noticed that at this stage, nothing corresponding to the order of arguments for a predicate function has been distinguished; no way has been supplied to tell the difference between "Brutus killed Caesar"

and "Caesar killed Brutus." These are recognizable as distinct members of the equivalence class that may be denominated "α killed β." But each of them is a result of applying that frame to the unordered set {'Brutus','Caesar'}. This is patently insufficiently discriminating for the purposes of codifying inferences; this point holds even before quantifiers, for it concerns the antecedently important implicit proprieties of inference that will be made explicit as the contents of assertions with the aid of quantifiers. So (to stick to the simplest sort of substitution inference) "Brutus killed Caesar" follows from "Brutus killed Caesar" but not from "Caesar killed Brutus." Finer discrimination is thus required.

This requirement should come as no surprise, if for the moment one thinks about substitution transformations in functional-categorial terms, rather than the other way around. Basic substitution transformations are ($T \rightarrow T$)s, functions substituting one term for another, which induce ($S \rightarrow S$)s. But it is seen in the body of the chapter that codification of simple inferences by substitution requires consideration not only of the inferential significance of substitution for terms but also of the inferential significance of replacement of predicates. This operation corresponds categorially to a function from predicates to predicates (hence inducing one from sentences to sentences), namely to a (($T \rightarrow S) \rightarrow (T \rightarrow S)$).

Understanding such an operation requires understanding predicates, ($T \rightarrow S$)s, not in their role as functions or frames but in their role as arguments and values of higher-order functions. (It is here that the connection Dummett rightly perceives between complex predicates and quantification emerges, for quantifiers are of course (($T \rightarrow S) \rightarrow S$)s.) For, in addition to their role as functions, the full-fledged derivative categories ($X \rightarrow Y$) of an unrestricted functional-categorial grammar can also serve as arguments for further, higher-order derived categories, such as (($X \rightarrow Y) \rightarrow Z$), and as values of such categories as ($Z \rightarrow (X \rightarrow Y)$). Talk of playing the role of argument and value is, in substitutional terms, talk of playing the substitutional roles of being substituted for and substituted in. The analog to being substituted for, for substitutionally and functionally derivative sorts of expressions, has been called here 'replacement'. Supposing that replacement can be defined, the role of sentence frames as values of functions—that is, as expressions that can themselves be substituted in (and therefore be understood as the result of broadly substitutional relations)—will follow straightforwardly. No new considerations are introduced by this further role, however, so it is not further considered here.

Defining replacement of one sentence frame by another is a more complex affair. This is the analog for substitutional frames of substituting one expression for another, which underlies the inference for instance from "Brutus killed Caesar" to "Brutus injured Caesar." (This is a propriety of practice that, in idioms expressively rich enough to count as logically articulated, can be made explicit in the principle $(x)(y)[(x$ killed $y) \rightarrow (x$ injured $y)]$.) It is with

respect to this operation that sentence frames must be individuated as finely as complex predicates, and not just as simple ones. Replacing "α killed β" by "α injured β" requires keeping the argument places straight. At this point structure is required that has no analog whatsoever at the level of simple substitution for basic expressions.

Understanding substitution for basic expressions requires that sentences be assimilated into equivalence classes corresponding to frames. Replacement of one substitutionally derivative frame by another requires not only those equivalence classes but a mapping from one to another that has special properties. In particular, there must be a bijection mapping the two equivalence classes onto each other so as to preserve the substitutional relations within each class. With respect to such a mapping, replacement of one predicate by another in a sentence exhibiting it then results in the element of the replacing equivalence class that is the image under that mapping of the first sentence. An example will make clear what is intended.

The set of sentences corresponding to "α killed β"—call it S—has the form {"Brutus killed Caesar," "Brutus killed Brutus,""Caesar killed Brutus,""Caesar killed Caesar," "The noblest Roman killed the conqueror of Gaul," . . .}. The set of sentences corresponding to "α injured β"—call it S'—has the form {"Brutus injured Caesar," "Brutus injured Brutus,""Caesar injured Brutus,""Caesar injured Caesar," "The noblest Roman injured the conqueror of Gaul," . . .}. The trouble is that these are unordered sets. Since at lower levels the occurrences of terms have been distinguished, it is already possible to specify that the result of replacing "α killed β" by "α injured β" in "Brutus killed Caesar" must be an element of S' in which the terms 'Brutus' and 'Caesar' both occur. So "Caesar injured Caesar" and "The noblest Roman injured the conqueror of Gaul" are ruled out. But nothing said so far makes it possible to choose between "Brutus injured Caesar" and "Caesar injured Brutus."

What is required is that the set S of sentences corresponding to "α killed β" be put in one-to-one correspondence with the set S' of sentences corresponding to "α injured β," so that $h($"Brutus killed Caesar"$)$ = "Brutus injured Caesar," $h($"Caesar killed Brutus"$)$ = "Caesar injured Brutus," and so on. Then to replace "α killed β" by "α injured β" in "Brutus killed Caesar," one simply applies the function h. The formal criterion of adequacy for a function h to be able to play this role is that:

> *If* $h(s)$ = s_1 and if Sub $\{s, s_2, t, t'\}$
>   (that is, s_2 results from s by substituting t' for t),
> *then* there must exist an s_3 such that Sub $\{s_1, s_3, t, t'\}$ and $h(s_2)$ = s_3.

In the example, since Sub {"Brutus killed Caesar," "The noblest Roman killed Caesar," "Brutus," "The noblest Roman"}, this means that if $h($"Brutus killed Caesar"$)$ is "Brutus injured Caesar," then there must be some sentence,

namely "The noblest Roman injured Caesar," such that it is the case both that h("The noblest Roman killed Caesar") is "The noblest Roman injured Caesar," and that Sub ("Brutus injured Caesar," "The noblest Roman injured Caesar," "Brutus," "The noblest Roman"). The notion of frame replacement makes sense only where such a mapping h has been defined from substitutional variants that are elements of one substitution-frame equivalence class to those that are elements of another.

Of course it is clear from this example that if there is one such mapping, there may well be others. For instance, h' could satisfy the condition if h'("Brutus killed Caesar") = "Caesar injured Brutus" instead of "Brutus injured Caesar." Selecting a substitution-structure preserving isomorphism h suffices to define the operation of predicate replacement that is employed in the semantic discussion of substitution inferences in the broad sense, which involves not only substituting for basic expressions but replacing substitutionally derived ones. This is all that is appealed to in the argument of this work.

To define the full functional-categorial hierarchy of derived categories, however, not only must frames be replaceable, but sentence-frame frames must be definable from them. This is part of playing the role of argument for higher-level functions. If the notion of predicate replacement is to be extended so as to be fully analogous to substitution for basic expressions (as the argument does not require), further structure still is needed. In particular, for this syntactic operation, one must be able to assimilate substituted-in expressions (sentences) accordingly as the same sentence-frame frame is exhibited—what "Kant admired Rousseau, and Kant wrote about Rousseau" has in common with "Kant lived longer than Rousseau, and Kant had a shorter name than Rousseau" and "Kant wrote more than Rousseau, and Kant wrote more carefully than Rousseau." Defining equivalence classes of sentences accessible from one another by replacing predicates with predicates requires more than the pairwise isomorphisms required to define replacement of predicates in the first place. It requires a set of such isomorphisms that link all the interreplaceable predicates into an equivalence class. This can be formulated as a requirement on a set of pairwise replacement-defining substitution-preserving isomorphisms. A structure $\langle R, H \rangle$ is a replaceability equivalence structure, in case:

1. $R = \{P_i/\text{each } P_i \text{ is a predicate of the same adicity } n\}$ and
2. $H = \{h(P_i, P_j)/P_i, P_j \text{ are elements of } R, \text{ and } h(P_i, P_j) \text{ is a substitution-preserving isomorphism between } P_i \text{ and } P_j\}$, and
3. H is reflexive, symmetric, and transitive over R, in that:
   (a) $h(P_i, P_i)$, an identity relation, is an element of H
   (b) $h(P_i, P_j)$ is the inverse of $h(P_j, P_i)$
   (c) $h(P_i, P_k)$ is the composition of $h(P_i, P_j)$ and $h(P_j, P_k)$.

Conditions (a), (b), and (c) need to be specially stipulated because of the potential multiplicity of isomorphisms mapping P_i onto P_j, and so qualifying

to be included in H as $h(P_i, P_j)$. This means that specifying such a structure amounts to picking one complex predicate from the set of those associated with each given simple predicate. Each structure $\langle R, H \rangle$ permits the definition of predicates as objects in the full-blooded sense of substituted-for expressions. Thus invariants of substituted-ins, across replacement within these classes R, as defined by the associated set of mappings H, permit the definition of genuine derived categories of higher-order sentence frames resulting from replacement of predicates. These same constructions of frames, by assimilation of substituted-ins, and of substituted-fors out of derived frames of lower levels, will be repeated at each level to generate the full hierarchy of functional categories.

Appendix II: Sentence Use Conferring the Status of Singular Terms on Subsentential Expressions—An Application

In Section IV of this chapter, an account is offered of what it is to use expressions as singular terms and predicates. That account is couched in terms of substitution-inferential relations among sentences. One consequence of the argument is accordingly that a theorist who analyzes some target system of linguistic practices by discerning the use of expressions as singular terms and predicates is obliged to show how that analysis is supported by appropriate features of the use of the sentences that contain them. The substitution-inferential structure described here puts substantial constraints on sentential practices, which must be satisfied if they are to be claimed to be sufficient to confer on subsentential components the semantic significance of singular terms and predicates. Where these constraints are not observed, erroneous conclusions will be drawn.

A prominent instance is what is often made of Quine's famous 'gavagai' example, from Chapter 2 of *Word and Object*. The example is forwarded as an argument for the thesis of the indeterminacy of translation. The significance of the example is typically understood to lie in its promise of a general recipe for generating alternate translation schemes by reindividuation. Specifically, wherever there is a 'straight' translation scheme, rendering a target-language sortal 'gavagai' by home-language sortal K, for instance 'rabbit', it is possible to produce a distinct and competing scheme that renders it instead by something that individuates more finely (or less, but it will suffice here to concentrate on the finer discriminations), for instance 'undetached rabbit part' or (temporal) 'rabbit stage'. The point is to be that since sentences are the smallest linguistic units that can be used to make a move in the language game, the evidence of linguistic practice directly constrains only the interpretation of sentences. This leaves considerable slack in how responsibility for the use of those sentences is indirectly apportioned between the subsentential linguistic units the theorist chooses to discern. The

considerations advanced in the body of this chapter do not provide reason to quarrel with the general conclusion but do give reason to quarrel with this example.

The idea is that what get construed in the straight translation scheme as predications addressed to singular terms governed by the sortal 'rabbit' are construed by the derived translation scheme as predications addressed to singular terms governed by the sortal 'undetached rabbit part'. Thus "There is a large rabbit" becomes something like "There is an undetached rabbit part of a large contiguous collection of such parts." As Quine indicates, "This is the same rabbit as that one" becomes "This undetached rabbit part belongs to the same contiguous collection of such parts as that one." The strategy is to take what appear to be sentences about rabbits, which predicate ordinary properties of them, as instead sentences about rabbit parts, which predicate of them gerrymandered properties involving the contiguous wholes they belong to.

From the point of view of the present analysis, the difficulty with such a derived scheme is that if the sentences as it construes them are to count as genuinely using some expressions as singular terms invoking parts, there must be some predications of them that do not address them solely through the wholes in which they appear. Not all the predicates that appear in the derived translated language can be of the sort that result from the recipe for retranslating what appear as predicates on the straight translation. That is, the use of sentences as translated must be governed by some symmetric simple material substitution-inferential commitments—which license substitutions of one part term for another—while insisting on a finer discrimination than that of their belonging to the same contiguous whole. These will be the symmetric SMSICs that could be (though they need not be) made explicit in the derivative translation by the use of genuine identity locutions such as "This is the same undetached rabbit part as that one." These will govern substitution inferences involving genuine predicates of undetached rabbit parts. Thus if the predicate P meant ". . . is a broken foot," a symmetric SMSIC governing terms a and b will license indifferently the inference from Pa to Pb and vice versa. It will be a commitment to the identity of the undetached rabbit parts a and b. It is one of the fundamental commitments of the present analysis that unless their use is such as properly to be governed by such symmetric SMSICs, a and b are not genuine singular terms. The point then is that derived translations of what are construed in the straight translation as predications applying to rabbits will not serve as contexts permitting genuine identity commitments regarding undetached parts.

This point can be seen intuitively, without appeal to the technical notion of simple material substitution inferential commitments. If 'gavagai' is to be construed as a genuine sortal, the language containing it must contain the apparatus needed to individuate the items it sorts. It must have a use for some notion that appears in the language as translated by the derived scheme

as 'same gavagai'. But the reindividuative strategy of construing apparent references to wholes as references to parts offers no assurance that the language being translated can be taken to have the apparatus needed to distinguish parts. Consider the suggestion that 'gavagai' means 'undetached organic molecule contained in a rabbit'. The natives presumably cannot identify and individuate molecules, and no amount of gerrymandering of their actual linguistic practice could construe it as already containing sufficient apparatus to do so.

The example of a reindividuative derivation of an alternative to the straight translation seems to work only because the theorist, working in a metalanguage rich enough to contain the full individuative and referential apparatus needed to make some expressions mean 'undetached rabbit part', or even 'undetached organic molecule contained in a rabbit', *stipulates* that a native expression is to be understood as used in the way such expressions are used. What the present considerations show is that this possibility does not ensure that the uses of the sentences attributed to the natives are themselves sufficient to *confer* that significance on the subsentential expressions they employ.[60] The result is a substantial asymmetry between the languages attributed to the natives on the straight construal and the derived construal. The straight construal attributes an autonomous language, in the sense that the use of the sentences attributed to the natives suffices by itself to make the subsentential expressions mean what they are taken by the theorist to mean. By contrast, the derived construal attributes a language that is not autonomous, in the sense that using the sentences in the way the natives are taken to is not enough to make the subsentential expressions mean what they are taken by the theorist to mean. Since no natural language could be like this—only an artificial language whose use is stipulated in some richer metalanguage could be—the straight construal is clearly theoretically preferable.

Thus the considerations advanced here concerning what it is for sentences to be related by substitution inferences in such a way that they count as containing occurrences of singular terms and corresponding predicates puts constraints on the theorist's discrimination of subsentential structure generating the use of sentences. These constraints are not satisfied by the proposed retranslations by reindividuation that would render what can be understood as 'rabbit' by 'undetached rabbit part' or 'rabbit-stage' (and dual considerations will apply to schemes that would move up to the less finely individuated 'rabbit-hood', rather than down to more finely individuated sortals).

To say this is to take issue with one (prominent) argumentative strategy, not with the indeterminacy thesis as such. For one thing, the present account begins with proprieties, including inferential proprieties, of the use of sentences, not with Quine's spare foundation of patterns of irritation of sensory surfaces (of theorist and native). Again, there are many other ways into indeterminacy not addressed by these conferral considerations, most notably

those Davidson develops involving the possibility of trading off attributions of beliefs and desires attributed to individuals, and the meanings attributed to their utterances. However, the fundamental point of this chapter has been to disagree with Quine's claim (offered at *Word and Object*, p. 53, as a lesson of the 'gavagai' example) that "terms and reference are local to our conceptual scheme," that "the very notion of term" is "provincial to our culture."

7

Anaphora: The Structure of Token Repeatables

Untwisting all the chains that tie
The hidden soul of harmony.
                    MILTON, L'Allegro

As fast as thou shalt wane, so fast thou grow'st
In one of thine from that which thou departest,
And that fresh blood which youngly thou bestow'st
Thou mayst call thine when thou from youth convertest
. . . . .
Let those whom nature hath not made for store,
Harsh, featureless and rude, barrenly perish
. . . . .
 She carved thee for her seal, and meant thereby
 Thou shouldst print more, nor let that copy die.
                    SHAKESPEARE, Sonnet 11

I. FREGE'S *GRUNDLAGEN* ACCOUNT OF PICKING OUT OBJECTS

1. Introduction

The first step in understanding why and in what sense claims represent or are about objects is to see what sort of conceptual content can be associated with the use of singular terms—the expressions that purport to refer to or represent objects—and correlatively with the predicates that denote the properties of objects and the relations they stand in. Only what has propositional (assertible or believable) content can serve as premise and conclusion—can both be given as a reason and have reasons given for it—and hence play a directly inferential role of the primary sort. But all sorts of conceptual content are essentially inferentially articulated. So the conceptual contents of singular terms and predicates must be understood in terms of their *in*directly inferential role—the contribution their occurrence makes to the inferential potential of sentences containing them. As Frege puts it: "We ought always to keep before our eyes a complete proposition. Only in a proposition have the words really a meaning . . . It is enough if the proposition taken as a whole has a sense; it is this that confers [*erhalten*] on its parts also their content."[1]

One of the principal technical conceptual debts the inferentialist semantic

tradition owes to Frege is the idea of using *substitution* to understand how the directly inferential articulation of what is expressed by sentences induces an indirectly inferential articulation of what is expressed by their semantically significant parts. The previous chapter investigated the substitutional species of inferential relations and showed how to understand singular terms and predicates in terms of the roles that expressions of those categories play in substitution inferences. It showed further the sort of independently characterizable expressive impoverishment to which a discursive structure is doomed unless its subsentential substitution-inferential structure takes the specific form of singular terms and predicates. Subsentential structure may be eschewed entirely, though the cost is substantial. For one then forgoes the expressive empowerment provided by the combinatory generation of novel interpretable sentences from familiar sentence-parts, which looms so large in our actual discursive practice. If semantically significant, essentially subsentential structure is discerned substitutionally, however, it can take a form other than that of singular terms and predicates only by relinquishing the full semantic explicitating expressive resources otherwise provided by sentential logical locutions, paradigmatically the conditional. This, it was claimed, is why there are singular terms (and so predicates, since the two categories come as a package). This same argument provides the ultimate reason why sententially atomic propositionally contentful claims are, or purport to be, about objects, and to represent those objects as having properties and standing in relations. The connection between singular-term usage and purported representation of objects can be filled in a bit by looking at how Frege, in the *Grundlagen*, understands the representation of objects in purely substitutional terms.[2]

2. Objects Are Given to Us by the Use of Singular Terms

One of Frege's concerns in the *Grundlagen* is to explain "how numbers are given to us."[3] In order to do that, he must consider the wider question of how particular objects are "given to us" cognitively. On the face of it, explaining what it is for our thought and talk to pick out or be directed at objects seems particularly difficult for the case of numbers, since, as he puts it, "we cannot have any ideas or intuitions of them."[4] Translated from the neo-Kantian idiom he is employing here, this means that the aboutness of numerical thought can be understood neither as derived from the supposedly more primitive aboutness of subjective mental pictures nor as a feature of the way in which thought about numbers is causally influenced by the numbers it is about. In fact in this context the abstractness of number is a philosophical boon because it requires Frege to address in its most general terms the question of what it is to pick out objects with our concepts—undistracted by such ultimately misleading features of some prominent special cases as the presence of mental pictures of or causal commerce with what is

thought and talked about. Instead, the role played by causal contact in the ability to pick out perceivable objects in thought and talk must be understood in terms of some more general conception of object-directedness.

Frege calls the grammatical category of expressions used to talk and think about objects "proper names": "The name of a thing is a proper name [*Eigenname*]."[5] This usage elides the distinction—of the first importance for Frege's project in the *Grundlagen*—between lexically simple singular terms, such as 'Frege', and definite descriptions formed from predicates and sortals, such as 'the author of the *Grundlagen*'. Frege's discussion focuses on the latter for two reasons: numerical expressions are formed in this way, and the definite article makes *explicit* the singular referential purport that is implicit in the use of other singular terms.

> We speak of "the number 1," where the definite article serves to class it as an object.[6]

> The definite article purports to refer to a definite object.[7]

The general question Frege is addressing is how expressions must be used for them to *succeed* as singular terms by referring to, picking out, or giving us a cognitive grip on definite objects—as "ways in which objects are given to us," ways of "arriving at determinate" objects, or "symbols signifying objects."[8] The issue of what it is to use an expression as a name of an object is ultimately a normative one; it is to be responded to by specifying proprieties of practice. Since the use of the definite article makes singular referential purport explicit, those proprieties can be brought out into the open by asking (in deontic scorekeeping terms) what sort of *commitment* is expressed by the use of the definite article, and what is required for *entitlement* to that commitment.

Frege insists that the issue of entitlement to singular referential purport is an important one. The use of the definite article stands in need of *justification* [*Rechtfertigung*].[9] The definite article is used in forming definite descriptions from predicates—what he calls "the definition of an object in terms of a concept under which it falls."[10] Frege is explicit about what is required for a justification of such a use of a definite article:

> If, however, we wished to use [a] concept for defining an object falling under it, it would, of course, be necessary first to show two distinct things:
>
> 1. that some object falls under this concept;
> 2. that only one object falls under it.[11]

These are the paired conditions, of existence and uniqueness, on which Russell later erected his theory of descriptions.

It may seem, that however it is with definite descriptions, explaining the object-directedness of thought (the way it puts us in touch with particular

objects our judgments are about) need involve attending only to the first of these. Showing what it means for atomic judgments to be analyzable in terms of the application of predicates would seem to suffice, for objects should then emerge as what the predicates are judged to apply *to* or be true *of*. In fact, however, understanding this 'to' and this 'of' requires mastery of the sort of practical issues of identity and individuation that are appealed to in the second condition. For in the absence of such considerations, one grasps only the use or application of whole sentences (what it is to take them to be true)—not yet what it is to apply them *to* something or take them to be true *of* something. That a judgment is directed toward an *object* is intelligible only in the context of practices of identifying objects as the same again, and individuating them as distinct.

3. Judgments Expressing Our Recognition of an Object as the Same Again Are Substitution Licenses

This is to say that the use of expressions as singular terms essentially involves, not only norms that could be made explicit as criteria of *application*, but also norms that could be made explicit as criteria of *identity*. Frege formulates this categorial point as the demand that "if we are to use a symbol *a* to signify [*bezeichnen*] an object, we must have a criterion for deciding in all cases whether *b* is the same as *a*, even if it is not always in our power to apply this criterion."[12] For what an expression makes cognitively available for us to "have a definite character" as an object our judgments are about, it is necessary that "it can be recognized again beyond doubt as the same, and can be distinguished from every other."[13] As indicated by the qualification "even if it is not always in our power to apply this criterion" in the previous passage, in spite of the epistemic flavor of "recognition" and "beyond doubt," the requirement is not that we in fact be able to apply the implicit criterion of identity or be infallible in our recognitions. It is just that a notion of *correctness* of identifications and discriminations must have been settled somehow. The normative status must have been instituted, even though any particular attitudes, attributions, and assessments may get it wrong.

The demand for an implicit criterion of identity associated with the use of a singular term is presupposed by the uniqueness condition on the application of definite descriptions, but it is not restricted to those singular terms in which the singular referential purport is marked overtly by the use of a definite article. In its absence, no sense could be made of the notion that terms (including those that are not definite descriptions) implicitly involve a specifically *singular* referential purport. An implicit criterion of identity provides the "authority to pick out [particulars] as self-subsistent objects that can be recognized as the same again [*selbständige, wiedererkennbare Gegenstände zu unterscheiden*]."[14] What does it mean for such authority or enti-

tlement (which could be made explicit in the form of a criterion of identity) to be in place? The key fact is that "objects too can change their properties without that preventing us from recognizing them as the same [*sie als dieselben anzuerkennen*]."[15] Recognizing an object as the same again is making a certain kind of judgment, what Frege calls a "recognition judgment." Thus "For every object there is one type of proposition which must have a sense, namely the recognition-statement."[16]

Indeed, the use Frege makes of the concept of a recognition judgment shows that he is committed to a much stronger claim. Not only is fixing the sense of recognition judgments *necessary* for entitlement to use an expression as singular term, it is *sufficient.* And once an expression has qualified as entitled to its singular referential purport, it is a way in which a determinate object can be picked out or given to us.[17]

> How, then, are numbers to be given to us, if we cannot have any ideas or intuitions of them? Since it is only in the context of a proposition that words have any meaning, our problem becomes this: To define the sense of a proposition in which a number word appears. That, obviously, leaves us still a very wide choice. But we have already settled that number words are to be understood as standing for self-subsistent objects. And that is enough to give us a class of propositions which must have a sense, namely those which express our recognition of a number as the same again . . .
>
> In doing this, we shall be giving a general criterion for the identity of numbers. When we have thus acquired a means of arriving at a determinate number and of recognizing it again as the same, we can assign it a number word as its proper name.[18]

That an expression is used as a singular term, and so has singular referential purport—that it is a way in which determinate objects can be made available to judgment ("arrived at," "given to us")—is a significance that performances can be accorded in the context of practices of keeping deontic score on special sorts of commitment and entitlement. It emerges from the passages quoted above that the central technical concept Frege employs to explain the commitments and entitlements that define singular term usage is that of *fixing the sense of a recognition claim.* The rest of this section is devoted to exploring how Frege uses this concept to elaborate his understanding of what it is to talk or think about particular objects.

Securing singular reference is for Frege "a matter of fixing the content of a recognition-judgment [*Wiedererkennungsurtheils*]."[19] Recognition judgments have the form of *identity* claims. Identity claims express recognition of an object as "the same again" when given or referred to in two different ways. To establish reference to a particular object by a given expression, one must settle what would make true or false various identities in which that expression occurs (even if one is not in a position to tell of each such identity

whether it is in fact true or false). Since the singular referential purport of terms amounts to claiming that recognition judgments involving those terms have a definite sense,[20] taking it that the significance of asserting an identity involving a term has been settled is treating the term as referring to or picking out an object. That is why "to use the symbol '=' is likewise to designate [something] an object."[21]

Frege's problem was set by the fact that the absence of causal contact with and mental images of numbers made the possibility of picking them out as objects of thought and knowledge seem particularly mysterious. Reconceiving the problem of securing singular reference in terms of recognition judgments yields the result that "to obtain the concept of Number, one must fix the sense of a numerical identity."[22] The general account of what it is to talk and think about particular objects accordingly shows how our cognitive and conceptual grasp on numbers can be made intelligible in terms of our capacity to take or treat sentences involving numerical terms as expressing identity claims.

> Our aim is to construct the content of a judgment [den Inhalt eines Urtheils zu bilden] which can be taken as [auffassen lässt] an identity such that each side of it is a number.[23]

> In the same way with the definitions of fractions, complex numbers and the rest, everything will in the end come down to the search for a judgment-content [beurtheilbaren Inhalt] which can be transformed [verwandelt] into an identity whose sides precisely are the new numbers. In other words, what we must do is fix the sense of a recognition-judgment for the case of these numbers.[24]

So Frege's explanatory strategy begins with the idea that particular objects are to be distinguished as what can be recognized as the same again—in the sense that the norms governing the use of terms referring to them would be made explicit by associating with them not only criteria of application but also criteria of identity. The recognition judgments that express the applicability of such norms are thus to be construed as identity claims. To carry this strategy through to completion, Frege must address two further issues, one quite general and the other specific to the case of numbers (as abstract objects). The general question is what it is to "fix the sense of an identity": How must a sentence be used, what sort of significance must it be accorded, in order to confer the content of an identity claim? The question specific to numbers is then what is required to confer such content on claims involving numerical expressions.

Frege's answer to the first question is straightforward, and just as it should be from the point of view of the discussion of using expressions as singular terms in Chapter 6. Identity claims make explicit *substitution* licenses. "In universal substitutability [allgemeinen Ersetzbarkeit] all the laws of identity

are contained."[25] Since identity claims are the form of recognition judgments, recognizing an object as the same again is itself to be understood in terms of substitutional commitments. "When are we entitled to regard a content as that of a recognition-judgment? For this a certain condition has to be satisfied, namely that it must be possible in every judgment to substitute without loss of truth the right-hand side of our putative identity for its left-hand side."[26] The consequences of application distinctive of identity claims consist in the undertaking of substitution-inferential commitments. What is made assertionally explicit as a claim of the form $a = b$ is commitment to a pattern of inferences requiring doxastic (assertional) commitment to the claim expressed by Pa whenever one undertakes doxastic commitment to the claim expressed by Pb, and vice versa. Frege understands particular objects as what we get cognitive and conceptual access to by using expressions as singular terms, and he offers a substitutional construal of what it is to use expressions as singular terms. The proprieties governing the circumstances and consequences of their application are those codified explicitly in identity claims, which have the significance of symmetrical substitution licenses.

4. The Maximal Substitutional Requirement on Using an Expression as a Singular Term

There is more to introducing a new term by "fixing the sense of an identity" involving it, however, than just understanding what one is committing oneself to in asserting such an identity. Ordinary cases of term introduction are special in a way that tends to obscure what more (beyond a general understanding of identity) is required to fix the sense of identities in which the new expression occurs. The sort of example that best highlights what else Frege takes to be needed is that of introducing not only new terms but new objects. The lesson appears most clearly from consideration of the role played by *abstraction* in his account of how we can become entitled to use numerals as names of definite objects.

The key point is that to be entitled to introduce a new term as the name of an object, one must settle when it would be *correct* to recognize the object picked out as the same again; in this way one distinguishes it from all other objects.[27] Frege officially insists that to do this one must see to it that the truth or falsity of *all* identities involving it has been settled. Doing so is settling when it would be correct to recognize the object picked out as the same again, and thus distinguishing it from all other objects. When what settles the truth-values of these identity claims involving a term is made explicit, it takes the form of a criterion of identity.[28]

In run-of-the-mill cases of term introduction, this requirement is quite easy to satisfy. For in the central cases a new singular term is being introduced to refer to an object that can already be referred to by using other terms

already available in the language. When a proper name is introduced for a person, place, or perceivable thing, there are typically already-individuating sortals in place appropriate to it, and it can be picked out by definite descriptions using those sortals, combined with specifications distinguishing it from others of its kind (for instance, spatiotemporal ones): 'the person who just came out of the front door of Jay's Bookstall', 'the Northwest corner of the intersection of Forbes and Meyran avenues', 'the black telephone in that corner', and so on. Assuming that the use of these antecedently available terms is already in order (as far as Frege's official requirement is concerned), all that is required to introduce a new term a is commitment to a reference-fixing identity. For under these circumstances, if the term '(the x)[Dx]' is already in use in the language, then by hypothesis the truth-values of all identities of the form (the x)[Dx] = t (where t is another term already in use in the language) have been settled. The introducing stipulation that a = (the x)[Dx] then automatically settles the truth-values of all the identities involving a and antecedent vocabulary: a = t just in case t = (the x)[Dx], and not otherwise. In these cases, then, committing oneself to the truth of a single-identity claim linking the novel term to a familiar one serves to fix the sense of *all* the identities involving the novel term, for it settles all their truth-values.

Clearly this technique is not available for introducing new terms for *new* objects—ones that cannot already be referred to in the antecedent vocabulary.[29] The problem of introducing numerical expressions referring to numbers, Frege says, is the problem of fixing the sense of numerical identities. He does this by the method of *abstraction:* a particular way of explaining the use of novel terms (referring to novel objects) by means of the use of familiar terms (referring to familiar objects). The idea is this: Where $\ulcorner a \urcorner$ and $\ulcorner b \urcorner$ are terms whose use is already established, new terms of the form $\ulcorner fa \urcorner$ and $\ulcorner fb \urcorner$ can be introduced wherever there is an equivalence relation R available defined on the old vocabulary. For one can then define the sense of identities involving f terms by stipulating that

$$fa = fb \quad \text{iff} \quad Rab.$$

In this way, if $\ulcorner a \urcorner$ and $\ulcorner b \urcorner$ are terms designating lines, one can introduce new terms of the form $\ulcorner$direction of $a \urcorner$ and $\ulcorner$direction of $b \urcorner$ (and hence the new sortal or object-kind *directions*) by appeal to the equivalence relation . . . *is parallel to___* defined on lines:

the direction of a = the direction of b iff a is parallel to b.[30]

In just the same way, if $\ulcorner a \urcorner$ and $\ulcorner b \urcorner$ are terms designating collections of already-available objects, one can introduce new terms of the form $\ulcorner$number of $a \urcorner$ and $\ulcorner$number of $b \urcorner$ (and hence the new sortal or object-kind *numbers*)

by appeal to the equivalence relation . . . *can be put in one-to-one correspondence with*___ defined on collections of objects:

> the number of a = the number of b
> iff a can be put in one-to-one correspondence with b.[31]

The claim that the relevant equivalence relation obtains between the familiar objects accordingly serves as the content of a judgment that can be taken as or transformed into an identity relating numerical (or direction) expressions, as Frege requires in the two passages quoted above.[32] That the judgment *Rab* can be reconstrued as an assertion of identity involving terms referring to novel abstract objects—rather than just as asserting a relation between familiar concrete (relative to this construction) ones—depends just on R being an equivalence relation; to be entitled to the reconstrual of such claims as putting us in cognitive and conceptual touch with abstract objects is just to be entitled to characterize R as reflexive, symmetric, and transitive. For since "in universal substitutability all the laws of identity are contained," it follows that

> in order to justify our proposed definition of the direction of a line, we should have to show that it is possible, if line a is parallel to line b, to substitute
>
> 'the direction of b'
>
> everywhere for
>
> 'the direction of a'.
>
> This task is made simpler by the fact that we are being taken to know of nothing that can be asserted about the direction of a line except the one thing, that it coincides with the direction of some other line. We should thus have to show only that substitution was possible in an identity of this type, or in judgment-contents containing such identities as constituent elements. The meaning of any other type of assertion about directions would have first of all to be defined, and in defining it we can make it a rule always to see that it must remain possible to substitute for the direction of any line the direction of any line parallel to it.[33]

Showing that the relation R on which the abstraction is based is an equivalence relation entitles one to regard *Rab* as an identity relating the new expressions *fa* and *fb* (circumstances of application). Regarding it that way is undertaking a substitutional commitment to the propriety of the inference from *P(fa)* to *P(fb)*, and vice versa, for any sentential context in which one discerns a primary occurrence of the new terms (consequences of application). The doctrine of abstraction Frege puts forward here is the claim that the significance of attributing this constellation of entitlement and commit-

ment is as taking the subject of those deontic statuses to be in a position to make judgments (to think and talk) about a new range of abstract objects—which may be thought of as equivalence classes of the old ones. This is how objects, paradigmatically mathematical ones, which we do not have causal commerce with (and need not be able to form mental images of) can be "given to us."[34]

As indicated above, Frege's official view is that to introduce a new term one must settle the truth-values of *all* identities relating it to other terms. This requirement leads to disastrous results in the later *Grundgesetze*, and Frege never does find an acceptable way to satisfy it for the introduction of terms referring to abstract objects.[35] In any case, the sort of abstractive definition just considered "fixes the sense of numerical identities" only in the sense of settling the truth-values of identities, both sides of which are numerical expressions—in the general case identities of the form $fa = fb$, but not of the form $fa = c$, where $\ulcorner c \urcorner$ is a bit of antecedent vocabulary, a term referring to an object that is concrete relative to the abstractive method of term-and-object introduction. The significance of the failure of abstractive definitions to meet the very strong condition Frege puts on term introduction—what one must do or show in order to be entitled to use an expression as a singular term—depends not only on whether it is possible to satisfy that condition in some other way but also on the reasons there are for endorsing that condition.

Frege's basic insight is that the essential singular referential purport involved in singular-term usage consists in the role such terms play in identity claims. Since he further analyzes what is expressed by identity claims in terms of the significance of such claims as intersubstitution licenses, this amounts to taking singular referential purport to consist in a structure of symmetric substitutional commitments. It is in terms of the undertaking and attributing of such substitutional commitments that the scorekeeping significance of using a singular term to express a claim is to be understood. Frege takes it that the strong condition he imposes on successful term introduction is a consequence of this substitutional analysis of what it is to use an expression as a singular term. For he takes it that unless the truth-values of *all* identities involving the candidate term have been settled, it has not been settled what one would be committing oneself to by employing it to make claims (for the identities merely make substitutional commitments explicit, that is, assertible). Abstractive definitions settle whether in using one of the new terms to make a claim of the form $P(fa)$ one is thereby committing oneself also to $P(fb)$, but they do not settle for arbitrary c whether one is committing oneself to $P(c)$.

Appealing to symmetric substitutional commitments (a species of inferential commitment) to explain what it is to use an expression as a singular term—the fundamental Fregean insight that is developed in detail in the previous chapter—does not necessitate the maximalist reading of what is

required for successful introduction of terms that Frege thinks follows from it. Frege thinks that there is something wrong with using an expression where it has been settled (whether or not anyone in particular is in a position to tell—a matter of status rather than of attitudes) that in endorsing a sentence in which it appears, one is thereby committed to the claims expressed by *some* substitutional variants of that sentence (and precluded from entitlement to commitments to the claims expressed by various other substitutional variants of that sentence), if there are some *other* substitutional variants on which one is *not* thereby counted as taking up a stance.[36] But what is wrong with its being settled that when I claim that the largest number that is not the sum of the squares of distinct primes is odd, I am thereby in some sense committing myself (whether I know it or not) to the claim that 17,163 is odd, am making a claim incompatible with the claim that 17,163 is even, and am not taking a stand on the question of whether England or the direction of the earth's axis is odd? Why would not such a situation count as one in which it had been settled exactly what I am and am not committing myself to (and similarly for entitlements), and so one in which a perfectly *definite* sense is associated with the numerical expressions involved, even though that sense is not *complete* in the way that Frege wants to require?

For many purposes it may be appropriate to insist on Frege's strong condition that the truth-values of all identities be settled; these may even include the purposes that motivate the development of the formal language Frege uses in the *Grundgesetze*. The issue is not even one of whether, relative to these purposes, a language in which this condition is imposed is *better* than one in which it is not. The question of interest at this point is rather whether there is some way of relaxing Frege's condition while maintaining the features of singular term use that make it appropriate to think of them as purportedly (and in favored cases successfully) picking out particular objects. Furthermore, it would be of interest to know just how weak the condition on the symmetric substitutional commitments associated with an expression could be made without endangering its singular referential purport. What is the *minimal* substitutional requirement (or necessary condition) on using an expression to pick out an object in thought, corresponding to the *maximal* substitutional requirement (or sufficient condition) that Frege imposes?

5. The Minimal Substitutional Requirement on Using an Expression as a Singular Term

A good place to begin in addressing this question is to notice that even according to Frege, to fix the sense of a novel term (for instance a numerical expression) it is *not* in fact sufficient merely to settle the truth-values of all the identities it can occur in. For it would not suffice for term

introduction to settle the truth-value of all the nontrivial identities—all those that relate the term to some *other* term—as *false*. No criterion of identity is implicitly associated with the expression ⌈a⌉ by stipulating that a = a, but that if ⌈b⌉ is any expression distinct from ⌈a⌉, then a = b is false. Settling the truth-values of all the identities involving the new expression in this way does not even implicitly involve associating with it an object that can be recognized as the same again. "All identities would then amount simply to this, that whatever is given to us in the same way is to be reckoned as the same. This, however, is a principle so obvious and sterile [*unfruchtbar*] as not to be worth stating. We could not, in fact, draw from it any conclusion which was not the same as one of our premises. Why is it after all that we are able to make use of identities with such significant results in such diverse fields? Surely it is rather because we are able to recognize something as the same again even though it is given in a different way."[37]

Objects are essentially things that can be recognized as the same again, even though given in different ways. That is why they are things for which the issue of identity arises—why using the identity sign with an expression is treating it as referring to an object.[38] To be an object is to be something that can be referred to in different ways; to associate an object with an expression as its referent requires settling what would count as *another* way of picking out that same object. Frege's maximalist claim is that introducing a term as picking out a definite object requires settling *every* other way of picking out that same object. The corresponding minimalist claim is that it requires settling at least *some* other way of picking out that same object.

The thought can be put more clearly by shifting from material mode to formal mode: from talk of objects to talk of the substitutional significance of singular terms by means of which talk of objects is officially to be understood. The basic idea is that unless the occurrence of a candidate term in the expression of a claim has *some* substitution-inferential significance (unless it commits one to *some* further claim that is expressed by a sentence resulting from the first by substitution of another term for the candidate), then the candidate is not functioning as a singular term at all. Its occurrence is not semantically significant in the way terms are; it is substitutionally idle, thus inferentially idle, and therefore semantically idle to discern its occurrence at all.[39] The minimal condition on using an expression as a singular term that emerges from understanding the characteristic substitutional role terms play is just that it must have been settled that the occurrence of the putative term have *some* (symmetric) substitution-inferential significance. As elaborated in the previous chapter, for the occurrence of an expression to have a significance of the kind characteristic of singular terms, its use must be governed by some simple material substitution-inferential commitments (SMSICs)—commitments of the sort that can be expressed explicitly as nontrivial identity claims or recognition statements. Where Frege demands a *complete* set of substitutional commitments associated with each term, the

minimal demand compatible with a substitutional understanding of singular terms (motivated by the observation that Frege would not permit all the nontrivial identities to be settled as false) is that a *nonempty* set of substitutional commitments be associated with each term; at least one nontrivial identity must be settled as true.[40]

Talk of objects as what can in principle be recognized as the same again—what can be given to us or referred to in different ways—reflects the structure of substitutional significance that the occurrence of bits of subsentential vocabulary must have for them properly to be understood as having the indirectly inferential content characteristic of the use of singular terms. The singular referential purport of such vocabulary consists in the fact that the deontic scorekeeping significance of its use is to be determined by symmetric substitutional commitments that link it to other vocabulary. These are the commitments that are made explicit by the nontrivial identity claims that Frege calls "recognition judgments"—which he takes to express the recognition of an object as the same again, though given in two different ways. This much of Frege's thought in the *Grundlagen* can be taken over without a consequent commitment to the requirement that the truth-values of *all* nontrivial identities must be settled in order for a singular term to have been properly introduced.

Even the minimal claim that settling the truth of some nontrivial identities involving a candidate singular term is a necessary condition for using it as a name of an object, however, has consequences that can seem mysterious unless the substitutional gloss on that claim is kept firmly in mind. For it follows that the idea of an object that can be picked out or referred to only in one way is not an idea of an *object* at all. (Recall the discussion above in 6.4.) A language cannot refer to an object in one way unless it can refer to it in two different ways. This constraint will seem paradoxical if referring to an object by using a singular term is thoughtlessly assimilated to such activities as using a car to reach the airport or using an arrow to shoot a deer: even if only one car or one arrow is available and impossible to reuse, what one is doing can still genuinely be driving to the airport or shooting the deer. Why should referring be different, something that cannot be done one way unless it can be done two ways? Understanding an expression's purporting to refer to an object in terms of its use being governed by proprieties articulating its significance according to substitution-inferential commitments dispels the puzzlement that can otherwise attend this phenomenon. An object that can be referred to in only one way is the sound of one hand clapping.

So for an expression to be used as a singular term, there must be *some* substantive substitutional commitment undertaken by the one who uses it. It is not necessary that either the one who undertakes that commitment or the one who attributes it—by attributing a doxastic commitment that would be avowed by the assertion of a sentence containing the singular term—be able to specify just what the content of that commitment is. But it is only

where the interpreter takes it that there *is* some such substitutional commitment included in the significance of the underlying doxastic commitment that the one who undertakes that commitment is interpreted as using a singular term to make a claim about an object. This is just the conclusion that was drawn in Chapter 6: the category of singular terms should be understood as comprising expressions whose proper use is governed by simple material substitution-inferential commitments (SMSICs) linking them to other such expressions. Taking an expression to be a singular term—taking it to purport to pick out a particular object—just is taking its use to be governed by some such SMSICs. When such a simple material substitution-inferential commitment linking two expressions is made propositionally explicit (as an assertible), it takes the form of a nontrivial identity claim. That is why to introduce an expression as a singular term, one must somehow settle the truth-value of at least one such identity (of what can be so expressed in an idiom with suitable explicitating resources—that is, logical vocabulary). Purported reference to objects must be understood in terms of substitutional commitments linking diverse expressions.

6. *Substitutional Triangulation*

This *substitutional holism*—according to which mastery of the use of one expression as a singular term involves mastery of the use of many—is the reflection at the subsentential level of the *inferential holism* according to which mastery of the use of one expression as a sentence (even one that can be used to make noninferential reports) involves mastery of the use of many.[41] Carving up sentences according to their substitutional relations to one another is just a method for extending the notion of content-conferring, inferentially articulated deontic significance to the subsentential level—to expressions that cannot themselves play the directly inferential roles of premises or conclusion of inferences. The conceptual content expressed by a sentence depends on its place in a network of inferences relating it to other sentences; the conceptual content expressed by a singular term depends on its place in a network of substitutions relating it to other terms. The substitutional roles that determine the pragmatic significance of the occurrence of singular terms are a kind of indirectly inferential role because substitutional commitments are a kind of inferential commitment.

Another topic this minimal substitutional requirement for using an expression as a singular term illuminates concerns picking out objects by conceptual *triangulation*. Triangulation strategies arise from consideration of a fundamental problem concerning the discrimination of a particular *stimulus* to which some sort of response is reliably keyed. In his discussion, Davidson introduces the familiar point this way: "Why say the stimulus is the ringing of the bell? Why not the motion of the air close to the ears of the dog—or even the stimulation of its nerve endings? Certainly if the air were

made to vibrate in just the way the bell makes it vibrate it would make no difference to the behavior of the dog. And if the right nerve endings were activated in the right way, there still would be no difference."[42] Typically, there is a whole causal chain of covarying events culminating in a response. In the standard case, the occurrence of one is accompanied by the occurrence of all the rest. Under these circumstances, the response being keyed to one of the event kinds is its being keyed to all the rest. How is one element of the chain to be singled out as the stimulus? What is the nature and source of the privilege that distinguishes one element from another?

One strategy for assigning such privilege, and therefore picking out as the stimulus one element from the whole chain of covarying event types that culminates in a response of the specified type, is (as Davidson goes on to suggest) to look to *proximity* to the eventual response. The justification for seizing on causal proximity of stimulating event to the response as what matters is maximizing the relative *reliability* of the connection between the occurrence of events of the distinguished stimulus type and the occurrence of events of the distinguished response type. The proximal element of the chain is the one that most reliably brings about the response. For prior occurrences in the chain elicit the response only in the cases where they succeed in bringing about an event of the proximal type, while events of that type can elicit the response regardless of whether they have themselves been brought about in the standard way. The trouble is that such a proximal theory of stimuli will always yield the result that the stimuli being responded to are at the sensory surfaces or within the nervous system of the responding organism.

In the context of the project of using reliable differential responsive dispositions as a model to understand which objects basic empirical concepts are being applied to, the adoption of such a policy for the discrimination of stimuli is disastrous. For what is classified by the protoconcepts that repeatable responses are going proxy for is not bells and tables and rabbits but only states of the responding organism. Nothing that looks like one of our ordinary empirical concepts, applying to ordinary observable objects, is within reach of such an approach. A distal strategy is required in order to get the protoconcepts represented by reliably differentially elicited noninferential response types to count as classifying and so applying to ordinary observable objects and properties. Understanding them this way involves respecting the language-learning situation in which these reliable differential responsive dispositions are established.

The most popular approach to identifying distal stimuli as what is classified by the exercise of reliable differential responsive dispositions is to appeal to *triangulation*. This is a strategy for picking out or privileging one bit of the causal chain of covarying event types that reliably culminates in a response of a distinguished type, by looking at the *intersection* of *two* such chains. The insight it develops is that the best way to pick a single *point* (the

stimulus) out of a *line* (the causal chain of covarying event-types that reliably elicit a response of the relevant type) is to *intersect* it with another line—another causal chain corresponding to another reliable differential responsive disposition.

One writer who employs such a triangulation strategy to address the problem of picking out distal stimuli as what a response is *about* is Dretske.[43] In order to pick out the distal stimulus he looks to the upstream intersection of two distinct "flows of information" (or causal chains of reliably covarying event-types) that reliably culminate in responses of the same type. A simple example of the sort of system he has in mind would be a thermostat that keeps the temperature of a room within a certain range by turning a furnace on and off. If the thermostat has only one way of measuring temperature—for instance by the bending of a bimetallic strip until it touches either the left electrical contact (too cold) or the right one (too warm)—there is no way, Dretske acknowledges, to say that what the system is responding to is the temperature of the room, rather than the temperature of the bimetallic strip or the curvature of the bimetallic strip or the closing of the circuit between the bimetallic strip and one or the other of the contacts. Notice that a pragmatist appeal to practical *consequences* of the response in question is of no help here; turning the furnace on affects not only the temperature of the room but also that of the bimetallic strip, its curvature, and so its relation to the electrical contacts.

The idea is that one can be entitled to such a description if the thermostat is slightly more complicated and has another causal route to the same response (turning the furnace on or off). If the thermostat has a second sensor—for instance a column of mercury supporting a float with an electrical contact that completes one circuit to turn the furnace on whenever the float is below one point (too cold) and turns it off whenever the float is above another point (too warm)—then the system has two ways of responding to the change in temperature in the room. Although for this second route by itself (just as for the first by itself) there is no feature of the system that entitles one to say it is responding to changes in the temperature of the room rather than to the temperature of the mercury or the length of the mercury column or the closing of the switches, when the two routes are considered together, they intersect in just two places—upstream at the change of temperature in a room (which is included in the "flow" or causal chain corresponding to each route) and downstream in the response of turning the furnace on or off.[44] Dretske shows how the general strategy of looking to the *intersection* of two reliable differential responsive dispositions might be funded from the resources of the responding system itself.

One might worry that Dretske has not in fact succeeded in responding to the general worry about how to justify describing the system as responding to a distal stimulus rather than a proximal one. For there is an objection available to his strategy that seems to reinstate the original worry. Why, it

might be asked, ought we not to conclude that even in the two subsystem case, what is responded to is a proximal stimulus, but a *disjunctive* one? The system turns on the furnace just in case *either* the temperature of the bimetallic strip is low enough *or* that of the mercury column is low enough, or alternatively, in case the curvature of the bimetallic strip is far enough to the left *or* the mercury column is short enough. (Again, pragmatic appeal to the practical consequences of entering this state will not solve the problem.)

This worry is connected to the complaint voiced already in Chapter 2, to the effect that mere differential responsiveness is not sufficient for identifying the responses in question as applications of *concepts*. The rationalist supplementation suggested there—that what is distinctive of the conceptual is the *inferential* role played by the responses that stimuli differentially elicit—is also what is required to exploit the triangulation strategy in connection with genuine concepts in a way that responds to the worry about disjunctive proximal stimuli.

Consider a man who reliably responds (as one wants to say) to the visible presence of rabbits by saying "Gavagai." Suppose further that he is reliably differentially responding not just to rabbits, but to the presence of the distinctive (according to him) rabbit flies that are for him decisive evidence of the presence of rabbits, or that the visual cue he is using, as determined by a physiologist of perception, is a glimpse of the fluff around the tail of the rabbit. What is it about the situation in virtue of which he can be said nonetheless to be reporting not the presence of the rabbit flies or of the fluffy tail but the presence of a rabbit? The inferentialist response is that the difference is not to be found in the reliable differential responsive dispositions, not in the causal chain of covarying events that reliably culminates in the response 'gavagai', to which not only the rabbit but the flies or the fluffy tail belong. It lies rather in the inferential role of the response 'gavagai'. For instance, does the commitment undertaken by that response include a commitment to the claim that what is reported can fly? Or is the claim expressed by 'gavagai' incompatible with the further characterization of the item reported as flying? If it is incompatible, then it is not the flies that are being reported. What determines which element of the causal chain of covarying events that reliably elicit the report is being reported is the *inferential* role of the report, what it *entails*, what is *evidence* for it, what it is *incompatible* with.

Assuming that the observable predicate corresponding to 'flying' has already picked out the things that fly, noticing that the report 'gavagai' could mean rabbit flies in case its applicability entails the applicability of 'flying' and could not mean rabbit flies in case its applicability is incompatible with the applicability of 'flying' is just what is wanted to pick out the distal stimulus the concept expressed by 'gavagai' is being applied to or is classifying. But the appeal to inference and incompatibility may seem just to put off the issue. How does 'flying' get to apply properly to flying things, and not

to whatever cues we in fact use in discriminating flying things—in short to one element of the causal chain of covarying event types that reliably culminate in its application? The answer must be that what the appeal to inferential role does is establish a sort of *triangulation,* or intersection of flows of information or reliable differential responsive dispositions. If 'gavagai' is used so as to entail 'flying', then whatever is properly responded to by the former expression must be properly responded to by the latter, so what is classified as gavagai must also be classified as flying, so 'gavagai' must apply to rabbit flies, and not to the rabbits that are their invariable (we are supposing) concomitants. In short, the appeal to inferential role, in addition to reliable differential responsive dispositions, involves triangulation of the sort that Dretske invokes, where two (or more) different reliable responsive dispositions of the system are invoked, so that their intersection can pick out a unique element of the causal chain of covarying events as the stimulus being classified by a response of a certain type. Because 'flying' will *not* be taken to apply to lots of things that merely hop, we can be sure that it does not mean *flying or hopping,* and so that 'gavagai' does not mean something disjunctive like *rabbit or rabbit-fly.*[45]

In sum, to make the triangulation approach to distinguishing distal stimuli work, one needs to look further 'downstream' from the response, as well as 'upstream'—just as orthodox functionalism would lead one to expect. What picks one kind of thing out as what is being *reported,* from among all those that are being differentially *responded* to, is a matter of the inferential commitments that response is involved in. These inferential consequences of going into a state make it clear that what is being classified is something outside the system. They are what determine that a physicist is reporting the presence of a mu-meson in a bubble chamber, and not simply a large hook-shaped pattern. For the consequences of classifying something as a microscopic mu-meson are quite different from those of classifying something as a macroscopic hook-shaped trace. It is the lack of such consequences that makes Dretske's dual thermometer liable to a disjunctive proximal interpretation. The conclusion is that *causal triangulation* by intersecting causal chains associated with reliable differential responsive dispositions must be supplemented by *inferential triangulation* associated with different concepts.

The minimal condition on singular reference that has been extracted from Frege in this section amounts to the demand that objects be picked out by *substitutional triangulation.* Taking it that an expression is being used to pick out an object is taking it that that *same* object could be picked out in some *other* way—that some commitment-preserving substitutions involving that expression are in order. Substitutional commitments are compound inferential commitments, corresponding to *patterns* of simple inferential commitments. Substitutional articulation is a kind of inferential articulation, and substitutional triangulation is a kind of inferential triangulation. The notion of substitutional commitments is what is needed to explain what

it is to take two distinct claims (whether responsively elicited or not) to be applications of concepts to the *same* object. The significance of causal triangulation is to be understood in terms of the supporting role it can play in this sort of substitution-inferential triangulation. It cannot by itself provide an analysis of picking out objects. And as Frege's discussion of picking out abstract objects shows clearly, however important a role it plays in the way perceivable objects are given to us, causal triangulation is not even a necessary component of the substitution-inferential triangulation that is what our cognitive grip on objects in general consists in.

7. Conclusion and Prospectus

This is by no means to say that the discussion of substitutional triangulation here and in the previous chapter suffices to understand what our talking and thinking about objects consists in. That discussion addresses primarily the issue of what it is for it to be *objects* (and their properties and relations) that our talk (and so our thought) purports to be about. To understand fully what it is for our thought and talk to purport to be *about* them requires an account of the crucial *social* dimension of the substitutional triangulation that structures the contents expressed by the use of singular terms (and predicates) and of the inferential triangulation that structures the contents expressed by the use of sentences. The way in which the social structure of the broadly inferential articulation of discursive practice bears on the nature of the conceptual contents that practice confers on the intentional states it institutes (and on the performances that express them) is already implicit in the discussion of discursive practice in terms of deontic scorekeeping, in Chapter 3. It is the task of Chapter 8 to make it explicit, and thereby to show how the representational dimension of conceptual content arises out of, and essentially depends on, differences in social perspective among the various discursive practitioners.

A further shortcoming in the account of picking out objects in terms of substitutional triangulation as adumbrated so far is that it is primarily addressed to the phenomenon of *purported* singular reference. Although general reasons have been offered motivating a direction of explanation that begins with the notion of representational purport, it remains to say something also about the *success* of such purport. To this end, the next section discusses what we are doing when we take it that a singular term succeeds in referring, in that the object the term purports to refer to actually *exists*. An account is offered of *existential* commitment as a kind of *substitutional* commitment. This story in turn permits an analysis of the commitments characteristic of the use of expressions as definite descriptions, and so shows how to extend the deontic scorekeeping model from languages with predicates and lexically simple singular terms to ones that contain definite descriptions as well.

The rest of the chapter then addresses the structures that make substitu-

tional triangulation (and so recognizing an object as the same again) possible when *unrepeatable* expression tokenings are involved. This means above all the *deictic* or *demonstrative* use of terms that is so important for understanding the role causal triangulation can play in substitutional and inferential triangulation, and so ultimately for understanding what is distinctive about *empirical* knowledge. Moving to the level of unrepeatable tokenings requires discerning a finer structure of *token recurrence* below that of *substitution*, just as the finer substitutional structure had to be discerned below that of *inference*. The key concept in this account is that of *anaphora*. Explaining the anaphoric linkage of tokenings in terms of the inheritance of the determination of substitutional commitments provides an official account in deontic scorekeeping terms of the phenomenon by means of which the traditional semantic vocabulary, 'true' and 'refers', was explained in Chapter 5, redeeming the promissory note issued there. So by the end of this chapter the full three-leveled structure of fundamental concepts in terms of which conceptual content is to be understood in the semantic portion of the present account will have been made available: *inference, substitution,* and *anaphora*. At that point the semantic raw materials will be on hand to be combined with the underlying pragmatics to yield in Chapter 8 an account of representation by conceptual contents, on the semantic side, and objectivity of conceptual norms, on the pragmatic side.

II. DEFINITE DESCRIPTIONS AND EXISTENTIAL COMMITMENTS

1. Forming Singular Terms from Predicates

To talk about the singular referential purport of singular-term usage is to talk about what kind of substitutional commitments one must attribute (and, as will emerge in the next chapter, acknowledge) in order for what one is doing—the practical deontic scorekeeping attitude one is adopting—to count as taking someone to be using an expression as a singular term.[46] To be doing that, one must treat the use of the expression as governed by proprieties determined by symmetric simple material substitutional commitments—commitments that in languages with sufficient logical expressive power are made explicit in the form of the nontrivial identity claims that Frege calls "recognition judgments." The substitutional commitments involving a singular term that a scorekeeper attributes and undertakes determine the pragmatic significance, for that scorekeeper, of each use of that term. That the significance for deontic scorekeeping of its occurrences is to be determined in this way is what treating it as a singular term (as purporting to pick out an object) consists in.

The referential purport that is in this way acknowledged or attributed concerns the committive antecedents and consequences of application of singular terms as such. Chapter 6, some of the points of which were reca-

pitulated in Frege's terminology in the previous section, showed how those committive circumstances and consequences of application can be understood substitutionally. But referential *purport* is one thing, referential *success* is another. What is the difference between taking it that an expression has been introduced as purporting to refer to a definite object and taking it that it in fact picks out or gives us a cognitive or semantic grasp on such an object? This is a question about a certain kind of *entitlement* to the substitutional commitments in which singular referential purport consists, and so in a broader sense about the appropriate circumstances of application of singular terms as such.

The deontic attitudes that constitute taking the singular referential purport characteristic of singular terms to be successful emerge most clearly from consideration of what is involved in taking someone to be entitled to use a definite description formed from a predicate. This is what Frege calls "the definition of an object in terms of a concept under which it falls."[47] As an example of a definite description that is defective—whose referential purport is not successful because it involves substitutional commitments the user cannot in the relevant sense be entitled to—he considers the expression 'the largest proper fraction'. The predicate that description is formed from is one that can be used to express commitments with appropriate entitlements.

> The expression "the largest proper fraction" has no content, since the definite article purports to refer to a definite object [*der bestimmte Artikel den Anspruch erhebt, auf einen bestimmten Gegenstand hinzuweisen*]. On the other hand, the concept "fraction smaller than 1 and such that no fraction smaller than 1 exceeds it in magnitude" is quite unexceptionable: in order, indeed, to prove that there exists no such fraction, we must make use of just this concept, despite its containing a contradiction. If, however, we wished to use this concept for defining an object falling under it, it would, of course, be necessary first to show two distinct things:
>
> 1. that some object falls under this concept;
> 2. that only one object falls under it.
>
> Now since the first of these propositions, not to mention the second, is false, it follows that the expression "the largest proper fraction" is senseless.[48]

Suppose that a predicate $P\alpha$ has been introduced and is in use. (This is a supposition that has been given definite content in substitutional terms by the discussion of Chapter 6.) The problem Frege is addressing is to make explicit what else is required for it to be proper to take someone to be entitled to use a definite description formed from it—a singular term of the form 'the P', or as it may be expressed more generally, $!x(Px)$. Treating $!x(Px)$ as a singular term is taking it that its use is governed by symmetric simple

material substitution-inferential commitments, that is, that there is some true nontrivial identity involving it. Thus, there must be some true recognition statement of the form $!x(Px) = a$.

As already indicated, the two conditions Frege imposes are that there be something that is P, and that it be unique:

1. Pa and
2. For any y, if Py, then $y = a$.

The second condition—uniqueness—can be straightforwardly parsed in the substitutional idiom already available. For it just amounts to saying that for any terms t, t', if Pt and Pt', then $t = t'$; the substitutional commitments associated with claims of these forms have already been explained. The first condition requires more discussion, however. In the *Grundlagen* Frege frames his dispute with the formalists in terms of the necessity of proving the existence of an object falling under the concept, or as he puts it, *"producing* something that falls under it."[49] What must one do to satisfy this requirement?

2. Substitutional Commitments Expressed by Quantifiers

It too can be understood in substitutional terms. Existential commitments are a kind of substitutional commitment, related to, but not identical with, the substitutional commitments involved in the use of quantifiers. As the discussion of the formation of complex predicates in the previous chapter indicates, universal and particular quantifiers are logical locutions that have the expressive function of making propositionally explicit conjunctive and disjunctive substitutional commitments. Attributing commitment to a claim of the form $(x)Px$ is attributing commitment to *all* claims of the form Pa. Such a substitutional rendering of quantification has been criticized as inadequate in cases where, for cardinality reasons, there are not enough singular terms to pick out all the objects one is quantifying over. It is very important that we be able to make claims about all real numbers— for instance that every one can be represented by converging sequences of rational numbers—even though we are in principle limited to the use of at most a countable number of singular terms referring to them.[50] In fact, however, that the stock of available substituends is in this way *limited* threatens a substitutional construal of quantifiers only if that stock is in addition conceived of as being *fixed*.

It is of the essence of singular-term usage that new terms can always be introduced—both new terms for familiar objects and terms that introduce unfamiliar objects, paradigmatically by description. We cannot indeed extend our language so as to have separate terms for *all* real numbers at once, but *each* real number can be picked out. For there is no real number that we cannot specify—for instance by a definite description in terms of converging

sequences of rational numbers. The substitution instances Pa one becomes jointly and severally committed to by committing oneself to a claim of the form $(x)Px$ include, not only those formed from terms a that are currently in the language, but all those that *could* be introduced (not necessarily simultaneously). The substitutional construal of quantificational commitments requires that the expressive powers of a set of discursive practices be conceived in the wider sense that takes account of the possibility of introducing novel expressions, rather than in the narrower sense that restricts attention to locutions already actually in use. This latter view amounts to freezing an idiom: taking a snapshot of it and evaluating its expressive capacities in abstraction from the process by which it develops. It is encouraged by thinking of languages as formal objects (perhaps set-theoretic structures) that have fixed vocabularies. If languages are instead conceived as living practices, then the ways in which new vocabulary is introduced take their place as fundamental aspects of those practices—as central as the ways in which new claims are made. Frege is the father of the formal approach to languages, but his project in the *Grundlagen* leads him to be vitally concerned with the process of introducing novel expressions functioning as singular terms, not only for unfamiliar objects of familiar kinds (by description), but even for unfamiliar objects of unfamiliar kinds (by abstraction).

Similarly, the use of a particular quantifier in connection with a complex predicate makes explicit a disjunctive substitutional commitment to the effect that for some term a, Pa. To be entitled to such a claim one may, but need not, be able to produce the relevant substitution instance. The vindicating substituend a need not even already be in the language; one is committed only to the possibility of introducing such a term. The point of the existence requirement Frege imposes on entitlement to introduce definite descriptions is that a certain kind of bare stipulation is not in general enough to entitle one to such term introduction. One is not permitted without further ado to introduce the expression $!x(Px)$ and then, relying on the fact that $P(!x(Px))$ (whenever use of the definite description is appropriate), to use that description as the substituend that vindicates the claim made by use of a particular quantifier. The large question of interest in this section is precisely what that existential condition on the employment of definite descriptions comes to.

Though it is common to do so, it is not necessary, however, to extend the existential condition Frege imposes on the use of definite descriptions to whatever counts as a vindicating substituend for a particular quantification. It is for this reason that the general formal notion of particular (that is, disjunctive) quantification should be distinguished from the substantive notion of specifically existential quantification. Free logics distinguish particular quantificational commitments from existential commitments so as to allow an idiom in which 'Pegasus is a winged horse' can count as true, even though Pegasus does not physically exist, and so in which 'Pegasus' can serve

as a substituend that vindicates a particular quantificational commitment to there being some winged horses. Frege, of course, does not want to talk this way, taking it that because 'Pegasus' has sense but no referent, 'Pegasus is a winged horse' cannot be true. For many purposes (certainly for Frege's), this policy is no doubt the best. Nonetheless there is nothing incoherent about scorekeeping practices that permit particular quantificational commitments to be vindicated by term substituends with respect to which the scorekeeper does not undertake existential commitments, and considering such ontologically relaxed idioms highlights certain important features of genuinely existential commitments.

The substitutional significance of particular quantification is entirely determined by features of discursive scorekeeping practices that have already been discussed if it is stipulated that a particular quantificational commitment with respect to the predicate $P\alpha$ is to be equivalent to the commitment expressed by $\sim(x)\sim Px$. The negation of a claim p was defined in Chapter 2 as its minimum incompatible: the inferentially strongest claim that is commitment-entailed by every claim incompatible with p. Thus the claim that for some x, Px is incompatible with any claim that for every term a entails some claim Qa that is incompatible with Pa. So the particular quantificational claim that for some x, Px is both commitment- and entitlement-entailed by any claim of the form Pa (but not necessarily just by these).

3. Sortally Restricted Substitutional Commitments

In his technical systems (both that of the early *Begriffsschrift* and that of the late *Grundgesetze*) Frege offers a substitutional account of the formation of complex predicates and of the formation of sentences from them by the application of quantifiers. In each case the scope of the quantifiers is unrestricted: *every* well-formed singular term can serve as the substituend that vindicates a particular quantification and can serve as a potential counterinstance to a universal quantification. One consequence of running these systems with their quantifiers wide open is that in order to give quantificational claims the force Frege wants them to have—above all for the assertion of claims formed by particular quantification to involve the undertaking of specifically *existential* commitment—Frege must ensure that it can be proven that every well-formed singular term has a referent. As Russell notoriously showed, another consequence—in the context of the expressive power provided by unrestricted formation of complex predicates or sentential functions by substitution—is that the resulting systems are inconsistent. That unpalatable result has prompted the investigation of how the various theoretical commitments that conspire to produce it might be relaxed so as to avoid it. One popular candidate is Frege's insistence on unrestricted quantification; the strategy of placing restrictions on the substitution instances relevant to the semantic evaluation of claims formed by the application of

particular and universal quantifiers is what lies behind the theory of types Russell develops in *Principia Mathematica* to evade Frege's difficulty.

Restricted quantification, however, is not of merely technical interest. Indeed, Frege's unrestricted version appears as an artificial extrapolation once it is realized that in natural languages ordinary quantificational tropes are sortally restricted.[51] The central uses of quantifiers are to make claims such as:

> Every integer is the sum of nineteen or fewer fourth powers.[52]
> Some nineteenth-century German philosophers did not care about ontological issues.
> All bank employees must wear neckties.
> A deer made those tracks less than an hour ago.
> Each of us has intentional states.

The central quantificational construction is *every K* or *some K*, where *K* is a sortal expression such as 'dog' or 'book'. 'Everyone' and 'someone' have the sense of 'every person' and 'some person', and even the apparently wide open 'everything' usually carries some restriction, either explicitly, as in

> Everything the author says about propositional content is confused,

or implicitly, as in

> Everything is a disappointment in the end.

The sortal restriction puts conditions on allowable substituends, so that even though 'the author of *The Stones of Venice*' is a perfectly good singular term, substitution instances formed from it are not relevant to the semantic evaluation of "Every integer is the sum of nineteen or fewer fourth powers."

As Frege indicates in the *Grundlagen*, sortals are like predicates, except that they have not only criteria and consequences of *application* but (like singular terms) also criteria (and so consequences) of *identity*. For many purposes, '. . . is a dog' functions predicatively, just as '. . . is large' does. But if *a* is a dog and *b* is a dog, it makes sense to ask whether *a* is the *same* dog as *b*. Sortals have associated with them practices of identifying and individuating the things they apply to, as nonsortal predicates do not. So in order to introduce the sortal 'number', Frege insists on "a general criterion for the identity of numbers [*Kennzeichen für die Gleichheit von Zahlen*]."[53] When made explicit in the form of a claim, such a criterion has the form:

> If *x* is a *K* and *y* is a *K* and *Rxy*, then *x* is the same *K* as *y*.

Introducing a sortal, like introducing a predicate, requires fixing the sense of claims formed by substitution into sentence frames of the form "α is a *K*"

("α is *P*"), but it requires in addition establishing a criterion of identity. Satisfaction of this additional constraint ensures that *K*'s can be *counted*.

In fact, establishing a criterion of identity (and so a sortal) is not only sufficient for countability, it is necessary as well. Unsortalized 'things' or 'objects' cannot be counted. There is no answer to the question how many things there are in this room; there is one number of books, another of molecules, another of atoms, another of subatomic particles. As Frege says: "If I place a pile of playing cards in [someone's] hands with the words: Find the Number of these, this does not tell him whether I wish to know the number of cards, or of complete packs of cards, or even say of honour cards at skat. To have given him the pile in his hands is not yet to have given him completely the object he is to investigate; I must add some further word— cards, or packs, or honours."[54] Counting is intelligible only with respect to a sortal concept.

'Thing' and 'object' are pseudosortals. They can occupy the syntactic positions occupied by sortals, but they do not individuate as sortals must. They are mere placeholders for sortals, used when for some reason—often the clumsy disjunctiveness of the sortal that would be required—one does not want to specify the relevant sortal explicitly. When we say something like "Put everything that is on top of the desk into the drawers," we usually mean all the middle-sized bits of dry goods: books, papers, pens, paper clips, and so on. We do not mean 'things' such as designs in the dust, cool spots, drops of water, and so on. One of the central uses of 'one' in English is as an *anaphoric prosortal*—an anaphoric dependent standing in for a sortal that is its antecedent—as in "John quoted an English philosopher, and I quoted a German one," or "Eric wants an ice cream cone, and Russell wants one too." In these examples 'one' is used in the 'lazy' anaphoric way, where it is replaceable by its antecedent sortal. Like the pronoun 'it', however, it is promiscuously available to stand in for a wide variety of antecedents. 'Thing' and 'object' are what one gets if one misunderstands this grammar and instead construes 'one' as expressing a genuine sortal.

Frege in fact makes exactly this objection to the attempt to press the term 'unit' (or 'one') into generalized duty in place of substantive sortals in theorizing about counting. His own view is that the invocation of substantive sortal concepts cannot in this way be avoided; he takes it rather that "a concept [is] the unit relative to the Number which belongs to it."[55] Not all concepts will do; only those expressed by sortals (rather than predicates without individuating criteria of identity): "The concept 'syllables in the word "three"' picks out the word as a whole, and as indivisible in the sense that no part of it falls any longer under the same concept. Not all concepts possess this quality. We can, for example, divide up something falling under the concept 'red' into parts in a variety of ways, without the parts thereby ceasing to fall under the same concept 'red'. To a concept of this kind no finite number will belong. The proposition asserting that units are isolated

and indivisible can, accordingly, be formulated as follows: Only a concept which isolates what falls under it in a definite manner, and which does not permit any arbitrary division of it into parts, can be a unit relative to a finite Number."[56] This insight ought to have led Frege to see quantifiers as coming with sortal restrictions on the admissible term substituends. For quantifiers quantify; they specify, at least in general terms, *how many*, and how many there are depends (as Frege's remarks about playing cards indicate) on *what* one is counting—on the sortal used to identify and individuate them. As Frege saw clearly, the use of quantifiers depends on the use of the singular expressions that provide their substitution instances. It is best therefore to think not only of quantifiers but of singular terms as properly introduced only in connection with some at least implicit sortal.

Definite descriptions should be explicitly sortally restricted: 'the *man* in the brown suit', 'the *book* that Carlyle had to rewrite because of Mill's maid', and so on. Individual proper names and demonstratives and other indexical expressions cannot properly be understood except in terms of their associated sortals. Thus one cannot simply point in the direction of a statue of a man on a horse and christen ^{s}it^s 'Lumpl'. It matters whether one is naming the statue or the lump of clay it is made of. If the former, reshaping it into a statue of a mother with a child destroys Lumpl; if the latter, not—for the transformed figure is the same lump, but a different statue.[57] If I hold up my copy of Kant's first Critique and ask "Has Eric read this?" my remark is susceptible to two different sorts of readings, depending on whether the demonstrative is associated with the sortal that individuates books according to the content of the text or rather (as might arise if I have just discovered a large jelly stain defacing the page that sets out the Table of Judgments) the sortal that individuates them according to particular physical copies. 'This' or 'that' used by itself should on these grounds always be understood as elliptical for 'this K' or 'that K'. Again, it is important that 'I' implicitly invokes the sortal that individuates persons—it is a *personal* pronoun. For I, who am flying to London, am the same person who last month flew to Philadelphia, while I am *not* the same passenger who did so.[58] The discussion in this vicinity about "relative identity" is prompted by this sort of observation. But it often takes the form of a mysterious thesis about *things*, rather than a clear one about the conditions that ought to be met to count as having introduced (or understood) a singular term (even a tokening of a demonstrative) as having a definite reference.[59] Such a confusion is the result of thinking of sortally unrestricted quantification and singular-term usage as conceptually fundamental, and seeing sortal restrictions as optional additions—rather than seeing the restricted case as fundamental, and unrestricted quantification as a dangerous and often unwarranted extrapolation based on a misunderstanding of the way pseudo- and prosortals such as 'thing', 'object', 'one', and 'item' function. Frege's requirement that to introduce a referring expression one must fix the sense of identities involving it—settle how it is

to be distinguished from other objects—turns out to require that it be associated, implicitly or explicitly, with an individuating sortal concept.

4. Existential Commitments

Existential commitment is a species of substitutional commitment. It can be thought of as a particular quantificational commitment that involves a special sort of restriction on the vindicating substituends that determine the content of that commitment. Generically, the restriction on substituends characteristic of existential commitments is akin to the sortal restriction involved in quantificational commitments—and more generally in the use of any expressions conveying singular referential purport. The structure of the restriction is different, however. What is distinctive of specifically existential commitments is the special role played by a distinguished class of admissible substituends, here called *canonical designators.* The difference between the substitutional function of canonical designators and that of sortally qualified substituends in ordinary quantification is what stands behind the principle that existence is not a property—that existent objects are not a *kind* of object.

The best way to appreciate the role an expression must play to be functioning as a canonical designator with respect to a species of existential commitment is by examples. Three different sorts of existential commitment are considered here, corresponding to *numerical* existence, *physical* existence, and *fictional* existence: the sense in which there is a number such that every number greater than it is the sum of distinct primes of the form $4n + 1$, the sense in which there is a pen on my desk, and the sense in which there is someone who keeps house for Sherlock Holmes, respectively. By looking at these different kinds of existential commitment, it is possible to see what they all have in common, in virtue of which they deserve to be understood as species of a genus—so that 'exists' can be understood as univocal, in spite of the important differences between commitment to the existence of particular numbers, of physical objects, and of fictional characters. The claim is that what these different sorts of existential commitments share—what makes them all varieties of *existential* commitments—is the way in which their pragmatic significance is determined by a set of expressions playing the role of canonical designators. The manifest differences between them are consequences of the very different sorts of expressions that serve as canonical designators in each case.

It is clear how the sense of the expression "the smallest natural number such that every larger one is the sum of distinct primes of the form $4n + 1$" is to be determined. The predicate this definite description is formed from has clear circumstances and consequences of application, and its inclusion of the specification 'the smallest' ensures that if that predicate applies to any natural number, it is to a unique one. (In this way it can be contrasted with

the predicates 'largest possible fraction' or 'most rapidly converging sequence', which Frege considers.) What does it mean to say in addition that in fact there is a number to which the predicate applies—that the smallest natural number such that every larger one is the sum of distinct primes of the form $4n + 1$ *exists?* What is it for the definite description not only to *purport* to refer to or pick out a particular number but actually to *succeed* in doing so?

In this case its success consists in the truth of this identity:

> The smallest natural number such that every larger one is the sum of distinct primes of the form $4n + 1 = 121$.

Entitlement to the existential commitment implicit in the use of the definite description, and so entitlement to use that description, can be secured by entitlement to *any* identity of this form. Of course the existential claim is not equivalent to the substitutional commitment that is made explicit by this particular identity—or indeed, any other of this form. The existential commitment is rather equivalent to the disjunctive claim that some identity of this form is true; the significance of the existential commitment is determined by its being incompatible with any claim that is incompatible with all claims of the form of this identity. The significance of such existential commitments is accordingly to be understood, and their propriety assessed, in terms of the class of vindicating substituends supplied by identities of the form of the one above.

What is the relevant form? Not just any identity will do. For instance

the smallest natural number such that every larger one is the sum of distinct primes of the form $4n = 1$	$=$	40 less that the smallest natural number such that every larger one is the sum of distinct primes of the form $6n - 1$

is an identity (substitution license) that does not, like the one above, settle it that the smallest natural number such that every larger one is the sum of distinct primes of the form $4n + 1$ *exists.* It just links that question to the question of whether the smallest natural number such that every larger one is the sum of distinct primes of the form $6n - 1$ exists. If the latter number exists, then so does the former. The claim is, however, that in identifying the first number with 121 (or the second with 161), one is doing more than merely settling this sort of conditional existential question. One is in that case settling the categorical existential question of whether the existential commitments implicit in the use of these definite descriptions are in order, whether those descriptions are successful singular referring expressions, whether the numbers they purport to specify exist.

To say this is to say that the issue of the success of their singular referential purport does not arise for expressions such as '121' and '161' in the same way that it does for expressions such as 'the smallest natural number such

that every larger one is the sum of distinct primes of the form $4n + 1'$. It is to take a frankly inegalitarian approach to referential purport and its success.[60] Numerals are semantically privileged ways of picking out numbers. By contrast to definite descriptions of numbers, the well-formedness of numerals suffices for their referential success, guaranteeing that they pick out a corresponding object. Furthermore, distinct numerals are guaranteed to correspond to distinct objects. According this privileged status to a class of singular terms is treating them as *canonical designators* of a kind of object.

In the paradigmatic case of natural numbers, numerals such as '121' can serve as canonical designators because they are systematic abbreviations of successor numerals: elements of the sequence

$$0, 0', 0'', 0''', 0'''' \ldots .$$

Peano's axiomatization ensures that numerals of this form cannot fail to refer to numbers, and further that their lexical distinctness (marked by the number of successor marks they bear) suffices for the distinctness of the numbers they refer to. Claiming that some noncanonical numerical expression succeeds in referring to a number is just claiming that it is intersubstitutable (saving discursive commitments) with some element in the canonical sequence of successor numerals. Existential commitment with respect to this sort of object, natural numbers, is a disjunctive substitutional commitment linking the candidate numerical expression to some canonical substituend. Saying which number a numerical expression refers to is producing the canonical designator that is intersubstitutable with it. (In this sense one has not yet said how many seconds there are in a century when one asserts the identity

$$\text{the number of seconds in a century} = 100 \cdot 365\frac{1}{4} \cdot 24 \cdot 60 \cdot 60.$$

One has only given a recipe that would make it possible, with some work, to say which number the definite description picks out—a recipe that guarantees at least that it does pick out some definite number.) Once the use of some expressions as canonical designators has been established, Frege's requirement that entitlement to use an expression as a singular term depends on its having been settled which object it refers to in a sense that includes distinguishing that object from others—the requirement that becomes explicit in the introduction conditions he imposes on definite descriptions—is satisfied for numerical expressions by settling it that there is some canonical designator linked to the expression in question by a true recognition statement: a nontrivial identity claim making explicit a simple material substitution-inferential commitment. Ensuring that novel singular terms are suitably substitutionally linked to canonical designators establishes both the existence and the uniqueness of the objects they pick out, and so secures the

success of the singular referential purport that distinguishes them as singular terms.

So the notion of canonical designators makes it possible to think of existential commitment as a kind of substitutional commitment. It can be thought of as a particular quantificational commitment in which the vindicating commitments that determine its content are restricted to canonical designators. It is important to notice, however, how differently such a restriction functions from the sortal restrictions associated with quantificational commitments generally. The sortal restriction is a restriction to kinds of *objects*, while the restriction to canonical designators is a restriction to kinds of *expression.* In substitutional terms, this means that if a sortal qualifies one of an equivalence class of intersubstitutable (that is, coreferential) singular terms, it qualifies all of them. If Kant is a person and Kant is the author of "Was ist Aufklärung?" then the author of "Was ist Aufklärung?" is a person. So the sortal restriction does not discriminate among different ways of referring to the same thing. Not so for the restriction to canonical designators. '9' is a canonical designator of a natural number,[61] and 9 is the number of solar planets, but 'the number of solar planets' is not a canonical designator of a natural number. It is of course a designator of a natural number: 'natural number' specifies a sortal, picks out a kind of thing. But 'canonical designator' picks out a kind of expression, not a kind of thing. So the structure of the restriction on admissible substituends involved in existential commitment is quite different from that involved in ordinary sortally restricted quantification. Existence is not a predicate or property, and existing things are not a *kind* of thing.

5. The Role of Canonical Designators

For a class of singular terms (for instance successor numerals) to have the status of canonical designators with respect to a kind of objects (for instance natural numbers) is a matter of the significance their use has according to the relevant discursive scorekeeping practices. The institution of that significance presupposes, rather than establishes, entitlement to use those expressions as singular terms, however. As Frege would be the first to insist, one cannot simply stipulate that the use of successor numerals as singular terms is in order. Like all expressions, their use must be governed by some nontrivial identities for it to count as the use of expressions as singular terms at all. The point of Frege's disagreement with the formalists is that merely laying down the Peano axioms is not enough to satisfy this requirement. Identities of the form

$$0''' = \text{the successor of } 0'',$$
$$0''' = \text{the successor of the successor of } 0',$$

and so on will not do because such identities are in the relevant sense trivial;

they do not link two different ways of picking out an object because the expressions flanking the identity sign are merely notational variants of each other. Thus they cannot serve as genuine recognition statements; they do not express substantive material substitution-inferential commitments. Frege's response to this situation in the *Grundlagen* is to link the use of the successor numerals to that of other expressions already in use, by means of the method of abstraction. In this way the use of successor numerals is determined by its relation to the process of counting the previously discriminated objects that fall under some sortal. His definition then permits versions of the Peano postulates to be *proven* to hold for expressions introduced in this way. Thus entitlement to the use of successor numerals as singular terms is secured, and they become available for duty as canonical designators. One must be entitled to use expressions as designators first, and only then can they serve as *canonical* designators, which can be appealed to in explaining what existential commitment consists in.

As Frege and others have shown, once one understands existence claims regarding expressions that purport to pick out natural numbers, one can systematically extend that understanding to existence claims regarding rational, real, and complex numbers, and so on. The canonical designators that give sense to the notion of existence for rational numbers can be pairs of successor numerals (corresponding to ratios of natural numbers), for real numbers they can be converging sequences of canonical specifications of rational numbers, for complex numbers pairs of such canonical specifications of reals, and so on. (Notice that in order to be entitled to use a sequence of canonical rational number designators as a canonical designator of a real, one is obliged to *prove* the convergence of the sequence of numbers those designators pick out.) The idea behind the use of successor numerals as canonical designators in explicating existential commitments as a species of substitutional commitment is that to say that some numerical expression succeeds in referring—to say that a number corresponding to it *exists*—is to say that it has some address in the structured space mapped out by the successor numerals. This idea can be carried over, with some differences, to the case of *physical* existence.

To say that some physical object expression succeeds in referring, that the object it designates exists, is to say that it exists *somewhere* in space and time, that it occupies some spatiotemporal region. This is to say that it has some address in the structured space of spatiotemporal coordinates centered on the speaker.[62] The speaker who takes it that Pegasus does *not* (and never did) exist, while P. T. Barnum's elephant Jumbo *does* (or did) is claiming that a continuous spatiotemporal trajectory cannot be traced out connecting the region of space-time occupied by the speaker[63] to one occupied by Pegasus, while such a trajectory can be traced out connecting the speaker's region with that occupied by Jumbo. It is not that Pegasus must be conceived as not taking up any space or surviving for any time; it follows immediately from

his being a horse that he does both. It is that the region he occupies is inaccessible from here and now—"You can't get there from here." He does not exist in *our* space and time, the one that defines physical existence for us. The analog in the case of physical existence of the structured address space defined by the successor numerals is the structured address space defined by egocentric spatiotemporal coordinate descriptions. Thus the term 'Pegasus' is not properly intersubstitutable with any expression of the form 'the (winged) horse located at $\langle x, y, z, t \rangle$ from here, while the term 'Jumbo' is intersubstitutable with an expression of the form 'the elephant located at $\langle x, y, z, t \rangle$ from here'.[64] Thus, like numerical existential commitments, physical existential commitments can be understood as substitutional commitments involving a class of canonical designators (again a kind of *expression*, not a kind of *thing*).

Of course there are also disanalogies between the way the spatiotemporal designators that are canonical for physical existence work and the way the successor numerals that are canonical for numerical existence work. Here-now centered coordinate specifications of accessible spatiotemporal regions are, like successor numerals, guaranteed to succeed in their referential purport. But the canonical designators of physical *objects*, as opposed to the regions they occupy, must include sortal information as well: the statue and the lump of clay *may* occupy just the same spatiotemporal region over the whole course of their existence. The sortals relevant in this case are those where identity (or difference) of the spatiotemporal regions occupied guarantees identity (or difference) of the objects within the sortal (as opposed to across sortals, as in the lump/statue case). Thus if horse$_1$ occupies region r and horse$_2$ occupies region r, then horse$_1$ = horse$_2$ (and if not, not).[65] The individuation of horses is parasitic on spatiotemporal individuation, in that if one has used the horse-specific criteria of application of 'horse' to stick labels only on horses, one then uses spatiotemporal coincidence or divergence to decide how many horses have been labeled and how many labels each horse has.

As in the case of numerical existence, these existential commitments can be understood as substitutional commitments involving physically canonical designators only where one is entitled to apply those canonical designators— which in this case are formed from sortals plus specifications of accessible space-time regions. In the case of successor numerals, this could be done wholesale—'producing' objects for them to refer to by abstraction, so that the Peano postulates could be proven. As in the case of using converging sequences of canonically designated rational numbers as canonical designators of real numbers, however, not only must one pick out a privileged general form as canonical, one must also settle which of the designators of that form are suitable for endorsement as canonical. Just as one must prove the convergence of each sequence of rational numbers that is put forward as a substituend that could vindicate a commitment regarding the existence of a

real number, so one must show in the physical case for each sortal-plus-region pair that the sortal properly applies to the region—that the region specified is occupied by an elephant. The variety of spatiotemporally individuating sortals means that there is nothing useful and general to say about how one becomes entitled to claims applying a sortal to a region. The appropriate circumstances of application for applying the sortal-derived predicate '. . . is an elephant', or '. . . is occupied by an elephant' to a particular space-time region are quite different from those of '. . . is (or is occupied by) an electromagnetic force field'. But these details concern the use of these particular sortals and predicates, not the notion of existence in general. The surplus significance of a commitment to physical existence lies in the *accessibility* to the one undertaking the commitment (via a continuous trajectory from here-now) of a spatiotemporal region to which the sortal (or its derived predicate) is properly applicable. For that reason appealing to the notion of a predicate or sortal being applicable to a region does not make this way of thinking about physical existence circular.

As a final example, *fictional* existence, existence in or according to a story, can be understood as having the same shape as that common to physical existence and the various sorts of numerical existence. To say that in or according to the Sherlock Holmes stories Holmes's housekeeper exists (or that the expression 'Holmes's housekeeper' succeeds in referring to an individual) is to say that that expression is intersubstitutable with some singular term that actually appears in the story (in this case a tokening of 'Mrs. Hudson'). The singular terms that appear in the text that defines the fictional context can be considered as the canonical designators. Thus the claim that according to those stories Holmes's archenemy exists but his fairy godmother does not involves undertaking a substitutional commitment regarding a canonical designator in the first instance, and a commitment incompatible with it in the second. Even if the phrase 'Holmes's archenemy' does not ever appear in the text, the fact that 'Professor Moriarty' does occur there and that it can be deduced from what is said about him that this term is intersubstitutable with 'Holmes's archenemy' ensures that 'Professor Moriarty' is a canonical designator that can vindicate the substitutional existential commitment. That no such term plays this role with respect to 'Holmes's fairy godmother' is the thrust of the denial of even fictional existence to that creature.

For some purposes it is useful to consider as canonical designators not only expressions that actually occur in the stories but also those, like 'Sherlock and Mycroft Holmes's maternal grandmother', whose applicability is entailed by what is said there, though they are never actually used. These boundaries are hazy, for it is not clear what auxiliary hypotheses one is entitled to appeal to in extracting the consequences of what we are told in the story. In most settings regularities of nature, even if not explicitly mentioned, seem safely carried over, but beyond that the matter seems one for decision rather than

discovery. To say this is just to say that the notion of fictional existence is itself hazy outside those objects actually mentioned in the text in question. Another regard in which fictional existence is ontologically indeterminate as far as singular reference goes is that a difference in canonical designators need *not* here entail a difference of objects. The story may simply be silent (even by implication) on the subject of whether the person who had last hired the hansom cab Holmes is riding in at a certain point was Mrs. Hudson (or Sherlock and Mycroft Holmes's maternal grandmother) or not. One can constrain such individuative issues by importing the physical spatiotemporal individuating apparatus into the fictional context; Pegasus was spatiotemporally accessible to Perseus, according to the story. Like the invocation of regularities of nature, these constraints only go so far, and certainly in typical cases far underdetermine the identity and individuation of the fictional objects referred to by canonical designators.

The point of this discussion does not reside in the particular choices that have been offered here as candidates for sets of canonical designators corresponding to different sorts of existence. It is that existential commitments can be understood as a special kind of substitutional commitment (akin to but distinct from sortally restricted particular quantificational commitments) by using the notion of a privileged set of substituends playing the special substitutional role of canonical designators. What one is doing in claiming that the largest number that is not the sum of the squares of distinct primes exists (its canonical designation is '17,163') is different in specific ways from what one is doing in claiming that the tallest edifice in Washington, D.C., exists; numerical existence is different from physical existence. But these different sorts of existence, or even senses of exist, have a structure in common that qualifies them both as notions of existence. It is that common structure that the notion of existential commitments as substitutional commitments restricted to canonical designators seeks to capture.

If something like this account of existential commitments is right, then kinds of existence are to be individuated by the associated sets of canonical designators. Picking different sets of canonical designators gives different senses of existence. So one who treats specifications of real numbers by *pairs* of sequences of specifications of rationals such that sufficiently late elements of one sequence are arbitrarily close to those of the other would, strictly speaking, mean something different by saying that a certain real number *exists*. Of course, as long as each canonical designator in the one set is intersubstitutable (coreferential) with a designator in the other, this difference would not make a difference. It would be significant only when, perhaps against all expectation, apparently equivalent specifications turn out to diverge in hitherto unconsidered cases.[66]

A set of canonical designators has a *structure*—paradigmatically that of the successor numerals or coordinatized specifications of spatiotemporal regions—which systematically provides *addresses* for all the existing objects

of the class in question. As was indicated above, this need not be a totality that is present in the language at any given moment; the structure may provide only recipes for producing suitable canonical designators, or even just criteria for recognizing them, as is the case for some definitions of real numbers in terms of converging sequences of rational numbers. It remains to say what it is to treat such a set of designators as privileged in the particular way necessary for them to be functioning as canonical designators defining a sort of existential commitment. This is a matter of the consequences of undertaking and attributing existential commitments. One of the consequences, of course, is that being entitled to an existential commitment regarding a definite description is a necessary condition of being entitled to use that description at all. Only scorekeepers who attribute an entitled commitment to the existence of some x such that Fx to an interlocutor take the use of a definite description of the form $!x(Fx)$ by that interlocutor to be appropriate.

More important, anyone who does *not* undertake a commitment to the physical existence of the object referred to by a term t cannot endorse any claims of the form Pt, where P is a physical predicate. Under these circumstances one can be entitled only to endorsements of the fictional, 'according to the story' truth of a claim, which differ from endorsements that are not in the fictional mode—for instance in the unavailability as premises for practical reasoning of the claims that appear inside the scope of the explicitating 'according to the story' operator. Similarly, one does not endorse numerical claims whose expression essentially involves terms with respect to which one is not prepared to undertake numerical existential commitments, and one does not endorse fictional claims whose expression essentially involves terms with respect to which one is not prepared to undertake fictional existential claims. It is this intimate connection between existential and doxastic commitments that led Frege to forbid nonreferring terms in his ideal languages.

Picking out the set of descriptions of accessible spatiotemporal regions as playing the role of canonical designators with respect to claims of physical existence includes offering a gloss on what someone is saying who denies that Pegasus (or any winged horse) existed, or ever flew over Greece. Negative existential judgments—claims to the effect that something or other does not exist—have been the source of considerable philosophical confusion over the years. (A paradigmatic example is the doctrine that nonexistent entities must at least have some sort of ⁵subsistence⁵ for us to be able to refer to them in denying that they exist.) But understanding existential claims as expressing substitutional commitments with respect to a class of canonical designators yields a straightforward reading of the sense of such negative existentials. Denying numerical existence to the largest integer, or physical existence to Bellerophon's flying horse, is committing oneself to something incompatible with *all* of the identities, one side of which is a canonical designator of the

relevant sort and the other side of which is the expression 'the largest integer' or 'Bellerophon's flying horse'—or expressions anaphorically dependent on them.

On this relaxed account, there is no reason to boggle at claims that numbers or other abstract objects exist. One must insist only that a determinate sense have been given to such claims, by specifying the relevant class of canonical designators. Once it has been settled that a class of expressions functions as singular terms, if some of the sentences in which they have primary occurrences are true, then it is a criterion of adequacy on the specification of a class of terms as canonical designators relative to such claims that there be canonical substituends for them. Where those conditions are satisfied, corresponding existential claims have been given definite sense and are themselves true. (Of course, how *interesting* they are is another matter.)[67]

III. SUBSTITUTION, TOKEN RECURRENCE, AND ANAPHORA

1. *Inference, Substitution, and Anaphora*

Discursive practice has at its center the game of giving and asking for reasons; social deontic scorekeeping practices confer specifically *propositional* contents because they are *inferentially* articulated. The previous chapter examined the *substitutional* fine structure discernible within that inferential articulation—the substructure in virtue of which subsentential expressions can play an indirectly inferential role by making a systematic contribution to the propriety of inferences in which the sentences they occur in serve as premises or conclusions. That claims are articulated according to substitution inferences in turn presupposes a further level of structure. For substitution is not definable for individual unrepeatable expression tokenings. It requires some notion of token *repeatability.*

To begin with, in order to be available as reasons, sentences used to express claims must be at least in principle repeatable—both within and across interlocutors. Furthermore, for an expression to be used as a singular term (or predicate), it must be possible for it to occur in different sentences, combined with different predicates (or terms).[68] The definition of substitution inferences requires that occurrences of the same expression, whether term or predicate, be identifiable in both premise and conclusion. The first section of this chapter showed how Frege's triangulation principle (according to which purporting to pick out a definite object depends on its having been settled what would count as recognizing that object as the same again) can be understood in terms of the substitutional construal of what it is to use an expression as a singular term. That discussion simply presupposed the availability of repeatable terms.

When analytic focus is sharpened from repeatables such as 'Benjamin Franklin' and 'the inventor of bifocals' to particular tokenings of such types

as 'this man' and 'it', however, more must be said about how such unrepeat-
ables are grouped together into term repeatables. The structure that emerges
as crucial to generalizing substitutional considerations so as to encompass
expressions of this sort is just that appealed to in Chapter 5 in explaining the
use of the traditional semantic vocabulary 'true' and 'refers'—namely *anaph-
ora*. The rest of this chapter investigates this phenomenon, adding a final
level of semantic analysis to the two considered already. The result is a
tripartite theoretical semantic structure whose key concepts are *inference,
substitution*, and *anaphora*.

2. Substitution and Repeatability

What sort of relations do the repeatable terms, predicates, and
sentences that have been under discussion thus far stand in to the unrepeat-
able tokenings that instantiate them in the actual performances that are
accorded pragmatic significance by discursive scorekeepers? What is it for
two tokenings to be occurrences of the *same* term or sentence? Up to this
point the examples considered have done nothing to discourage the supposi-
tion that those repeatables are just lexical *types*: equivalence classes of lexi-
cally cotypical tokenings. But cotypicality is neither necessary nor sufficient
for tokenings to be occurrences of the same term in the semantically relevant
sense of 'same term'. It is not sufficient because cotypicality cannot guaran-
tee the correctness of substitution inferences of the form:

> This organism is a mammal,
> therefore this organism is a vertebrate.

For the two tokenings of the type ⟨this organism⟩ may involve different
demonstrations, and so be governed by different material substitution-infer-
ential commitments. In that case they are not guaranteed to be coreferential
(= intersubstitutable), as different (primary) occurrences of a single term are.

Cotypicality is not necessary for tokenings to be occurrences of the same
term in the semantically relevant sense of 'same term' because the relation
between the predicates *does* guarantee the correctness of substitution infer-
ences of the form:

> This organism is a mammal,
> therefore it is a vertebrate.

For the tokening of the pronoun is guaranteed to be governed by the very
same material substitution-inferential commitments as the tokening that is
its anaphoric antecedent. In the sense required by the discussion of the
previous chapter, the latter of these is a proper substitution inference, and
the former is not. The latter involves what are in the semantically relevant
sense two occurrences of the same term, while the former does not.

These examples emphasize that the idea of substitution inferences (and

hence the assimilation of sentences as substitutional variants of one another)
presupposes that repeatable expressions can *recur*—occurring now in one
context, now in another. This makes no sense applied directly to unrepeat-
able tokenings. One might, of course, literally cut a token out of one written
context and paste it into another, but this is just the sort of case that points
to the need to focus on token*ings* or particular unrepeatable *uses* of tokens,
rather than on the tokens themselves. For the same token may be used to
perform various different speech acts at different times. One sign saying "Dig
here" may be moved from place to place on the campus during a treasure
hunt; the different tokenings of 'here' have different referents (for instance,
they are intersubstitutable with different definite descriptions of locations)
at different times, even though just one token is involved. Token recurrence
may be determined by cotypicality, but that is not the only structure of token
recurrence. Indeed, as will emerge, if what it is for a term to have a cotypi-
cality recurrence structure is for *all* tokenings of a given type to be tokenings
of that same term, and hence be guaranteed to be coreferential, then *no* sort
of expression has such a recurrence structure—not even proper names or
definite descriptions.

Practical acknowledgment or attribution of expression recurrence (treating
some tokenings as tokenings of the same term or sentence) is an attitude that
is implicit in the adoption of substitutional, and hence of inferential com-
mitments. A particularly vivid reminder of the implicit presupposition of the
reidentifiability of *terms* that stands behind the substitutional reidentifiabil-
ity (and hence identifiability) of *objects* is provided by consideration of the
way identity claims work in making substitutional commitments explicit.
Identity locutions make explicit the claim that two terms pick out the same
object. Their defining use is accordingly in explicit substitution inferences of
the form:

$$\Phi a,$$
$$a = b$$
$$\text{therefore } \Phi b.$$

The correctness of an inference of this form depends on the tokening of
the type $\langle a \rangle$ occurring in the first premise and the tokening of that type
occurring in the second premise being tokenings of the same term, and
similarly for the tokening of type $\langle b \rangle$ occurring in the second premise and in
the conclusion. Clearly identity locutions cannot be used to make *this* sort
of implicit presupposition explicit in turn, on pain of an infinite regress. For
the use of identity locutions to license substitutions presupposes the possi-
bility of tokening recurrence between elements of identity claims and ele-
ments of the sentences substitutionally governed by them. Of course, like
the substitutional commitments made explicit by identity locutions, com-
mitments to treating some tokenings as recurrences of others can be made
explicit if suitable further vocabulary is introduced. If the two premises and

the conclusion of the inference above are denominated (i), (ii), and (iii), respectively, then the metalinguistic apparatus for talking about tokenings that was introduced in Section IV of Chapter 5 can be used to do just that.

The first tokening of type ⟨a⟩ can then be picked out as $/a/_i$, the second as $/a/_{ii}$, and the claim that they are tokenings of the same term (not just of the same type) expressed by saying "$/a/_{ii}$ is a recurrence of $/a/_i$," abbreviated as Recur($/a/_i$, $/a/_{ii}$). The reason the previous substitution inference is a good one is then that Recur(/This organism/$_{premise}$, /it/$_{conclusion}$), and the reason the prior inference need not be a good one is that it can happen that ~Recur(/This organism/$_{premise}$, /This organism/$_{conclusion}$). So it is possible to express the sort of inference that is licensed by explicit identity claims by saying that the inference from

$$\Phi a, \quad \text{and}$$
$$a = b, \quad \text{to}$$
$$\Phi b$$

is a good one provided that Recur($/a/_i$, $/a/_{ii}$) and Recur($/b/_{ii}$, $/b/_{iii}$). Of course, such a further explicitation is of use only insofar as one can implicitly recognize the different occurrences of 'a', '$/a/_i$', and so on as recurrences of the same expressions. It is not in principle possible to use explicit stipulations to eliminate the need for reliance on implicit capacities to recognize recurrences.[69] For this reason, an implicit token recurrence relation is appealed to in what follows (much as an implicit substitutional variation relation was appealed to in the previous chapter).

3. Token Recurrence

There are two varieties of substitutional equivalence. These are *intra*term and *inter*term, or de jure and de facto equivalences of tokenings. The former are (taken to be) binding on all interlocutors; the latter vary from doxastic repertoire to doxastic repertoire, according to the particular substitutional commitments undertaken by or attributed to an individual. Each attributor takes recurrence to bind all, in keeping track of significances of identificatory commitments and invocations of them by term use. But the identificatory or substitutional commitments themselves vary from individual to individual.

Substitutional structure requires both sorts. They cannot be defined separately, apart from their role in such a structure; one cannot have the one sort of equivalence without the other. What the intraterm equivalences are *for* is to be the vehicles of interterm substitutional commitments. These latter in turn presuppose them, in that they could not otherwise have content. The model of invariance of something under substitution involves changing something and preserving something else. The changes that do preserve the appropriate something define interterm equivalences.

The notion of recurrence or repetition without change that is presupposed by (is part of, as complementary to) the notion of change invoked in that definition is intraterm or de jure intersubstitutable equivalence. The fact that tokenings can have deictic and anaphoric significances means that the notion of recurrence presupposed by substitutional (and so inferential) relations cannot be reduced to that of lexical cotypicality. This raises three fundamental questions. First, once this issue has been distinguished from that of their being of the same lexical type, what does it mean to say that two tokenings are tokenings of the *same* term or sentence? Second, granting that expressions may exhibit a recurrence structure that can be represented by an equivalence class of cotypical tokenings, what *other* sorts of recurrence structure are there? Third, what difference does it make to the expressive power and function of a repeatable term (or expression of another grammatical category) *which* sort of recurrence structure it exhibits? That is, what sort of expressive impoverishment does an idiom suffer from if it lacks one or another of the different sorts of recurrence structure?

The terms 'term' and 'sentence' are usually thought of as picking out items that are lexically individuated—that is, by character rather than content, in Kaplan's typology. The *semantically* relevant recurrence classes of tokenings correspond not to this sort of repeatable but to what particular tokenings *express.* At the propositional level, it is possible to talk about different sentence tokenings that (in different contexts) express the same *claim.* It is in this sense, of shared content rather than character, that tokenings of "You are tired," uttered by me, and "I am tired," uttered by you, can under suitable circumstances be said to express the same claim. In the next chapter, where same-claiming is investigated in more detail in connection with the use of the 'that' clauses that specify the propositional contents of explicit ascriptions of propositional attitude, this semantic equivalence relation over sentence tokenings is grounded in practical proprieties of *deference* (see 8.4.3). This is a central structure of giving and asking for reasons as a *social* enterprise—of assertion as making premises available for *others* to use in inferences.

The semantically relevant recurrence classes of term tokenings considered here similarly correspond not to the characters of terms as lexical items but to the contents their tokenings express. Awkwardly, there is in ordinary philosophical parlance no generally agreed-upon expression that stands to 'term' as 'claim' or 'propositional content' stands to 'sentence'. 'Individual concept' and 'singular sense' are sometimes used, but each carries theoretical baggage it is best not to import here. The more neutral notion of the *conceptual contents* expressed by term use will be used here to indicate what is wanted. The use of such substantives is misleading in the present context in any case, for it suggests that what links the tokenings that make up the recurrence class of a term is their *shared* relation to some one content. This presupposes that recurrence classes are equivalence classes, and that is im-

portantly not always the case. It is best instead to look at how recurrence classes work and then to adjust talk of contents to fit their expressive function.

For an expression to be used as a singular term is for the inferential potential of sentences containing it to be determined in part by a set of material substitution-inferential commitments linking that term to others. This principle obviously applies to *repeatable* expressions, the tokenings of which can recur from sentence to sentence. Sharpening the focus to the level of unrepeatable tokenings requires that the idea of one term appearing in different sentences—presupposed by its being linkable to others by substitution-inferential commitments—be translated in terms of tokenings that count as recurrences of other tokenings. An unrepeatable tokening can occur in only one sentence tokening. So what it is for an unrepeatable *tokening* to have the significance of an occurrence of a singular term must be cashed out in terms of the inferential potential of the sentence tokening in which it occurs, and of other sentence tokenings in which tokenings occur that qualify as recurrences of the original tokening. For two tokenings $/t/_i$, $/t'/_j$ to be linked by a substitution-inferential commitment means that if $/\Phi t/_k$ is any sentence tokening containing a singular term tokening $/t/_k$ that is a recurrence of $/t/_i$, what it expresses has an inferential consequence that can be expressed by a sentence tokening $/\Phi t'/_l$, where $/t'/_l$ is a recurrence of $/t'/_j$. Thus substitution-inferential commitments should be thought of as linking tokening recurrence structures. For a tokening to be used or treated as an occurrence of a singular term is accordingly for its significance to be determined by SMSICs relating the recurrence structure to which it belongs to the recurrence structures to which other tokenings belong.

These recurrence structures *may* be equivalence classes of term tokenings all of which share a single lexical type. In that case a substitution-inferential commitment linking an equivalence class of tokenings of the type ⟨Benjamin Franklin⟩ and an equivalence class of tokenings of the type ⟨the inventor of bifocals⟩ licenses inferences from what is expressed by tokenings of the form $/\Phi⟨$Benjamin Franklin$⟩/_i$ to what is expressed by tokenings of the form $/\Phi⟨$the inventor of bifocals$⟩/_j$, and things look just the way one would have expected from the discussion in the previous chapter. But recurrence structures not only need not be restricted to lexically cotypical tokenings; they need not be equivalence classes at all. Given the role that recurrence structures of tokenings play in defining substitution-inferential relations, recurrence must be a reflexive relation; each tokening must (trivially) count as a recurrence of itself. After all, the significance of each tokening is guaranteed in advance to be governed by the same set of substitutional commitments that it is itself governed by. More substantively, it must be a transitive relation: if $/t/_k$ is a recurrence of $/t/_j$ and $/t/_j$ is a recurrence of $/t/_i$, then $/t/_k$ is a recurrence of $/t/_i$. For from the point of view of substitution, to say that $/t/_k$ is a recurrence of $/t/_j$ is to say that $/t/_k$ inherits the SMSICs determining its significance from

/t/ⱼ. If /t/ⱼ in turn inherits the SMSICs determining its significance from /t/ᵢ, then /t/ₖ inherits the SMSICs determining its significance ultimately from /t/ᵢ.

But it does not follow from the role of recurrence relations in defining substitution-inferential commitments that recurrence must be a *symmetric* relation. /t/ⱼ may inherit the SMSICs determining its significance from /t/ᵢ without its being the case therefore that /t/ᵢ inherits the SMSICs determining its significance from /t/ⱼ. If the inheritance runs in either direction, both tokenings have their significance determined in the end by the *same* SMSICs. Nonetheless, as will become clear, the direction of inheritance can make a difference in counterfactual situations: if the SMSICs governing /t/ᵢ *were* different, so *would* those governing /t/ⱼ be, but not in the same sense vice versa.[70] Where the recurrence structure is not symmetric, it has the form of a *chain* or *tree*, rather than of an equivalence class.

4. Anaphoric Recurrence Structures

This asymmetric structure of recurrence (inheritance by one token of the substitution-inferential potential of another) is *anaphora*. For one tokening to be anaphorically dependent on another is for it to inherit from that antecedent the substitution-inferential commitments that determine the significance of its occurrence. Different interlocutors may disagree about what those commitments are, but to take it that there is an anaphoric connection is to take it that the use of the anaphoric dependent is correctly evaluated according to whatever substitutional commitments govern the use of its antecedent. One cannot settle the substitutional commitments determining the significance of the occurrence of a singular-term tokening without computing the recurrence class to which it belongs.[71] If one takes a certain tokening to be an anaphoric dependent, attributing a determinate substitutional significance to it accordingly requires identifying some other tokening as its antecedent.

Thus

> (p) Carlyle wrote his brilliant satire of Hegel, *Sartor Resartus*,
>   in part to show that he was an important thinker

has two readings, depending on whether /he/ₚ is taken to be anaphorically dependent on the tokening /Carlyle/ₚ or on the tokening /Hegel/ₚ. Those who understand the claim as involving the latter anaphoric commitment thereby take it to entail and be entailed by its substitutional variant

> (p′) Carlyle wrote his brilliant satire of Hegel, *Sartor Resartus*,
>   in part to show that Hegel was an important thinker,

for both of these tokenings of ⟨Hegel⟩ are, as /he/ₚ is on this reading, recurrences of /Hegel/ₚ. It does not follow from this account, however, that dox-

astic commitment is always preserved by "replacing pronouns by their antecedents."

Whether it is or not depends precisely on how that phrase is understood. If "replacing a pronoun by its antecedent" means generating a substitutional variant in which another tokening of the same *type* as the antecedent is put in the place originally occupied by the pronoun, the principle is false. It is true only if it means generating a substitutional variant in which a tokening that is a *recurrence* of the antecedent is put in the place originally occupied by the pronoun. Thus

> (q) An influential British author wrote his brilliant satire of Hegel, *Sartor Resartus*, in part to show that he was an important thinker

is not, even on the first reading of the anaphoric commitment it involves, equivalent to

> (q') An influential British author wrote his brilliant satire of Hegel, *Sartor Resartus*, in part to show that an influential British author was an important thinker.

For the latter, but not the former is entailed by

> (q") An influential British author wrote his brilliant satire of Hegel, *Sartor Resartus*, in part to show that his friend John Stuart Mill was an important thinker.

In light of the discussion (in Chapter 5, Section IV) of definitization transformations in connection with Chastain's similar examples, it is clear that what is wanted is rather

> (q''') An influential British author wrote his brilliant satire of Hegel, *Sartor Resartus*, in part to show that *the* influential British author [or *that* author] was an important thinker.

Though grammatically a definite description, ⟨the influential British author⟩ is not semantically a definite description, failing as it does of unique designation; and /the influential British author/$_{q'''}$ should be understood as anaphorically dependent on, and hence as a recurrence of /an influential British author/$_{q'''}$ (itself a tokening of a type other tokenings of which function quantificationally, rather than as singular terms). Similar remarks apply to the variant in which a tokening of ⟨that author⟩ is used as the dependent (again a tokening of a type other tokenings of which function differently, in this case deictically rather than anaphorically).

According to this account, understanding one tokening as anaphorically dependent on another is attributing (or in one's own case, acknowledging) a certain kind of commitment. In the language employed here, it is a commitment to the dependent tokening being a *recurrence* of the antecedent tokening. Recurrence commitments of this sort can be understood in terms of the

inheritance of substitutional commitments, which has already been discussed.[72] It should be emphasized that this is an account of what anaphoric relations consist in—or better, given the methodological phenomenalism about normative statuses that governs the theoretical idiom employed here, about the practical attitudes that constitute *taking* or *treating* two expressions as anaphorically linked. It does not pretend to address the questions about anaphora that linguists and cognitive psychologists have been most concerned with—namely, when it is *correct* to adopt this attitude and treat one expression rather than another as the anaphoric antecedent of another, what *lexical* or *syntactic cues* there are for adopting this attitude, or how audiences *in fact* go about deciding which of various possible readings to adopt. The question of interest here is what it is to do the trick—what counts as doing that trick—rather than when it is called for or how it can be brought off.[73]

The basic claim is that tokenings can have their recurrence classes determined in different ways. The recipe for calculating the recurrence class may be one that looks first to the lexical *type* of the tokening, including as basic recurrences other possible tokenings of the same type. Anaphora, by contrast, is a way of computing the recurrence class of a tokening by reference to another *tokening*. It is from that tokening that the anaphoric dependent inherits or borrows its recurrence class (which in turn determines the substitutional commitments that determine its significance). The first recurrence mechanism is *symmetric*. It results in a core class of cotypical tokenings, each of which is related to the others in the same way: by means of their type similarity. The second is *asymmetric*. It depends on the distinct roles played by the anaphoric antecedent and the anaphoric dependent tokenings. The distinction crucial to understanding the way in which the notion of *substitutional* structure makes it possible to extend the notion of *inferential* significance from the sentential to the subsentential level turned out to be the distinction between the *symmetric* substitutional commitments that govern the use of singular terms and the *asymmetric* substitutional commitments that govern the use of predicates. Just so the distinction crucial to understanding the way in which the notion of *token-recurrence* structure makes it possible to extend the notion of *substitutional* commitment from the level of repeatable types to that of unrepeatable tokenings turns out to be the distinction between the *symmetric* recurrence structure governing the use of expressions such as (some occurrences of) genuine definite descriptions and the *asymmetric* recurrence structure governing the use of anaphoric dependents.

5. The Significance of Asymmetric Recurrence

The asymmetry of recurrence reflected in the fact that interchanging the expressions playing the roles of anaphoric antecedent and dependent in general preserves neither the identity of claims nor doxastic commitment

pays substantial expressive dividends. Discursive practices that lack this structure are expressively impoverished in a variety of ways. The two most important concern the *empirical* dimension of discourse and its social or *communicational* dimension. Under the first heading, it is because they encompass noninferential reporting practices (with their attendant reliability structure of authority, as introduced above in 4.2–3) that languages can express empirical claims. Indexical or token-reflexive constructions, particularly deictic or demonstrative ones, play an essential role in such reports. But no language can have such constructions unless it also has asymmetric recurrence structures: *deixis presupposes anaphora.*

Under the second heading, the capacity to pick up another interlocutor's reference by using a pronoun is one of the central mechanisms by which communication is secured across the interpersonal gap created by differences in doxastic commitments (which induce differences in the inferential circumstances and consequences of application expressions are taken to have by different interlocutors). The significance of anaphora in this context is that it permits each interlocutor to produce utterances employing tokenings that have been stipulated to be recurrences of arbitrary tokenings by others. At both the level of sentences and the level of terms, such recurrence provides the basic points of contact between different repertoires of commitments (including inferential ones). The rest of this chapter is concerned with exploring these two sorts of expressive capacity that anaphora bestows.

The definition of anaphoric dependence allows many tokenings to be (treated as) anaphorically dependent upon the same antecedent. As the mention of the significance of anaphoric links among tokenings for securing communication among interlocutors suggests, the antecedent tokening may be uttered by someone other than the one who produces the tokenings that are anaphorically dependent upon it.[74] It is also possible for that antecedent itself to be anaphorically dependent on some prior antecedent. Since recurrence and inheritance of substitutional commitments is transitive, so is anaphoric dependence. It is in this way that anaphoric chains or trees are formed. They can be anchored or initiated by tokenings that are not themselves anaphorically dependent on other tokenings. These are *anaphoric initiators.*

This role can be played by expressions of two different kinds. Tokenings that acquire their substitutional significance symmetrically (recurring by type), such as some uses of definite descriptions and proper names, can serve as initiators; they are potential anaphoric antecedents that need not themselves in turn be anaphoric dependents. Also, indexical and demonstrative tokenings—which acquire their recurrent tokenings asymmetrically in the form of anaphoric dependents, without themselves being recurrences of any prior tokenings—can serve as initiators; indeed, if their occurrence is not to be cognitively and semantically idle, they *must* so serve. Being an anaphoric antecedent or dependent is a role that individual tokenings can play. Though

their lexical and syntactic types can suit them for those roles, tokenings are not *compelled* to adopt one or the other role by those types. It is crucial to the communicative function of anaphora that any term tokening (or sentence tokening—indeed any tokening that is not syncategorematic) whatsoever can function as an anaphoric antecedent.

Less obviously, with the exception of indefinite descriptions, there seem to be no expression types that preclude their tokenings from functioning as anaphoric dependents. That is, there are no other types *all* of whose tokenings function as anaphoric initiators—expressions *none* of whose tokenings can have the significance of anaphoric dependents. Definite descriptions, for instance, can be used anaphorically—typically as the result of definitization transformations of indefinite descriptions. Demonstrative constructions can also be used to form anaphoric dependents, as in:

> Kant admired Rousseau, but that writer admired only himself

and

> Fichte fought for political tolerance in Germany, and partly because of his efforts this precious end was eventually achieved.

Anaphorically dependent uses of proper names are discussed below in 8.5.6. Since indefinite descriptions always also have quantificational and predicative uses (as in "Carlyle was an influential British author"), there are *no* expression types *all* the tokenings of which are recurrences of one another—or indeed (for that reason) even all coreferential. There can, of course, be sets of cotypical tokenings all of which are recurrences of each other (and hence coreferential). But they all belong to recurrence structures that can also contain anaphoric recurrences of other lexical types, and they all exclude some (possible) tokenings of that type.

IV. DEIXIS AND ANAPHORA

1. Demonstratives

Substitutional commitments relate token repeatables. So unrepeatable tokenings must be sorted accordingly as some count as recurrences of others in order for any of them to have the sort of indirectly inferential significance in virtue of which their production can contribute to making a move in the language game. That sorting may be based only on the lexical or syntactic type instantiated by the tokenings, or it may be based on further features of the individual tokening. Within tokening-based approaches to explaining the significance in discursive practice of expressions and complex components of expressions, the two central phenomena appear to be anaphora and deixis.

In recent years demonstrative and indexical constructions have received a

lot of attention. Perry and Lewis and others have shown that such expressions are not eliminable in favor of expressions that recur by type, such as standard uses of definite descriptions (those that do not themselves contain indexical elements).[75] Evans has provided an elegant account of how the capacity to have demonstrative thoughts and the practical capacity to locate one's actions and perceptions simultaneously in egocentric and public space and time are interdependent.[76] The result of these discussions (and of the many others addressing the same topic) is a robust sense that deictic tokenings provide expressive resources that are essential to our conception of ourselves as empirically situated knowers and agents.

Discussions of the use of token-reflexive or indexical expressions usually do not include treatments of anaphoric dependency. Indeed, pronouns do not fit well with the paradigms that lead to this terminology for tokenings whose use depends on features other than their types. The idea behind talk of *token reflexiveness* is that what tokenings of types such as 'I' refer to depends on who utters the tokening, for types such as 'now' and 'here' it depends on when and where the tokening was produced, and so on. What a tokening of 'it' refers to indeed depends on features of the tokening itself (just which features depending on the type or character of the pronoun). But the relevant feature of the tokening is just its anaphoric antecedent, and this is not a feature that can be specified independently and in advance of semantic and pragmatic interpretation, the way speaker, time, and place can. (Indeed, whether it is to be counted a semantic or pragmatic feature on traditional ways of dividing up these issues is itself a nice question.) Similarly, thinking of pronouns as *indexical* demands that one be able to specify the index to which their semantic evaluation is relative, and once again the anaphoric relation of a tokening to an antecedent tokening is not happily assimilated to general indices such as speaker, time, and place, which are specifiable in nonsemantic terms. Anaphoric chains running through bits of discourse are not naturalistic features of them like which organism produces the tokening, or when or where it is produced. They are *normative* features attributed to the discourse by deontic scorekeepers, matters of conditional commitment or commitment inheritance—of the obligation that the significance assigned to, or score kept on, one part of the discourse answer in systematic ways to the significance assigned to, or score kept on, another.

It is interesting to consider the use of demonstratives in connection with this contrast. A common strategy is to assimilate demonstratives to the indexical paradigm. The semantic interpretant correctly assigned to a demonstrative tokening depends not only on the lexical type it instantiates but also on a further index, which is taken to function in a way analogous to the way the indices of speaker, time, and place function in the semantic interpretation of tokenings of 'I', 'now', and 'here'. That additional index in the case of demonstratives is the object indicated or demonstrated by the one using the

demonstrative. What sort of an index is this? Is what a speaker is indicating or demonstrating a naturalistic fact, specifiable in nonnormative, nonsemantic terms, as speaker, time, and place are? Or is it a normative, scorekeeping matter of how various commitments should be understood as related to one another, as anaphoric dependence is? Or is the dependence of the significance of demonstratives on what is demonstrated of yet a different sort, not to be assimilated to either sort of relation?

What is it to indicate or demonstrate an object? The idea that animates discussions of demonstratives in terms of the indexical paradigm is that the core phenomenon is *pointing:* the demonstrated object is the one pointed at. This idea encourages a picture of indication as a physical matter of picking out an object by extending the line formed by the knuckles of the index finger out until it intersects something opaque. Of course things are not so simple. Wittgenstein reminds us that even such a practice of pointing requires a great deal of social stage-setting—the untrained may be unable to transfer their attention beyond the tip of the pointing finger, or may perversely trace the line of indication in the wrong direction, from finger tip to base, and so take it that something behind the one pointing has been singled out. Again, he reminds us of the emptiness of 'bare' demonstration. The use of 'this' or 'that' must at least implicitly be connected with some sortal, for the same physical gesture can have the significance of pointing to a book or to its cover, its title, its color, its shape, and so on. Although everyone would concede that for reasons such as these the actual practice of pointing is a complex affair, the picture of the virtual line extending from the finger nonetheless exercises considerable force. It seems to promise a way of picking out objects by causal, rather than conceptual triangulation (though as the point about sortals indicates, one of the intersecting beams is quite hazy in the picture).

But this is a false promise. The one using a tokening of 'that pig' demonstratively in many cases need not *do* anything in order to have indicated or demonstrated one particular pig in the barnyard, provided that the unique salience of that pig has somehow already been established—whether through the efforts of the one employing the demonstrative or not. The requisite sort of salience is a motley; it can consist in distinguishing features of the perceptible environment, properties that are highlighted by their relation to psychological factors, conceptual stereotypes, background beliefs, the previous course of the conversation, and perhaps much else. As with anaphora, it is helpful to put to one side the difficult psychological question of how scorekeepers in fact determine salience and so often correctly discern what object is being indicated or demonstrated, and also the difficult conceptual question of what makes it *correct* to take or treat one object rather than another as the one being demonstrated. For the fundamental question concerns the scorekeeping attitude itself: What is it to *take* or *treat* some object *as* what is being indicated or demonstrated in connection with a demonstrative use

of some expression? What does taking or treating a singular-term tokening as having a demonstrative significance that ties it to a particular object consist in?

The function of demonstration is to pick out an object and attach a tokening to it—to settle that the tokening refers to that object (notice that the second 'that' is functioning anaphorically, not demonstratively in this sentence). In light of the general discussion of picking out or using terms to refer to objects, this means that the effect achieved by successful demonstration must be understood in terms of the *substitutional* role of the demonstrative tokening.[77] But a demonstrative tokening as such is unrepeatable; substitutional commitments govern repeatable expressions—those that can occur in more than one sentential context and that can be replaced in each context by others. Only accidentally and in very special circumstances would a speaker be in a position to repeat a demonstration of an object. Even where repeated demonstration is possible, it results simply in the production of another unrepeatable tokening, one that can be seen not to have the significance of a recurrence of the original by the fact that there is always the possibility that, unbeknownst to its author, it in fact picks out a different object. If demonstrative tokenings could not recur, then they could play no substitutional role, hence no inferential role, and so would be semantically and cognitively idle. They would in that case not be ways of picking out or talking about objects at all, but mere noises.

Of course the recurrence structure in virtue of which demonstrative tokenings can play a conceptual role is not far to seek. Demonstrative tokenings can be picked up anaphorically. Because they can have anaphoric dependents, demonstratives can figure in substitution inferences:

> *That pig* is grunting, so *it* must be happy. I'm glad, because *it* is our champion boar, Wilbur.

Anaphoric chains that the demonstrative initiates are available to figure in substitutional commitments and the inferences they govern just as repeatable term types such as ⟨Wilbur⟩ are. Because they are, uttering the demonstrative can be understood as contributing to making a move in the language game, in particular as indicating an object. It follows that the capacity of pronouns to pick up a reference from an anaphoric antecedent is an essential condition of the capacity of other tokens (which can serve as such antecedents) to have references determined deictically. Deixis presupposes anaphora. No tokens can have the significance of demonstratives unless others have the significance of anaphoric dependents; to use an expression as a demonstrative is to use it as a special kind of anaphoric initiator.

One might choose to assimilate the use of indexicals generally to deictic uses in this regard. I can make a claim by uttering a token of the type ⟨John should be leaving the house *now*⟩, only because I and another can later utter tokens of types such as ⟨If he had left *then*, he would have been at the

meeting on time⟩, in which the tokening of ⟨then⟩ should be understood as anaphorically dependent on the earlier tokening of ⟨now⟩.[78] The claim would be that in such a context 'then' functions much as it does in contexts such as

> Clothes were made differently in *the seventeenth century* than they are these days; *then* they were made by hand, while now they are made by machines.

The tokening /then/ here clearly is used anaphorically, with /the seventeenth century/ as its antecedent. It could be claimed that 'then' is *always* used in this way (though only at the price of allowing virtual or merely possible antecedent tokenings). In the same way, tokenings of 'you' could be understood as anaphorically dependent on (perhaps merely possible) tokenings of 'I' by others.

In these cases, however, by contrast to that of demonstratives, one can understand how recurrences of indexical tokenings are possible without invoking specifically anaphoric connections. The character ⟨I⟩, for instance, is systematically linked to ⟨you⟩ (and ⟨he⟩ or ⟨she⟩), in such a way that tokenings of the one can count as recurrences of tokenings of the other (so that the same content is expressed by their use) without that linkage having to be understood anaphorically. The way in which uses of ⟨here⟩ and ⟨now⟩ can be picked up by uses of ⟨there⟩ and ⟨then⟩ respectively seems to work in a way that is intermediate between the systematic interchange of speaker and audience that governs ⟨I⟩ and ⟨you⟩, on the one hand, and the demonstrative element that is the essence of ⟨this⟩ and ⟨that⟩, on the other. But in these cases, too, index-matching rules connecting the contexts of utterance can be formulated that specify which tokens of the one type count as recurrences of which tokens of the other. Nothing like this, though, is possible in the case of purely deictically significant tokenings. There recurrence can be understood *only* on the anaphoric model.

It has been argued that deixis presupposes anaphora. But if anaphora could be given a deictic analysis, then it would not follow that anaphora is the more fundamental phenomenon. Here one would look for a converse of the familiar cases in which apparently deictic or demonstrative expressions are actually playing the role of anaphoric dependents, as in "Kant was a Prussian Pietist, but *that philosopher* did not always think like one." It is hard to assess the promise of this idea in the absence of a detailed working-out of it. One idea would be to assimilate the relation between an anaphoric dependent and its antecedent to the relation between a demonstrative and what is demonstrated. Anaphoric dependents would be understood as indexical tokenings that referred to their antecedents. In a tokening /Hegel understood Kant's argument, but he did not refute it/$_i$, the token /it/$_i$ would be understood as meaning what a token /that/ would mean, if it could be arranged that what was demonstrated was the antecedent tokening /Kant's argument/$_i$.

But as it stands this cannot be how such a story would go. For in that case the tokening /it/$_i$ would be conceived of as intersubstitutable, not with other tokenings co-identified with /Kant's argument/$_i$, but with tokenings that (presystematically) would be said to refer to the *tokening* /Kant's argument/$_i$—such as /the very tokening of type ⟨Kant's argument⟩ that was just uttered (or tagged with the index i)/. That is not what Hegel understood but failed to refute, for he never heard of that tokening. Some mechanism would be required to get us from the demonstrated anaphoric antecedent as referent of the supposed demonstrative-anaphoric dependent to the referent of that tokening. It should be obvious from the discussion of Chapter 5 what sort of mechanism that is. It is precisely this job that defines some locution as meaning what 'refers' means. An operator that forms indirect definite descriptions of objects from direct definite descriptions (or other specifications) of term tokenings is just what 'refers' is. If the natural anaphoric analysis of such locutions is to be avoided (as pursuing this strategy would require), then an alternate demonstrative or more broadly deictic reading must be offered of 'refers' and locutions performing a cognate function, in order to offer a deictic analysis of other anaphoric locutions. If this is not forthcoming, then anaphora must indeed be seen as presupposed by but not presupposing deixis.

To summarize, then. The recent tradition has focused on the cognitive centrality and irreducible importance of deixis. Concern with "direct reference" has been developed in part by exploiting intuitions about the basic nature of the word-world link established by the demonstrative use of expressions—what we used to think of as the use of demonstrative expressions, before it became apparent that anything can be used demonstratively (just as anything can be used anaphorically). From this point of view it has seemed natural to make the point that linguistic significance is always the significance of a (possible) event or uttering, that is, of a tokening—which may then be connected to semantically relevant types according to various models, by insisting on the priority and irreducibility of unrepeatable demonstrative and indexical uses over descriptive repeatable ones. Anaphora, as another tokening-based phenomenon, has not seemed of the essence, for it deals only with intralinguistic continuations or preservations of something one must in principle be able to understand already without it. Thus if one cared about tokenings as well as types, it seemed natural to care about deixis first, and about anaphora, if at all, only later. It now appears, however, that such an attitude, natural as it seems, is not only strategically wrong-headed but actually incoherent.

2. Deictic Mechanisms Presuppose Anaphoric Mechanisms

Deictic uses presuppose anaphoric ones. One cannot coherently describe a language in which expressions have demonstrative uses but no pronominal uses (although the converse is entirely possible). For indexical

uses generally, like deictic ones, are essentially unrepeatable according to types. Different tokenings of 'this' or 'here' or 'now' are not in general recurrences of each other, or even co-identifiable. Yet it is only as repeatable (that is, as elements of recurrence classes) that they can be substituted for. Recurrence is presupposed by the possibility of substitution, and the possibility of substitution is presupposed by picking out occurrences as semantically significant (that is, as indirectly inferentially and so assertionally significant). Since deictic uses as such are not type-recurrent, that recurrence must be understood as token recurrence—in particular, as anaphoric.

In short, unless one could pick deictic uses up anaphorically to generate recurrence classes, one would not be able to involve such deictic tokenings in (undertaken or attributed) identificatory substitutional commitments, and so could not treat them as involving occurrences of singular terms. Without the possibility of anaphoric extension and connection through recurrence to other tokenings, deictic tokenings can play no significant semantic role, not even a deictic one. Deixis presupposes anaphora. Anaphora is the fundamental phenomenon by means of which a connection is forged between unrepeatable events and repeatable contents. No semantically significant occurrence of a subsentential expression can be discerned unless it is governed by substitution inferences, which requires token recurrence: no (semantically significant) *occurrence* without (the possibility of) *recurrence*.

One consequence of this claim is an argument that is similar in many ways to one of the most important arguments of "Empiricism and the Philosophy of Mind." There Sellars argues that it is not possible to conceive of a language consisting only of noninferential reports. Noninferential responsive reporting does not form an autonomous language fragment, does not constitute a set of practices one could engage in without also engaging in specifically inferential practices. Not every claiming or sentence tokening can be a noninferential reporting, the exercise of a reliable disposition to respond differentially to features of the nonlinguistic environment. For what distinguishes reports from any other response produced according to a reliable responsive disposition is precisely its possession of propositional content, that it means something that can be understood by the responder and by others. Possession of such content is in turn a matter of the inferential significance of the reporting response. Understanding reports requires being able to distinguish what further claims follow from them and what claims would provide evidence for them. This inferential articulation, which is the possession of conceptual content by the responses produced noninferentially, presupposes the possibility of making the inferences involved, drawing the conclusions or offering the justifications indicated. These will not be reporting uses of sentences, since they are arrived at inferentially. So not all claimings can be reportings, if even the noninferential language entries are to have the significance of reportings.

The conceptual dependence of deictic mechanisms on anaphoric ones

suggests a way of extending this argument from the level of sentences to that of singular terms. Just as reportings as events are not autonomously significant (they depend for their semantic significance on the possibility of connection with other claimings by inference), so demonstrations are not autonomously significant (they depend for their semantic significance on the possibility of connection with other term tokenings by anaphora). The same sort of mistake that Sellars diagnosed in those who were tempted to see claimings acquiring their content solely through the circumstances in which they are appropriately produced as noninferential reports or *Konstatierungen*—a content that is then only reflected in (consumed or presupposed by) inference—can then be discerned in those who would assign a similar priority (presupposing the possibility of autonomous significance) to deictic mechanisms of "direct reference" in securing references that are then merely reflected in or preserved by (and so are consumed or presupposed by) anaphora. The phenomenon here is exhibited for terms, but the significant extension of Sellars's point is better thought of as one that moves from lexical-syntactic types to tokenings:

$$\text{inference : anaphora}$$
$$::$$
$$\text{semantic significance of types : semantic significance of tokenings}$$
$$::$$
$$\text{reporting uses of sentences : deictic uses of terms.}$$

Just as it is their potential for inferential involvements that makes sentence repeatables bearers of contents, so it is the potential for anaphoric involvements that makes unrepeatable tokenings bearers of contents.

Conceptual articulation, it was claimed in Chapter 2, is in the first instance, *inferential* articulation. *Sub*sentential expressions, it was claimed in Chapter 6, can be understood as *indirectly* inferentially articulated, and hence as conceptually articulated—in spite of the fact that they cannot play the role of premise or conclusion in inference—in virtue of the significance of their occurrence for *substitution* inferences. Unrepeatable tokenings, paradigmatically demonstratives, can now be seen to be *conceptually* articulated, for they can stand in *anaphoric* relations to other tokenings, and the chains thus formed can be involved in substitutional, and hence inferential, commitments. The use of a demonstrative may be elicited noninferentially as a response to an environing stimulus. What makes it a term referring to an object—rather than a mere conditioned response like "Ouch"—is its role as an anaphoric initiator of chains that can be the subjects of substitutional commitments. It is in virtue of those anaphoric connections that a demonstrative tokening can play a conceptual role.

Equipped with this thought, it is possible to address the question of what it is to take it that some particular object has been demonstrated (whether by an actual gesture or by implication). For now this attitude of taking it that

some particular object has been demonstrated appears just as a special case of taking the use of any expression at all to have picked out a particular object. It is to be understood in terms of the sort of cognitive, conceptual, or semantic triangulation discussed in Section I—recognizing an object as the same again by undertaking a substitutional commitment. Such commitments are made explicit in the form of identities employing expressions that are anaphoric dependents of the demonstrative tokening. The anaphoric relations may be implicit, as in:

> *It* is Wilbur (or the big one),

or they may involve explicit use of what might be called (after the terminology of Chapter 5) anaphorically indirect demonstrative descriptions, as in:

> *The pig he pointed to (indicated, demonstrated) saying 'that pig', is Wilbur (or the big one).*

Both sorts of tokenings, /It/ and /The pig he pointed to (indicated, demonstrated) saying 'that pig'/, should be understood as pronouns, anaphorically dependent on a prior utterance /that pig/.[79] Thus the concept expressed explicitly by locutions such as 'what one is *pointing* to', like what is expressed by locutions such as 'what one is *referring* to', must be understood anaphorically. Thinking of referring in the way that has been developed here (by appeal to the concepts of inference, substitution, and anaphora), one could say that referring cannot be explained in terms of pointing, because pointing must be understood in terms of referring.

Anaphora has been presented here as a kind of token recurrence—a relationship among tokenings that is presupposed by, and hence not analyzable in terms of, substitutional commitments. Taking one individual's tokening to be anaphorically dependent on another is not attributing a substitutional commitment; it is attributing a more primitive sort of commitment, one that determines which substitutional commitments regarding other tokenings are relevant in assessing the substitutional significance of the one treated as anaphorically dependent. It may seem that this is an unnecessary shuffle, that a further level of analysis need not be broached. For anaphorically related tokenings are coreferential, and treating two expressions as coreferential has been explained in terms of the practical deontic scorekeeping attitude of attributing substitutional commitments. The reason such an account will not do has already been indicated: substitutional commitments govern the use of *repeatable* expressions. Anaphora is required to generate repeatables from unrepeatable tokenings, paradigmatically deictic ones, where cotypicality does not carry even a defeasible presumption of coreference, hence not of (co-)recurrence.

Identity claims make substitutional commitments explicit as the contents of possible assertions and judgments. It was pointed out above that the use of these explicitating locutions as inference licenses evidently depends on a

prior notion of term recurrence, since terms occurring on the two sides of the identity must be reidentifiable in the other premise and in the conclusion (respectively) of such inferences. It is only by means of anaphoric recurrence that deictic tokenings become accessible to identity claims at all:

> That cat is watching television.
> It is (=) the most spoiled cat in the room.
> Therefore the most spoiled cat in the room is watching television.

The tokening /it/ here is not replaceable with another tokening of the type ⟨that cat⟩, saving the goodness of the inference. The closest thing along these lines that will do is to replace /it/ with a tokening of the type ⟨that *same* cat⟩ or ⟨that *very* cat⟩. These look like deictic uses but are in fact anaphoric dependents, as tokenings of ⟨the same cat⟩ or ⟨the cat just mentioned (or pointed at)⟩ would be in this context. The function of such locutions is precisely to make available tokenings that are recurrences of the one they (anaphorically) refer to.

3. Rigidity is an Anaphoric Phenomenon

Recognizing that locutions such as 'that very *K*' play the role of operators that form anaphoric dependents sheds light on another aspect of demonstratives: their modal *rigidity*. Kripke introduced the term "rigid designator" to distinguish expressions (such as proper names) that pick out the same individual in all possible worlds, from those (such as definite descriptions) that do not. The pretheoretical phenomenon this theory-laden description addresses is the observation that even if Archie is the most spoiled cat in the room, Archie *might not* have been the most spoiled cat in the room (one even more indulged may just have stepped out for a snack); but it is not possible that Archie not have been Archie. Evaluated with respect to the actual situation, the claim that Archie is the most spoiled cat in the room is true. The coreference (that is, intersubstitutability) of these terms is a commitment actually undertaken by the evaluating scorekeeper. Evaluated with respect to other possible situations, the claim that Archie is the most spoiled cat in the room would not be taken to be true, since that description would then pick out another cat, say Ana; a different set of commitments undertaken on the part of the scorekeeper would entail that Ana, not Archie, is the most spoiled cat in the room. Against either background set of commitments, however, Archie would still be Archie. For the coreference of different tokenings of the proper-name type ⟨Archie⟩ is guaranteed by their being recurrences of one another. Definite descriptions are not rigid designators in that when modal contexts (which make explicit the inferential potentials of expressions when evaluated with respect to diverse sets of background commitments) are taken into account, different tokenings of the same definite description type

are not guaranteed to be coreferential. It follows that at least in these contexts, cotypicality of definite descriptions is not sufficient for tokenings to belong to the same token recurrence structure. Of course, cotypical tokenings of grammatically definite descriptions are not guaranteed to be coreferential in any case, because such descriptions have uses in which they play the role of anaphoric dependents; two tokenings of the type ⟨the author⟩ would not be expected to be coreferential if they were drawn from different anaphoric chains appearing in reviews of different books.

Rigidity is an anaphoric phenomenon. Instead of repeating the proper name, the modal claim above can be expressed by saying:

> Archie is the most spoiled cat in the room, but *he* (or *that cat*) might not have been the most spoiled cat in the room.

The pronoun has as its antecedent, and so is a recurrence of, the tokening /⟨Archie⟩/ that appears in the claim about how things in fact are. The anaphoric chain to which it belongs is then available to specify that same cat in other possible situations. The anaphoric chain, in other words, denotes rigidly.[80]

Kaplan introduces a rigidifying operator 'dthat' with the stipulation that while !xDx is not in general a rigid designator (since its denotation varies from world to world), dthat(!xDx) is to be rigid, picking out in each world whatever !xDx picks out in the actual world.[81] His operator in effect forms a type all of whose tokenings are stipulated to be anaphoric dependents of a tokening in the actual world. Thus 'dthat' does systematically what expressions like 'he' and 'that very cat' do informally in specifying other possible situations involving the same objects that are picked out in the actual situation by contingently associated expressions. The idea that rigidity has something special to do with the use of demonstratives—which lies behind the choice of 'dthat' (a homonym of the demonstrative 'that') to express the rigidifying operator—arises precisely because demonstrative tokenings can recur *only* anaphorically, and hence rigidly.[82] Thus someone who says

> *That bar* is the standard meter stick. If *it* had been heated, *it* would have been more than one meter long.

is using the anaphoric dependent to make the demonstrative recur, thereby displaying the same rigidity that would be made explicit in saying

> *That bar* is the standard meter stick. If *it* had been heated, *Dthat(that bar)* (or *Dthat(the bar the speaker indicated)*) would have been more than one meter long.

Kripke's original discussion pointed to the distinction between the modal rigidity of proper names and the nonrigid behavior of the definite descriptions that often serve to fix the reference of those names.[83] In Kaplan's terms, proper names can be understood as rigidified, 'dthat'ed descriptions or dem-

onstratives. Understanding rigidity in anaphoric terms accordingly suggests that tokenings of proper names be themselves understood as anaphoric dependents—elements in an anaphoric chain that is anchored in some name-introducing tokening.[84] Causal-historical theories of proper names then appear as dark ways of talking about the sorts of anaphoric chains that link tokenings of proper names into recurrence structures. Concern about the nature of baptism or name introduction is concern about how this special sort of anaphoric chain or tree can properly be initiated, so that all the anaphorically linked tokenings it contains pick out one object. Concern about transmission of names is concern about what is involved in earlier tokenings being picked up as anaphoric antecedents by later ones. It is clear in these terms what is happening when different tokenings of the same lexical type are used as names of different people—how there can be more than one person called 'George' or 'Aristotle'. In such cases there are just multiple anaphoric chains; the multiplicity of people who can be referred to as 'George' is a phenomenon to be understood by analogy to the way in which many people can be referred to as 'she'. Investigations of the roles played by social-linguistic context and practices, or conventions, on the one hand, and individual name-user's intentions, on the other, in determining what previous uses a particular tokening ought to be considered beholden to should be understood as investigations of which anaphoric chain a particular tokening ought to be considered to be part of.

The claim is not that assimilating these questions about the use of proper names to scorekeeping questions about recurrence commitments (commitments regarding the inheritance of substitutional commitments) solves at a stroke all the questions that have vexed causal-historical theories. On the contrary, those questions, when transposed into the anaphoric framework, typically address issues about what determines when it is correct to treat one tokening as anaphorically dependent on another. The present discussion pretends to address only the issue of what it is for a scorekeeper to take one tokening to be anaphorically dependent on another, not the specific practices by which such scorekeeping attitudes are assessed as correct or incorrect. Nonetheless, seeing these issues concerning proper names as special cases of general issues concerning anaphoric links pays certain explanatory dividends. . (Some of these are exploited below in 8.5.5–6, in the discussion of Kripke's puzzle regarding the behavior of proper names in ascriptions of propositional attitude.) For instance, it becomes clear that the issue of whether a name user remembers where the name was picked up from or is disposed to defer to some other individual more knowledgeable in its use is relevant to assessing uses of the name only insofar as they bear on the *commitments* that name user has undertaken; one may be obliged (according to a scorekeeper) by the circumstances under which one acquired a name even if one is ignorant or mistaken about them, and may be committed to defer to various authorities without being disposed to do so.

Fans of causal-historical theories of reference do not typically restrict these accounts to proper names. Natural-kind terms and many predicates are thought to function according to this model as well. Indeed, it is important to remember that the opposition between causal and descriptivist theories of reference that Kripke appealed to in his influential arguments against the latter arises specifically in addressing the reference of proper names. No one was ever a descriptivist about the reference of the predicates used to form descriptions—on pain of an obvious regress. Some other mechanism of reference was always envisaged for those predicates.

Thus it is natural to understand expressions such as ". . . is red" and ". . . has a mass of twelve grams" as having the denotations they do in virtue of their links to authoritative episodes of calling things red and measuring their mass in grams. According to the present suggestion, then, these expressions ought also to be understood as functioning anaphorically. Insofar as this is the right way to look at things, then, there is in fact only one primitive recurrence structure, namely the anaphoric one. Apparently type-recurrent expressions such as proper names and basic predicates should in fact be understood as having their tokenings linked by relations of anaphoric dependence. The relations between tokenings in these structures can be symmetric, by contrast to the asymmetry of paradigmatic pronoun-antecedent links, because and insofar as those tokenings owe allegiance to a common antecedent.

Such an approach, according to which *everything* works anaphorically, may seem to explain too much. If all these sorts of expressions have anaphoric recurrence structures, and such structures act rigidly in modal context, what room has been left for expressions that are modally flaccid, varying from context to context in the substitutional commitments determining their significance? Although basic predicates should be understood to recur anaphorically, compounds of them—in particular definite descriptions formed from them—need not. Although the recurrence structures determining the inheritance of substitutional commitments of their basic parts link them to antecedents in the actual world—that is why the answer to Lincoln's question "If we agree to call the tail a 'leg', how many legs would horses have?" is still "Four"—the existential and uniqueness commitments involved in the use of definite descriptions are in modal contexts evaluated with respect to alternative situations.

Archie, the cat in front of the television, does exist; but he, that very cat, might not have—even though another cat much like him might have been named 'Archie' in circumstances much like those in which Archie was baptized, and might be sitting in front of the television. Evaluated with respect to a different set of commitments undertaken by the one keeping score (that is, with a different set of claims being taken-true, and so functioning as the facts), a description might pick out a different object than it does according to the actual commitments of the scorekeeper. For what it is

intersubstitutable with depends on what other substitutional commitments are in play. This is compatible with there being an underlying anaphoric recurrence structure governing the use of the components of the description, which determines which of the substitutional commitments in play are in fact relevant to determining the proper use of the description in question.

The foregoing remarks are not intended to present a theory of the use of demonstratives, proper names, and definite descriptions. They are meant rather to suggest reasons for recasting the standard problems regarding such expressions into the terms of the common framework provided by the notion of anaphoric structures of tokening recurrence. According to the picture being presented, taking someone to have used an anaphorically dependent tokening in making a claim is attributing an *anaphoric commitment*. An anaphoric commitment is a commitment to treating the dependent tokening as a recurrence of the tokening that is taken (by the one attributing the commitment) to be its antecedent. Recurrence commitments, the genus of which *anaphoric* commitments are a species, are commitments regarding the *inheritance* of *substitutional* commitments. In the same way, *inferential* commitments, the genus of which *substitutional* commitments are a species, are commitments regarding the *inheritance* of *doxastic* or assertional commitments.

The story accordingly has three layers. At the top, sentences can be understood as propositionally contentful in virtue of their use in expressing *claims*—that is, assertional commitments. The key concept at this level is *inference*, for what makes the contents expressed *propositional* is the role of sentences in giving and asking for reasons. Inferential connections among claims are understood in turn pragmatically, in terms of *consequential* relations among the attitudes by means of which score is kept on commitments and entitlements to commitments—how attributing one commitment entails attributing others, precludes entitlement to others, and so on. At the next level, subsentential expressions can be understood as indirectly inferentially contentful, in virtue of the significance their occurrence has for the inferential involvements of the sentences in which they occur. The key concept at this level is *substitution*, for taking subsentential expressions to be contentful consists in distinguishing some inferences as substitution inferences, some inferential commitments as substitutional commitments. The substitutional structure of the inferences sentences are involved in is what the contentfulness of their subsentential components consists in. At the lowest level, unrepeatable tokenings (paradigmatically deictic uses of singular terms) can be understood as involved in substitution inferences, and so as indirectly inferentially contentful, in virtue of their links to other tokenings in a recurrence structure. The key concept at this level is *anaphora*. For taking an unrepeatable tokening to be contentful requires associating it with a repeatable structure of the sort that can be the subject of substitutional commitments. Anaphoric inheritance by one tokening of the substi-

tution-inferential potential of another does just that. The articulation characteristic of specifically *discursive* commitments is to be understood most broadly in terms of *inference,* the details of which require attention to *substitution,* the details of which in turn require attention to *anaphora.*

V. INTERPERSONAL ANAPHORA AND COMMUNICATION

1. Communication

Linguistic studies of anaphora typically distinguish between intrasentential anaphora and discourse anaphora.[85] The study of discourse anaphora is addressed not only to inter*sentential* anaphora, where an anaphorically dependent tokening occurs in a sentence different from that in which its antecedent occurs, but also to inter*personal* anaphora, where those tokenings are uttered by different interlocutors. The account of anaphora in terms of token recurrence—that is, as consisting in inheritance by one tokening from another of the structure that determines which substitutional commitments are relevant to its semantic assessment—applies to interpersonal (and hence intersentential) anaphora as well as to the fundamental kind of intrasentential anaphora. Indeed, certain features of that account stand out more sharply in the interpersonal case. Thus looking at examples of this kind a bit more closely provides an opportunity to clarify the sort of scorekeeping involved in attributing recurrence commitments, by specifying further the sort of inheritance they involve. More important, a new dimension of the expressive role of anaphoric connections among tokenings comes into play in the interpersonal case.

Anaphora as here conceived contributes two crucial sorts of expressive power to the idioms in which it is operative. First, as discussed in the previous section, anaphora makes possible the construction of repeatable expressions from unrepeatable tokenings. It is the mechanism by means of which unrepeatable tokenings are picked up and made conceptually—that is, ultimately, inferentially—significant. It is only because deictic and other indexical tokenings can recur anaphorically that their occurrence contributes to the inferential role played by sentences containing them, and hence that their occurrence can be counted as semantically significant at all; no occurrence without recurrence. Even where no overtly indexical expressions occur, this expressive capacity is crucial to the functioning of empirical languages.

Empirical languages are those that include noninferential reporting practices, and the authority of such reports is essentially tied to the particular unrepeatable tokening that is elicited as a response by the exercise of a reliable differential disposition. For such tokenings to have cognitive authority (for them to be available as expressing commitments to which interlocutors can be entitled and to which they can appeal in entitling themselves to further conclusions), their significance must be governed by substitutional

commitments. Such commitments relate recurrence repeatables, which in this case must be constructed anaphorically. Thus the report

> The traffic light has just turned red

can serve as a premise from which to draw conclusions according to inferential patterns that can be made explicit in the form of such conditional principles as:

> If it was red then, it will be green soon.

This expressive capacity is important even if it is the same interlocutor who makes the noninferential report and draws conclusions ("It will be green soon") from it; after the first glimpse the reporter may no longer be in a position to report the color noninferentially. But there is a second sort of expressive power anaphora contributes to discursive practice that arises *only* in the social context of inter*personal* communication of information.

For information (whether true or false) to be communicated is for the claims undertaken by one interlocutor to become available to others (who attribute them) as premises for inferences. Communication is the social production and consumption of reasons. So communication (giving and asking for reasons) involves the interaction of the *inferential* articulation of *contents* that is at the center of the *semantics* presented here and the *social* articulation of discursive *commitments* that is at the center of the *pragmatics* presented here. The nature and significance of this interaction of the inferential and the social dimensions of discursive practice is a large and important topic. It is the subject of the next chapter, which argues that the representational dimension of propositional contents is a reflection of the essential role played in their specifically inferential articulation by differences of social perspective—that is, differences between the point of view of the one who undertakes a commitment and the points of view of those who attribute it. As a result, the contents of the claims that are deployed monologically in intrapersonal reasoning in soliloquy must be understood as having been conferred by public practices of deploying claims dialogically in interpersonal reasoning in conversation. Meditation is made possible by disputation.

In advance of that fuller discussion of communication, the treatment here of the specifically social dimension of the expressive function of anaphora can only be preliminary. Nonetheless, some of the cardinal points are sufficiently detachable to be available already at this point. Interpersonal anaphora plays an important role in securing the possibility of communication across the doxastic gap created by the differing commitments of speaker and audience. The capacity of those in the audience to pick up a speaker's tokening anaphorically, and so connect it to their own substitution-inferential commitments, is part of what makes it possible for them to understand the speaker's utterance by extracting information from it. Anaphoric connec-

tions among tokenings that are utterances by different interlocutors provide a way of mapping their different repertoires of substitutional commitments onto one another—a structure scorekeepers can use to keep track of how each set of concomitant commitments relates to the others.

Such correlation of the substitution-inferential commitments (and hence doxastic commitments) undertaken by a scorekeeper with those attributed to others is a necessary part of the *interpretation* that is the uptake by a scorekeeping audience of some speaker's claim. It is an essential part of being able to use others' judgments as reasons, as premises in the scorekeeper's own inferences (even just hypothetically) to assess their significance in the context of those collateral commitments. Interpretation in this sense is necessary even in the case where all parties share a language.[86] The reason communication requires interpretation of this sort is twofold. First, speaker and audience typically have *different* sets of collateral commitments—if they did not, communication would be superfluous. Second, the inferential significance of a claim (what its consequences are and what would count as evidence for it) depends on what *auxiliary hypotheses* are available to serve as collateral premises. So differences in background beliefs mean that a remark may have one inferential significance for the speaker and another for each member of the speaker's audience.

2. Frege and Kant on Fruitfulness

This point reaches deep into inferentialist approaches to semantic content. It will be recalled that in the definition of conceptual content that opens the *Begriffsschrift*, Frege acknowledges the role of collateral commitments serving as auxiliary hypotheses: two judgments have the same content if and only if "all inferences that can be drawn from the first judgment when combined with certain other ones can always also be drawn from the second when combined with the same other judgments." The 'always' here signifies universal quantification over auxiliary hypotheses. It is not enough if there is *some* set of further judgments that yields the same set of consequences when combined with each of the candidates whose contents are being assessed. Such a requirement would obliterate distinctions of content, since for *any* two claims such a set of auxiliary premises can be found—in Frege's systems, for instance, any two claims have the same consequences when conjoined with a logical contradiction.[87] This quantification over possible sets of background beliefs accordingly is an acknowledgment that what follows from a claim depends on which further claims one is allowed to assume in extracting those consequences. (The dual point also holds, of course; for what constitutes evidence for a claim—its inferential circumstances, rather than consequences, of application—equally depends on the available auxiliary hypotheses.)

Frege makes more of this relativity of inferential significance to available

auxiliary hypotheses in the famous opening paragraph of "On Sense and Reference": "$a = a$ and $a = b$ are obviously statements of differing cognitive value [*Erkenntniswerte*]; $a = a$ holds *a priori* and, according to Kant, is to be labelled analytic, while statements of the form $a = b$ often contain valuable extensions of our knowledge and cannot always be established *a priori*. The discovery that the sun is not new every morning, but always the same, was one of the most fertile [*folgenreichsten*] astronomical discoveries." *Folgenreichsten* here is literally *richest in inferential consequences* (in what follows from it). The cognitive value of a statement is to be assessed by its inferential significance, by the difference that adding it to one's repertoire of endorsed judgments makes to what else one is committed or entitled to.[88]

This idiom, and in particular the understanding of analytic identities it expresses, is borrowed directly from Kant. In the section of his *Logik* entitled "Logical Perfection of Cognition as to Quantity,"[89] he says: "The magnitude of cognition may be understood in a twofold way, either as extensive or as intensive magnitude. The former refers to the extension of cognition and therefore consists in its volume and manifoldness; the latter refers to its content [*Gehalt*], which concerns the manifold validity [*Vielgültigkeit*] or logical importance and fruitfulness [*Fruchtbarkeit*] of a cognition, as far as it is considered as a ground for many and great consequences [*großen Folgen*] (non multa sed multum)." Here content is understood in terms of fruitfulness in the sense of leading inferentially to many consequences. Kant's definition of analyticity to which Frege is appealing is similarly couched in these terms:

> The identity of concepts in analytic judgments can be either explicit [*ausdrückliche*] (explicita) or non-explicit [*nicht-ausdrückliche*] (implicita). In the former case analytic propositions are tautological.
>
>> Note: Tautological propositions are virtualiter empty or void of consequences [*folgeleer*], for they are of no avail or use. Such is, for example, the tautological proposition Man is man. For if I know nothing else of man than that he is man, I know nothing else[90] of him at all.
>
> Implicitly (implicite) identical propositions, on the contrary, are not void of consequences or fruitless [*folge- oder fruchtleer*], for they clarify the predicate which lay undeveloped (implicite) in the concept of the subject through development (explicatio).
>
>> Note: Propositions void of consequences must be distinguished from propositions void of sense.[91]

Implicitly identical propositions have an expressive role—namely developing the content of a term by making explicit some of its inferential consequences, as in analytic claims such as "The oldest living mammal is a vertebrate."

The circumstances under which claims of the form $a = b$ "contain valuable extensions of our knowledge" are those in which, first, they are not

analytic (and hence a priori) in the explicative sense and, second, the knowledge they are extending includes further claims expressed using the terms *a* or *b*. For it is only in the presence of such auxiliary hypotheses that the identity licenses nontrivial substitution inferences. What it is for a claim to have a nontrivial cognitive value or content is accordingly defined by a *particular* quantification over possible sets of collateral commitments; there must be *some* context in which adding the claim has nontrivial inferential consequences. So the notion of content is being defined in terms of a more basic notion of the inferential significance of adding a claim to a set of antecedently endorsed claims. For a claim to have a nontrivial content at all is for the inferential significance of its endorsement to include nontrivial inferential consequences in *some* doxastic context, and for two claims to have the *same* cognitive or conceptual content is for their significances to comprise the same inferential consequences in *all* doxastic contexts. The primitive notion of inferential significance (of what follows from a claim and what is evidence for it) is explicitly relativized to a set of background claims—namely the set of those that are available as auxiliary hypotheses or collateral premises in extracting inferential consequences.

So even though Kant and Frege do not talk about the social dimension of inferential articulation, their elaborations of inferential conceptions of conceptual or cognitive content implicitly acknowledge that the inferential *content* of a claim manifests itself in different inferential *significances*—different claims counting as its consequences and potential evidence—from the perspectives provided by various sets of concomitant commitments. When Frege speaks of "extensions of our knowledge," he is comparing the perspective available before a claim is added to the repertoire of commitments we undertake (and take ourselves to be entitled to) with the perspective available afterward. What is a fruitful (inferentially significant) addition from one point of view may not be so from another. Though Frege does not discuss the consequences this observation has for understanding synchronic communication connecting different doxastic perspectives, looking at that case is helpful in understanding the diachronic cases he does appeal to.

3. Quine, Communication, and Reference

The underlying point is that what a given endorsement of claim commits one to, is entitled by, and is incompatible with depends on what else one is committed to, on what collateral information is available as auxiliary hypotheses for the inferences in question. Quine appeals to this Duhemian relativity of evidential significance to total evidential context in the closing sections of "Two Dogmas" to enforce a constraint on theoretical concepts of meaning (that is, claim content). Transposed into the idiom in use here, his holist argument for relativizing the meaning of a claim to the "total theory" of which it is a part is that:

—the meaning of a claim is what must be grasped to understand it, and

—what is understood must at least determine the inferential significance of endorsing what is understood, but

—what follows from a claim depends on what other claims are available as auxiliary hypotheses, so

—any difference in collateral commitments means a difference in inferential significance, hence meaning.

The fact that the inferential significance of endorsements is always and in principle relative to collateral commitments available as auxiliary hypotheses shows just what Quine wants it to, and thereby gives a definite sense to the claim that "the unit of meaning" is the whole theory or set of concomitant beliefs.

What effect does this relativity of inferential significance have for understanding *communication*? Quine does not explicitly raise this issue because he systematically waffles on the question of whether his "webs of belief" or "total theories" are individual or communal, whether we each have a different "total theory" or all share one. The account being unfolded here of the social-practical structure of inferential articulation—and hence of propositional contents—is one way of trying to take account of the motivations that push him now to talk one way, now another. This is an important issue in the context of an argument for not distinguishing changes of meaning from changes of belief; what is one to make of the consequence Harman extracts from this theory, that when I notice a cloud pass in front of the sun, the meaning of all my words changes? It must be granted that the noninferential addition of this new commitment alters (at least slightly) the inferential significance of all the claims that I do endorse, and all those I might. At the very least, conditionals of the form "If there is a cloud in front of the sun, then *p*" clearly would come in this way to have a different potential for transforming my commitments, and this would in turn alter the inferential significance of any claim that could appear as the consequent of such a conditional (or the conclusion of the inference it makes explicit)—and that is any claim whatsoever.

But must this alteration of the inferential *significance* different claims have for me be understood as involving an alteration in the inferential *content* they express? The view developed in the next chapter is one according to which the *inferential* holism that requires the pragmatic significance of doxastically endorsing a propositional content to be relativized to a repertoire of concomitant commitments must be understood in the context of a *social* holism. That social holism requires the grasp of the semantic content whose endorsement has such a significance to depend on scorekeepers' abilities to exploit relations among the different perspectives constituted by the different commitments undertaken by and attributed to those whose deontic scores

they keep track of. The significance for the understanding of communication of a holism that relativizes to a repertoire of background commitments either the inferential significance or the conceptual content of claims depends, like the significance of all commitments, on the auxiliary hypotheses that are available to serve as collateral premises in drawing inferential consequences from it. In this case, the background commitments it is important to be aware of take the form of a model of communication. Holism about inferential significances has different theoretical consequences depending on whether one thinks of communication in terms of *sharing* a relation to one and the same *thing* (grasping a common meaning) or in terms of *cooperating* in a joint *activity* (coordinating social perspectives by keeping deontic score according to common practices).

Communicating is naturally conceived of as *conveying* something. According to such a conception, before an episode of communication takes place only the communicating agent possesses what is to be conveyed; after successful communication the recipient possesses it as well. Overt performances serve as the *vehicles* by which what is communicated is transported from speaker to audience. In the Lockean version of this transportation model of communication, what is transferred is ideas—which are related to words by conventions, which are in turn reflected in the associations of the various interlocutors. Upon having an idea, the speaker associates a spoken or written word with it, and upon hearing or seeing the word, the audience associates the corresponding idea. Communication is a way for speaker and audience to achieve a shared idea.

The framework conception of communication as conveyance of something can be filled in by various particular notions of what is conveyed. Rather than ideas, it might be propositions, meanings, or information that speaker and audience are understood as sharing. Of course not every sort of performance that brings about a similarity between its maker and its taker provides a candidate for a conception of communication. The concept of communication involves that of *understanding*. What is to be communicated by an utterance is what its audience is to understand by it. What the producer of a meaningful performance has initially and what in the case of successful communication its consumers eventually acquire is something—a content or meaning determining the significance of the remark—that is understood by both parties. How the details of the conveyance model are filled in depends on a further conception of what it is to grasp or understand what is conveyed.

A problem arises if this commonsensical model of communication is combined with the inferentialist account of discursive practice in terms of deontic scorekeeping presented in these pages. According to that account, the fundamental communicative performance (making a claim) is acknowledging or undertaking a doxastic commitment. The sort of understanding or uptake of such a performance required for successful communication is for

the audience to figure that performance correctly in its score: to attribute the right commitment to the one making the claim. What makes the commitment a discursive commitment is its inferential articulation. In particular, the propositional contentfulness of a doxastic commitment consists in the material norms governing its role as premise and conclusion in inferences.[92] In the paradigmatic case of communicating by claiming, the audience's understanding of a claim must determine the inferential significance that adopting or believing that claim would have—that is, what one would be committing oneself to by endorsing it, what other commitments might entitle one to that endorsement, what other commitments are incompatible with it (and so preclude being entitled to such an endorsement), and so on. It is only insofar as the audience assigns some such significance to an utterance that a claim content is communicated or conveyed thereby.

Given the relativity of the inferential significance of a claim to the context of concomitant commitments available to serve as auxiliary premises, it follows that inferential significance is not preserved in communication—is not conveyed or transported from producer to consumer of communicational performances. For any difference in collateral commitment may involve a difference of inferential significance, understanding, and appropriated meaning. If I believe that Zoroaster is the sun and that its shining is his beatitude, then an utterance of "The sun is shining" means something different in my mouth than it does in your ears. If it is nonetheless possible for us to agree or disagree about that claim, that cannot be because it has the same significance for us. Inferential significance can be determined only relative to a total belief-set, so if what audiences understand must determine such significances, it cannot be independent of the context of collateral commitments. Since, as pointed out above, communication is superfluous in the case in which all commitments are shared (which alone would guarantee sameness of inferential significance), if inferential significances were what needed to be conveyed for communication to take place, communication would be impossible in all cases in which it was not otiose.

This is the line of thought that led theorists (such as Feyerabend[93]) who took Quine's inferential holism seriously to worry about the incommensurability of different theories or sets of commitments. Corresponding to the transportation model of communication is an accumulation model of progress—acquiring epistemic rights to more and more true claims fabricated out of a common stock of meanings or candidate belief contents. If these must at least include inferential significances, then since those significances can depend upon any collateral commitments, meanings are not shareable across theories, and so not establishable cumulatively as theory develops and changes. Given the very different background beliefs quantum theory has given us to govern our inferences involving the word 'electron', how can we so much as understand Rutherford's turn-of-the-century claim that electrons are particles with definite boundaries, orbiting atomic nuclei with definite

boundaries? Given that 'electron' meant something so different for him than it does for us, how is it possible for us to deny the very claim that he was making—as opposed to denying some *other* claim whose content appeals to *our* concept of electrons? 'Incommensurability' is the name given to this threat to our understanding of what communication is and how it is possible. It is a threat that arises for inferentialist approaches to meaning and understanding once the sensitivity of inferential significances to background beliefs is appreciated. Although this challenge is more pointed in the case of diachronic conceptual change in the history of scientific theories, the corresponding difficulty evidently confronts inferential role theories of the sort of content that is grasped and conveyed in synchronic, face-to-face, intralinguistic communication among interlocutors with different repertoires of doxastic commitments.

Quine himself quickly drew the conclusion that what matters semantically is not *meaning* but *reference*—what we are representing or talking about rather than just what we are saying about it. Although he does not put the point in the context of communication, this move reflects the realization that even if (in virtue of my Zoroastrian beliefs) the observation that the sun is shining means something different in my mouth than it does in your ears, you can still learn something from me that you can use in your own inferences—if and insofar as you understand me to be talking *about* the sun, and saying *of* it that it belongs in the class of shining things, that I am representing that thing as being in that class. Again, though Rutherford's many false background beliefs make his claim that electrons orbit around an atomic nucleus mean something to him that is unintelligible in the context of post-quantum-theory background beliefs, we can still understand him to have been talking *about* electrons (the same things we refer to) and to have been representing them as having certain properties and standing in certain relations. The information communicated consists in the purely extensional content of the claim made. A difference in inferential significance and commitment is compatible with identity of referential commitment and achievement. Where common reference of terms and extensions of predicates can be secured in spite of inferentially different employments, progress is comprehensible as talking about more and more objects, invoking more and more predicate-extensions, and coming to say more and more true things about those objects, for instance by classifying them under the predicate extensions. This is the conclusion that comes out of the debate between Feyerabend and Sheffler.[94] It is clearly the lesson that the early, realistic Putnam drew and conveyed to students and admirers such as Field, Boyd, and Devitt. It even sets the terms for Putnam's recantation of realism in "Realism and Reason." Thus an important motivation for the emphasis on semantic extensions—the referential dimension of discourse—can be found in the concern with making intelligible the possibility of communication.

This strategy gives up on contents as inferential roles, in favor of a differ-

ent sort of primitive. Inference can then be reinstated at two levels. First, some inferential proprieties can be read off of inclusion relations among the extensions of expressions. More important, recognition of the relativity of extensions to various elements of context yields the notion of intensions, as functions from indices to extensions. Such intensions are a more robust sort of content, which can be seen to be shared by speaker and audience in favored cases. Also, more finely grained inferential proprieties can be read off of inclusion relations among the sets that serve as domains and ranges of the intension functions. This is a two-leveled scheme, starting with extensions and ascending to intensions as functions defined on them. Inferential significances (the inferential potentials of particular claims in particular doxastic contexts) play no systematic role. Instead of inferential significances varying from speaker to speaker, there are extensions varying from possible world (together perhaps with other indices) to possible world.

4. Intensions

Analogy with this appeal to functions suggests that in the inferential case one might treat the inferential content expressed by a sentence tokening as a function, assigning to each repertoire of concomitant commitments an inferential significance. Such significances could be (crudely) thought of as ordered pairs of circumstances and consequences of application. The first element then might consist of sets of inferentially sufficient antecedent claims (those from which the claim in question can be inferred) and the second of a set of inferentially necessary consequent claims (those that can be inferred from the claim in question).[95] Since what is evidence for or commits one to a claim, and what it is evidence for or commits one to, depends on what background commitments are available as auxiliary hypotheses, inferential *contents* could then be thought of as *functions*. The content of each claim would be represented by a function that takes sets of concomitant background commitments as arguments and yields inferential significances as values.

The theoretical advantages of such a picture would accrue from taking inferential contents so construed as what is shared and communicated within a discursive community. Differences between the inferential significance that a claim has in the mouth of a speaker and the ears of an audience would then be compatible with a common understanding of what is being claimed. It would not then be necessary to concede the counterintuitive claim that the meanings of all one's words change (at least slightly) whenever one acquires a new belief. This explanatory advantage would be bought at a significant price, however. Unless the theorist is content with stipulative semantics—associating intension functions with expressions by fiat—an account must be offered of what it is about the way expressions are *used* that *confers* such contents on utterances and the states and attitudes

they express. What is required in this case is an answer to the question, What is it for an expression to be so used as to have associated with it one rather than another intension determining a function from doxastic context to inferential significance?

The generic difficulty with answers to this question stems from the very features that make an intensional response attractive in the first place. For functions of the sort in question are individuated so finely that it is hard to see how the use of an expression could determine that one rather than a slightly different one should be associated with it. In different forms this is the worry underlying Quine's rejection of intensions, Lewis's discussion of the relation of linguistic behavior to formal semantics for artificial languages in "Languages and Language,"[96] and Kripke's "finiteness" version of Wittgenstein's skeptical arguments concerning the underdetermination of use by meaning if meanings are conceived in standard ways. In each case the difficulty arises because one can in general construct a function that differs from a given one only for arguments that are in one way or another beyond the reach of behavioral dispositions. Where this is so it becomes difficult to see what is being envisaged (never mind how one could know that it is true) when it is said by the theorist that one rather than another of these behaviorally indistinguishable functions is nonetheless to be associated with a particular claiming.

This difficulty is particularly pressing in the case of communication across generations. Thinking of the communication of content in terms of shared intensions, functions from context of collateral commitments to significance, is most plausible as a response to worries about incommensurability for a synchronic linguistic community. Sharing intensions is speaking the same language in a strong sense. It is not clear how plausible such an account is in the diachronic case, where what is at issue is the possibility of incommensurability produced by conceptual change within a scientific tradition. Surely Rutherford or even Bohr did not and could not have shared the intensions contemporary physicists associate with such expressions as 'electron', 'mass', 'particle', and so on. It is not just that our views have changed substantially during the twentieth century, but that they have changed in ways unforeseeable by our conceptual ancestors of a few generations ago. It would require considerable argument to show that they had nonetheless used their expressions according to intensions that left room for all of our radical rethinkings, which could accordingly be represented just by differences in the context of assertional commitments with which each claim is conjoined. It is not possible to rule out such an approach a priori, but it is not surprising that it is hard to find a champion for an intensional transportation model of diachronic communication.

These cases provide one of the strongest motivations for adopting a different strategy: one that breaks with the conclusion Quine arrives at in "Two Dogmas" by distinguishing between a kernel of inferential (and perhaps also

doxastic) commitments that must be shared by those who count as grasping the content, concept, or meaning in question, and a shell of peripheral beliefs, which could differ without alterations of content. One privileges *some* of the inferences a concept is involved in as constitutive of it, treating the rest as warranted by collateral information. Grasping the concept then involves mastering only these essential inferences, and these are what interlocutors must share on pain of misunderstanding one another. Acquiring a new peripheral belief—for instance that a cloud now obscures the sun—would not then count as altering the concepts expressed by such words as 'sun'.

The difficulty faced by this approach is just the one Quine emphasized: saying what it is about the practices of using expressions that deserves to be characterized as treating some claims and inferences involving a concept as essential to it, and others as providing merely ancillary information about what it applies to. In constructing artificial languages, one might simply stipulate that some commitments are to be in the first class, while others are in the second. Even then one would be obliged to say how the proprieties of using expressions then differ depending on how particular commitments are classified. But insofar as this apparatus aspires to contribute to the analysis of natural languages or languages in use, those features of discursive practice that confer such a distinction of status between conceptual and merely empirical commitments must be specified. In the present context, one would need to explain in scorekeeping terms the different roles played by the practical attitudes of taking or treating commitments as conceptual and empirical. Of course, the fact that Quine can find no trace in our discursive practice of an analytic/synthetic distinction by looking at such candidate attitudes as treating as unrevisable or as a priori is hardly decisive. Other possibilities are not far to seek. (One that has not gotten the attention it deserves is Sellars's suggestion that the practical status that privileges concept-constitutive inferences is their counterfactual robustness.)[97]

5. A Three-Leveled Approach

Nothing rules out such a strategy, but it is not the one pursued here. The present account substitutes a three-leveled approach for the standard two-leveled one.[98] Instead of beginning with extensions and defining intensions as functions from indices (including possible worlds, which provide a background of endorsed claims serving as the facts) to extensions, the story begins with the inferential significances of claims. The theory then moves down, defining the extensional dimension of discourse in terms of substitution-inferential commitments. Those commitments in turn determine equivalence classes of expressions corresponding to what is represented—what is talked and thought *about*. Various features of the interpretive scorekeeping practices appealed to in this move *down* from perspectival inferential significances to *extensions* then make it possible to

move *up* from those significances to propositional or conceptual *contents* (corresponding in some ways to *intensions*), which systematically relate the distinct perspectives responsible for the different significances claims have to different interlocutors. What has been left out of the traditional formalism of extension and intension, from this point of view, is precisely the *interpersonal communicational* dimension. Yet it is this that gives the discerning of extensions and intensions its connection to discursive practice, and hence the explanatory role in virtue of which alone it can be appropriate to call what is discerned *semantic* correlates.

The way in which concern with what is talked about arises in the process of mapping the repertoire of commitments of an interpreted interlocutor onto the repertoire of commitments of an interpreting interlocutor is discussed in the next chapter. That chapter also seeks to explain the sort of perspectival propositional contents that coordinated scorekeeping practices confer. The paradigm of communication as joint possession of some common thing is relinquished in favor of—or modified in the direction of—a paradigm of communication as a kind of cooperation in practice. What is shared by speaker and audience is not a *content*-as-function but a scorekeeping *practice*. Contents as functions from repertoires to inferential significances can be seen as implicit in such practices, but the practice can retain its identity even though the functions implicit in it are different (at different times, and from different doxastic points of view).

For what is implicit can be made explicit in various, not always compatible ways. From each doxastic point of view on a speech act there can be a content common to the one undertaking a commitment and the scorekeepers attributing it, but what is taken to be shared may be different from the points of view of different scorekeepers. Thus inferential *contents* are essentially perspectival—they can in principle be specified only from a point of view. What is shared is a capacity to navigate and traverse differences in points of view, to specify contents from different points of view.[99] Explaining this capacity is explaining what it is to take or treat (understand or interpret) someone's remark as representing or being about one thing rather than another. So what appear theoretically as distinct moves *down* from inferential significances to *extensions* by assimilating expressions as intersubstitutable (= coreferential), on the one hand, and *up* from those significances to *intensions* by relativizing them to repertoires of background commitments, on the other, correspond to aspects of a single interpretive activity of understanding, grasping a meaning—the cognitive uptake of communication that is deontic scorekeeping.

The perspectival nature of propositional contents and the way in which their essential representational dimension emerges from communicative scorekeeping practice is approached in the next chapter by considering what is made explicit in *de dicto* and *de re* ascriptions of propositional attitude. The role of anaphora in securing coreference across differences in perspective

can be considered here as an introduction, however. Anaphora serves to link the equivalence classes of expressions that are intersubstitutable according to one interlocutor to the classes generated by the substitutional commitments of others. The need for such a mechanism arises in the interpersonal context because the speaker may have different substitutional commitments from the audience. If the speaker believes that the first postmaster general of the United States is the inventor of bifocals, and the audience does not, the inferential significance of the claim "The inventor of bifocals spoke French well" is different from their various perspectives. The question then arises how those in the audience can manage to have an attitude toward the same claim the speaker is making, can agree or disagree with *it*, rather than some variant of it that they associate with the same noises or inscriptions. Given that speaker and audience disagree about whether the claim is about the first postmaster general of the United States, how can they nonetheless secure a common topic of conversation in order to argue about whether or not *he* spoke French well?

This way of putting the question contains the answer. (Compare the way the fact that one can assert the modal nonrigidity of the description 'the first postmaster general' by saying "Benjamin Franklin was the first postmaster general, but *he* [the man just referred to] might not have been" points to the central expressive function played by anaphoric relations in understanding that phenomenon.) Use of an anaphoric proform implicitly *stipulates* coreference with the anaphoric antecedent upon which it is semantically dependent. Thus differences in the substitutional commitments that determine the propriety of inferences involving 'the inventor of bifocals' according to speaker and audience can be bracketed and a common topic of conversation secured by using a tokening that is anaphorically dependent on the speaker's tokening. To respond to the speaker by saying "*He* did not speak French well" is to disagree with the claim made, *whoever* the inventor of bifocals might turn out to be. Indeed, if more than the object referred to is in question, the claim can be affirmed or denied by using an anaphoric dependent on the whole sentence, rather than just picking up one of its singular terms: the audience can say "That is true" or "What you claim is false." Interpersonal anaphora achieves just the effect that matters for securing communication in the face of differences in collateral commitments.

The capacity to use a pronoun that anaphorically picks up another's tokening is also a cardinal component of another important ability, one whose cognitive significance is often underrated. For pronouns enable us to talk without knowing what we are talking about. Thus a speaker can come late into a conversation in which someone is already being referred to as 'he' and can jump in—continuing that conversation with a remark such as "If he did that, he deserves whatever he gets." The speaker may under such circumstances have no idea at all of who it is that is being talking about. The form in which the later claims are expressed nonetheless commits the speaker

anaphorically to their being about whoever it is the others were already talking about. That is, a scorekeeper will assess the doxastic commitment the latecomer has undertaken according to whatever substitutional commitments that scorekeeper takes to govern the antecedent of the anaphoric tokening 'he'. Anaphora is a mechanism that permits undertaking and attributing commitments concerning objects that one need not be able to specify (nonanaphorically) if challenged. Thus one is not obliged to know or accept the descriptions by means of which the utterer of the anaphoric initiator might pick out the subject with respect to which both are undertaking and attributing commitments.

6. Speaker's Reference

These cases of interpersonal anaphora show that one must be careful in thinking of anaphora as inheritance of substitutional commitments by one tokening from another. The anaphoric antecedent is what determines the substitutional commitments relevant to the assessment of the significance of its dependents. But in using a pronoun that is anaphorically dependent on a tokening uttered by another, one is not thereby bound by whatever substitutional commitments the other happens to acknowledge as governing that tokening. An interlocutor who disagrees with the speaker's assertion "The inventor of bifocals spoke French well" by saying "*He* did not speak French well" is not making an incompatible assertion by adding "And he was not the first postmaster general, either," even though the utterer of the antecedent of those pronouns *is* committed to the intersubstitutability (that is, coreference) of 'the inventor of bifocals' and 'the first postmaster general.' Although the divergence of perspective that makes the point evident did not arise in the case of intrapersonal anaphora,[100] the substitutional commitments to be inherited anaphorically by one token from another are assessed by the scorekeeper who attributes the anaphoric commitment, that is, who takes or treats the one tokening *as* anaphorically dependent on another. To take one tokening to be anaphorically dependent on another is to take it that it should be understood as governed by whatever substitutional commitments govern its antecedent.

Different scorekeepers may disagree about what these are, and they may disagree even with the ones producing the performances whose significance they are assessing. They may nonetheless all agree in attributing an anaphoric commitment, that is, in interpreting one tokening as being anaphorically dependent on (hence a recurrence of) the same antecedent tokening. A scorekeeper who takes it that the inventor of bifocals is the inventor of the lightning rod will take it that the first speaker claimed *of* the inventor of the lightning rod that he spoke French well, and that the second speaker claimed *of* that same individual that he did not speak French well. That is, a scorekeeper who undertakes such a substitutional commitment and attributes

that anaphoric commitment is obliged to take it that what the first speaker said is *true* just in case the inventor of the lightning rod spoke French well, and that what the second speaker said is true just in case he did not. For what a scorekeeper takes to be true is just what that scorekeeper endorses. The scorekeeping significance of attributing an anaphoric commitment is accordingly just that the significance of the dependent tokening is to be assessed according to the *same* substitutional commitments by which its antecedent tokening is assessed—whatever those are. Where the scorekeeper is concerned with when a given claim is true, it is the substitutional commitments that scorekeeper *undertakes* that matter, rather than those attributed to the utterer of the antecedent.

This is to say that according to a scorekeeper who undertakes a commitment to the intersubstitutability of 'the first postmaster general' and 'the inventor of the lightning rod,' one who asserts "The first postmaster general spoke French well" has thereby in a certain sense undertaken a commitment to the claim that the inventor of the lightning rod spoke French well. And this is true even in the case where the one making the original assertion would deny that the first postmaster general is the inventor of the lightning rod. The speaker is, according to such a scorekeeper, committed to that further claim just in the sense that what he has said is true if and only if the inventor of the lightning rod spoke French well. In this sense, what someone is committed to may (according to a scorekeeper) not only outrun, but even conflict with, what that interlocutor is prepared to acknowledge. The scorekeeper must keep two sets of books.

The necessity for this dual score follows from the fact that there are in principle two places a scorekeeper can draw auxiliary hypotheses from in extracting the inferential consequences of (and so the commitments consequentially undertaken by) a set of commitments some individual is taken to acknowledge. Those auxiliary hypotheses may be other commitments the individual acknowledges, or they may be commitments the scorekeeper *undertakes* (acknowledges), rather than attributing as acknowledged. Since these latter represent the *facts* (facts being just true claims), according to the scorekeeper, these latter consequences are those that, according to the scorekeeper, *actually* follow from the claims made (given how things really are), regardless of whether the one making the claims realizes that they follow or not. The relations between these two sets of books, and the way their interaction constitutes the representational dimension of propositional content, is the topic of the next chapter. As an introduction to the perspectival character of claim contents that is investigated there, it is helpful to consider the phenomenon of *speaker's reference* in terms of interpersonal anaphora.

What Kripke called "speaker's reference," by contrast to "semantic reference"—a distinction closely allied to Donellan's distinction between "referential" and "attributive" uses of definite descriptions[101]—is a phenomenon that depends on the possibility of taking up identificatory or substitutional attitudes toward a tokening that is not treated as functioning in a type-recur-

rent way. It is a matter of the significance (substitution-inferential potential) that an audience attributes or ought to attribute to a particular tokening, by contrast to the significance that would otherwise be associated with it on the basis of its type. What is fundamental is the way an audience interprets or keeps score on the tokening. Once what it is to *take* someone to be speaker referring is understood, it will be possible to understand what it could be for some situation conventionally to call for or warrant the use of this sort of interpretation. So the account is in terms of audience uptake, not what the speaker does or intends. Like the cases just considered, the situation in which an audience counts as treating someone as having "speaker-referred to someone other than the one semantically referred to" by a remark is always one in which the identificatory commitments in the vicinity of the recurrence class of the uttering to be interpreted that the audience attributes to the speaker are different from those that the audience undertakes itself.

Adopting the subsentential forms a bit so as better to accord with tradition, a case might go like this. The speaker, Fred, says, "The man in the corner with champagne in his glass is very angry." According to Wilma, in the audience for this remark, Fred claims that Barney is the man in the corner with champagne in his glass. So according to Wilma (that is, the commitments she *attributes*), Fred might just as well have expressed his claim by saying, "Barney is angry." But according to Wilma (that is, the commitments she *acknowledges*), Barney is the man in the corner with ginger ale in his glass, and the man in the corner with champagne in his glass is Nelson. (She takes it that Fred does not see Nelson and does not believe that he is in the corner at all.) Then we can say that, according to Wilma, Fred has speaker-referred to Barney and attributed anger to him but has semantically referred to Nelson and attributed anger to him.

These two different ways of interpreting the claim that Fred has expressed by his utterance correspond to assessing his assertional commitments with respect to the identificatory commitments that Wilma attributes to him, and to assessing those commitments with respect to the identificatory commitments that Wilma herself undertakes. It is essential that, according to Wilma, there is some expression that Fred could (compatibly with the commitment Wilma attributes to him) have used to semantically pick out his referent, in order that he be able to speaker-refer to it by another expression. For what Wilma is doing when she assesses his remark as true by taking him to have speaker-referred to Barney is treating his tokening /the man in the corner with champagne in his glass/ as an anaphoric dependent whose antecedent is another tokening that Fred *could* have used (and would have used had he realized that there was a dispute about the matter): perhaps a tokening of the type ⟨Barney⟩, or ⟨the man in the corner with bubbly liquid in his glass at whom I am looking⟩, or just ⟨that man⟩. Attributing speaker-reference rather than semantic reference is assessing the substitutional commitments a tokening owes its allegiance to anaphorically rather than by type.

Having this interpretive or scorekeeping strategy available is useful for

reasons of charity. Wilma can make more of what Fred says come out *true* (according to her) by taking some of his remarks this way. Charity of this sort is necessary only where there is a relevant difference in perspective between audience and speaker—that is, where it makes a difference whether the commitments taken to be available as auxiliary hypotheses in drawing infer-ential consequences from a claim are those *undertaken* by the scorekeeper or those *attributed* to the one whose performances are being assessed. When such a difference in social perspective becomes explicit in ascriptions of doxastic commitment (the fundamental propositional attitude), it appears as the difference between ascriptions *de re* and ascriptions *de dicto*.

Appendix: Other Kinds of Anaphora— Paychecks, Donkeys, and Quantificational Antecedents

Section III offers an account of the practical attitude a discursive scorekeeper must adopt in order to count as treating one tokening as anaphorically dependent on another in the most basic sense—that is, taking the substitu-tional commitments that determine the significance of the dependent to be inherited from those that determine the significance of its antecedent. This suffices to show how anaphoric relations can be introduced into or diagnosed in the simplified discursive practices described here. In actual natural lan-guages, anaphora is an immensely complex phenomenon; many more sophis-ticated tropes have been built up around the asymmetric token-recurrence structures identified here as the fundamental anaphoric phenomenon. Dis-cussion of these goes beyond the scope of the present project, but perhaps a few signposts are in order.

   One important issue that is put to one side here concerns the thorny problem of paycheck cases. Understanding a sentence like

> (a) The man who gives his paycheck to his wife is wiser than
>   the man who gives it to his mistress

requires treating the anaphor /it/$_a$ as replaceable by another tokening of the same type as its antecedent, /his paycheck/$_a$, even though these tokenings will be governed by different SMSICs (have different referents). /His pay-check/$_a$ will not have the same referent as /it/$_a$, any more than the two tokenings of ⟨the man⟩ do. So this sort of anaphora cannot be dealt with in terms of the establishment of token-recurrence structures. Nonetheless, it is clear enough how to understand this sort of lexically 'lazy' anaphora: /it/$_a$ is replaceable by another token of type ⟨his paycheck⟩, and the antecedents of the two tokens of type ⟨his⟩ are the different, noncoreferential tokens of type ⟨the man⟩. The hard question (which is important to linguists and for some projects in artificial intelligence) is not understanding the correct reading but

telling when that sort of reading is called for. It is how to tell when one ought to understand anaphoric dependence in terms of token-recurrence, and when it should be understood rather in terms of the sort of type-recurrence that paycheck cases demand. The explanatory task undertaken here is finished, however, when the differences between the two sorts of readings have been made clear.

Another large issue passed by here concerns anaphoric dependents whose antecedents are *quantificational* expressions. In the simplest cases, the interpretation of claims formed in this way follows from the general account of quantifiers; the anaphoric chains determine what count as substitution instances of particular and universal quantifiers, the significance of those substitution instances is determined by the token-recurrence model, and the significance of the quantificational claim is determined disjunctively or conjunctively by those instances. This is how anaphoric dependents on quantificational initiators should be understood when those dependents behave like the bound variables of the predicate calculus. As Evans has pointed out, however, not all anaphoric dependents on quantificational antecedents are happily assimilated to this model.[102] Thus on the most natural reading,

> John bought *some donkeys*, and Harry vaccinated *them*

entails that Harry vaccinated *all* the donkeys John bought, whereas the bound reading

> [some x: donkeys x] (John bought x & Harry vaccinated x)

requires only that there be some donkeys that John both bought and Harry vaccinated. In his excellent discussion, Neale points out further that

> John bought *exactly two donkeys*, and Harry vaccinated *them*.
> *Few politicians* came to the party, but *they* had a good time.
> *Just one man* drank rum, and *he* was ill.

entail that John bought exactly two donkeys, few politicians came to the party, and just one man drank rum, respectively—consequences that are lost on the bound reading.[103] As he concludes: "The upshot . . . is that among pronouns anaphoric on quantifiers we need to distinguish between those that function as bound variables and those that do not."[104]

The interpretation wanted has already been alluded to, in the discussion above of 'definitization transforms' in connection with Chastain's treatment of anaphoric chains (see 5.4.2).[105] The quantificationally unbound anaphoric dependents of quantificational expressions in examples such as those above go proxy for definite descriptions formed from their antecedents. So the sentences above make the same claims as:

> John bought *some donkeys*, and Harry vaccinated *those donkeys*.

John bought *exactly two donkeys,* and Harry vaccinated *those donkeys.*
Few politicians came to the party, but *those politicians* had a good time.
Just one man drank rum, and *the man who drank rum* was ill.

Here the quantificational antecedent determines the class of relevant substitution instances, and the significance of the clauses in which the anaphoric dependent appears is determined by that class. The bound cases differ just in that the clauses in which the anaphoric dependent appears also function to constrain the class of substitution instances with respect to which *both* clauses are evaluated. The difference between the bound and unbound cases accordingly corresponds to a difference in the order of application of the two operations of determining a class of substitution instances and making anaphoric connections. In the bound case, the anaphoric connections govern the inheritance of substitution-inferential significance by one clause from another within each quantificational substitution instance; in the unbound case, they govern rather the inheritance of a class of quantificational substitution instances by one clause from another.

Thus what might be called 'definitizing' anaphora—which governs the inferential significance of quantificationally unbound anaphoric dependents having quantificational antecedents—is another sort of sophisticated anaphora. It is distinct both from lazy, type-recurrent ('paycheck') anaphora and from the basic case of token-recurrent anaphora discussed in the body of the text. It is clear, at least in broad outlines, how such anaphora should be understood in the discursive scorekeeping idiom developed here. As before, the difficult task is formulating rules codifying when it is appropriate to adopt one sort of reading rather than another. As before, no stand is taken here on this difficult problem—and so none on Evans's suggestion that a pronoun anaphorically dependent on a quantificational expression behaves like a variable bound by it just in case the pronoun is c-commanded by the quantifier.[106] In the idiom suggested in the text, these issues are all taken to concern when it is appropriate to do the trick (construe the substitution-inferential significance of anaphoric dependence according to one model rather than another) rather than what it is to do the trick (keep score according to one reading rather than another), which is all that is of concern here.

Geach's original donkey sentence was

Any man who owns a donkey beats it.[107]

Here one does not want the definitized reading, for those who own two donkeys are being accused of beating both of them. The trouble is that apparently then 'a donkey' must be understood as expressing a particular quantifier relativized to the universal quantifier expressed by 'any man'. But

neither of the two ways of putting this relation in standard first-order quantificational language seems right. Unlike the original,

[every x: man x] [some y: donkey y] (Owns (x, y) → Beats (x, y))

is compatible with

Some man who owns a donkey does not beat it,

while

[every x: man x] ([some y: donkey y] (Owns (x, y)) → Beats (x, y))

is syntactically incoherent, having the second quantifier, which occurs in the antecedent of a conditional, binding variables that occur in the consequent of that conditional.

This causes a problem, however, only for those concerned to provide a uniform way of mapping quantificational expression-types in natural languages onto operators in the first-order predicate calculus. Those not concerned with rules determining when it is appropriate to interpret tokens of the type ⟨a K⟩ or ⟨some K⟩ one way rather than another can rest content with understanding Geach's donkey sentence as having the inferential role of

[every x: man x] [every y: donkey y] (Owns (x, y) → Beats (x, y)).

Further anaphoric dependents on these quantificational expressions may then act either as quantificationally bound anaphors or as definitized ones.

A final sort of example that deserves mention is Bach-Peters sentences, such as

A boy who was fooling her kissed a girl who loved him.
The pilot who shot at it hit the MiG that chased him.

The difficulty here is that the anaphoric chains cross; each dependent inherits its substitution-inferential role from an antecedent that inherits its role in turn from the original dependent. So these anaphoric circles do not settle what is to be counted as the anaphoric initiator. As has long been recognized, however, these surface forms are ambiguous; they have two nonequivalent readings, depending on which expression is treated as an initiator (which includes a dependent), and which as a dependent.[108] With definitization, quantificational cases such as the first example reduce to those involving definite descriptions, like the second, and the two readings of those are not far to seek. As Neale puts it:

If 'the pilot who shot at it' is given wider scope, 'him' is bound and 'it' is D-type:

[the *x:* pilot *x* & [the *y:* MiG *y* & *y* chased *x*] (*x* shot at *y*)] ([the *y:* MiG *y* & *y* chased *x*] (*x* hit *y*)).

If 'the MiG that chased him' is given wider scope, 'it' is bound and 'him' is D-type:

[the *y:* MiG *y* & [the *x:* pilot *x* & *x* shot at *y*] (*y* chased *x*)] ([the *x:* pilot *x* & *x* shot at *y*] (*x* hit *y*)).[109]

What Neale calls 'D-type' anaphoric dependents are those to be interpreted by definitization transforms of their antecedents. Once again, there is no special problem with interpreting each of these readings in discursive score-keeping terms, so long as care is taken to distinguish anaphora determining the inheritance of substitution-inferential significance within quantificational substitution instances from anaphora determining the inheritance of classes of quantificational substitution instances (in the case of definite descriptions, singletons). The defining symmetry of the Bach-Peters sentences ensures that in this case there is no residual problem of determining when one reading rather than the other is appropriate; the *only* task is making sense of the two readings.

In conclusion, although there are other sorts of anaphora in play in natural languages besides the one taken as fundamental in the discursive scorekeeping semantics, there are strategies available for making sense of them within the model as developed in the text.

8

Ascribing Propositional Attitudes: The Social Route from Reasoning to Representing

> Quot homines tot sententiae: suo quoque mos.
> (So many men, so many opinions; his own a law to each.)
>
> TERENCE, Phormio

I. REPRESENTATION AND *DE RE* ASCRIPTION OF PROPOSITIONALLY CONTENTFUL COMMITMENTS

1. Introduction

At this point the deontic scorekeeping model has been developed sufficiently to show how discursive practice can confer *propositional* conceptual content on the repeatable *sentential* expressions used in performances manifesting doxastic and practical normative statuses, in virtue of their proper use being governed by *inferential* commitments. It has also been shown how discursive practice can further confer conceptual content on the repeatable *subsentential* components of those expressions, in virtue of their proper use being governed by *substitutional* commitments. Finally, it has been shown how discursive practice can confer conceptual contents on unrepeatable *tokenings*, whether sentential or subsentential, in virtue of their proper use being governed by *anaphoric* commitments. In this chapter these raw materials are assembled to address several important bits of unfinished business.

The chief task is to explain the *representational* dimension of thought and talk. Such an explanation is required to redeem a promissory note that was issued in Chapter 2, when commitment to an inferentialist order of semantic

explanation was undertaken. The complementary representationalist order of explanation, dominant since the seventeenth century, presents propositional contentfulness in representational terms from the outset—appealing to the notion of states of affairs involving represented objects as the worldly conditions of the truth of judgments. This approach is objectionable if it is pretended that an account in these terms gives one an independent grip on what is expressed by the declarative use of sentences—as though one could understand the notions of states of affairs or truth conditions in *advance* of understanding claiming or judging. The representational semantic tradition embodies an undeniable insight: whatever is propositionally contentful does necessarily have such a representational aspect; nothing that did not would be recognizable as expressing a proposition. The point of the inferentialist order of explanation is not to *object* to using representational locutions to talk about semantic content. Inferentialism must be understood instead as a strategy for *understanding* what is said by the use of such locutions. The objection is only to treating representational locutions as basic in the order of semantic explanation.

The inferentialist idea is to start with a preliminary understanding of conceptual content in terms of inferential articulation—to approach semantic contentfulness by means of the functional role claims play as premises and conclusions of inferences. Working out this idea along pragmatist lines focuses attention on inferring as a *doing*. An account is then sought of the practice of giving and asking for reasons, an account that can be expressed without the use of explicitly representational locutions. The aim is to be able to explain in deontic scorekeeping terms what is expressed by the use of representational vocabulary—what we are doing and saying when we talk about what we are talking *about*. Thus a criterion of adequacy for the inferentialist program is that it be possible to say without using specifically representational vocabulary what would count as introducing into discursive practice locutions that make *explicit* the *implicit* representational dimension of the semantic contents that claims acquire in virtue of their role in the game of giving and asking for reasons.

The claim developed and defended here is that representational locutions should be understood as making explicit certain features of *communicating* by claiming—the *interpersonal* giving and asking for reasons. The context within which concern with what is thought and talked *about* arises is the assessment of how the judgments of one individual can serve as reasons for another. The thesis is that the *representational* dimension of propositional content is conferred on thought and talk by the *social* dimension of the practice of giving and asking for reasons. Logicians typically think of inference as involving only relations among different propositional contents, not as also potentially involving relations among different interlocutors. However, discursive practice, the giving and asking for reasons, from which inferential relations are abstracted, involves both inter*content* and inter*personal*

dimensions. The claim is that the representational aspect of the propositional contents that play the inferential roles of premise and conclusion is to be understood in terms of the social dimension of communicating reasons and assessing the significance of reasons offered by others. The conceptual contents employed in *monological* reasoning, in which all the premises and conclusions are potential commitments of one individual, are parasitic on and intelligible *only* in terms of the conceptual contents conferred by *dialogical* reasoning, in which the issue of what follows from what essentially involves assessments from the different social perspectives of interlocutors with different background commitments. Representationally contentful claims arise in the social context of communication and only then are available to be employed in solitary cogitation.

Because, first, there is no propositional (and hence no conceptual) content without this representational dimension, and, second, that representational dimension is the expression of the social articulation of inferential practice, it follows that propositional and hence conceptual content can be conferred *only* by *social* practice. The social analysis of representational content will accordingly vindicate the claim of Chapter 3, as to the essentially social nature of conceptual content. It provides the promised argument for the claim that *discursive* practice must be understood as a fundamentally *social* practice. Appreciating the intimate connection between the representational dimension of conceptual content and the interpersonal or communicative dimension of the inferential practice of giving and asking for reasons brings into relief features of the inferential approach to semantics and the deontic scorekeeping approach to pragmatics that depend on their interaction.

On the semantic side, what appears is the social-perspectival character of propositional contents, and hence of conceptual contents generally. The semantic contents of discursive commitments, attitudes toward those commitments, and the linguistic performances that express those attitudes can in principle only be specified from the perspective provided by some repertoire of background commitments and attitudes; how it is correct to specify any particular content varies from one such repertoire to another. This relativity to doxastic point of view makes conceptual contents fundamentally unlike ordinary nondiscursive things.

On the pragmatic side, what appears is the social-perspectival character of the distinction between normative status and normative attitude—between what someone is *really* committed or entitled to and what anyone, including even the subject of those statuses, *takes* that individual to be committed or entitled to. Putting these semantic and pragmatic phenomena together yields an account of the *objectivity* of conceptual norms—a way of understanding how our scorekeeping practices can confer conceptual contents about whose proper applicability and consequences we can not only *each* but *all* be in error. It takes the form of an account of the structure that must be exhibited by the practices a community is interpreted as engaging in for those practices

thereby to be understood (both by those who are interpreted and by those who are interpreting them) as socially instituting implicit norms according to which the truth of claims and the correct use of concepts answer to how things objectively—rather than how things subjectively, or even intersubjectively—are taken to be.

These issues concerning the relations between pragmatics and semantics can be brought into simultaneous focus by considering the use of explicitating locutions that express the adoption of the sort of deontic attitudes studied by the pragmatic theory of scorekeeping in the form of the sort of propositional contents studied by the semantic theory of inferential articulation. The attribution of discursive commitments is an attitude that is implicit in deontic scorekeeping practices. It is something that scorekeepers *do*. The introduction of a sentential operator that functions as "*S* believes that . . ." or "*S* is committed to the claim that . . ." does in English makes it possible, not merely implicitly or in practice to *take* someone to be committed to a claim, but explicitly to *say* that someone is committed to a claim, and to which claim. The explicit is the claimable, what can be given as a reason and have reasons demanded for it; ascriptional locutions make implicit attributions explicit as the contents of claims.

The notion of explicitating locutions was introduced in Chapter 2 with the example of conditionals, which make *inferential* commitments explicit in the form of *assertional* commitments. Negation was then explained as playing the expressive role of explicitating incompatibilities. The paradigm of subsentential explicitating vocabulary was provided by identity locutions, which make *substitutional* commitments explicit. The expressive role of quantifiers and the operators that form definite descriptions then appeared as making explicit various kinds and combinations of substitutional commitments. All these are *semantically* explicitating locutions, which make explicit the *inferential* involvements of the nonlogical expressions they operate on. Traditional semantic vocabulary, paradigmatically 'true' and 'refers', has been shown to play the expressive role of making explicit the *anaphoric* relations that make it possible for unrepeatable tokenings to play indirectly inferential roles and so to have conceptual content. These too figure in the present scheme as semantically explicitating locutions (though the question of whether anaphora should be classed as a semantic or pragmatic phenomenon is equivocal at best according to traditional ways of dividing up those domains).

Vocabulary for ascribing propositional attitudes plays a corresponding explicitating expressive role, but it is *pragmatic* rather than semantic features of discursive practice that are made explicit by its use. For it is the adoption of a deontic *attitude* (the attribution of a deontic status) rather than the inferential articulation of a semantic content that is in the first instance made explicit by using these locutions. For this reason the traditional denomination of them as locutions for ascribing propositional *attitudes* is

somewhat misleading in the present context; it is propositionally contentful commitments—a kind of deontic *status*, not a kind of deontic *attitude*—that is in the first instance ascribed (that is, explicitly attributed) by the use of these tropes. An attitude is *expressed* by asserting an ascription, but it is not in general an attitude that is *ascribed*. Indeed, as will emerge, *two* attitudes are thereby expressed—for in asserting an ascription one *attributes* (usually to someone else) one propositionally contentful attitude, namely the one ascribed, and undertakes (oneself) another, namely the ascribing one.

However, ascriptional claims can not only be undertaken but also attributed, and even ascribed. Thus once the expressive resources they provide are available, it becomes possible to do what one could not do without them—to attribute not just statuses but attitudes. Only by considering what ascriptional locutions express is it possible to understand the relation between the *status* of *being* committed and the *attitude* of *acknowledging* a commitment—a distinction that has been systematically kept out of sight up to this point. Thus like the vocabulary that makes semantic (inferential) features of discourse propositionally explicit, the vocabulary that makes pragmatic (deontic) features of discourse propositionally explicit not only makes it possible to *say* things one could not *say* before but makes it possible to *do* things one could not *do* before, *by* saying those things. In virtue of this explicitating expressive role, propositional-attitude-ascribing locutions deserve to count as *logical* vocabulary, albeit in an extended or generic sense that includes pragmatic as well as semantic species.

2. De Re *Ascriptions Are the Fundamental Representational Locution of Natural Languages*

In natural languages there are two sorts of explicitly representational vocabulary—locutions whose expressive role is to make explicit the representational dimension implicit in the inferential articulation in virtue of which discursive social practice is conceptually contentful. Central to the first are expressions like 'true' and 'refers'. The way these work was considered in Chapter 5, where they were analyzed as anaphoric proform-forming operators. The sort of anaphoric dependence they display was then explained in Chapter 7 in terms of asymmetric token-recurrence and the consequent inheritance of substitution-inferential commitments. These expressions of natural language provide the basis from which technical semantic vocabulary is elaborated. Philosophers' uses of terms such as 'denotes' and 'satisfies' in formal semantic investigations, and 'represents' in less formal epistemological theorizing, originate in and—except where a misunderstanding of their grammar has led to commitment to hypostatized relations and properties—can be explained in terms of these anaphoric idioms.

There is another, and in many respects more important, sort of representational locution in ordinary language. It comprises the idioms we typi-

cally use to express the intentional directedness of thought and talk—the fact that we think and talk *about* things and states of affairs. Words such as 'of' and 'about' play their characteristic intentional, semantic, or representational expressive role in virtue of the way they figure in *de re* ascriptions of propositional attitudes. It is these ascriptions that we use to *say* what we are thinking and talking of or about. 'Of' is used in many ways that have nothing to do with intentionality or representational aboutness, for instance its use to indicate possession, as in 'the pen of my aunt' or 'the positive square root of two'. 'About' is used in many ways that have nothing to do with the ofness of thought and talk, as in "The book is about six hundred pages long" or "It is about time for the meeting to start." The core locutions in terms of which the intentionality-expressing senses of these words must be distinguished are *de re* ascriptions, such as "Thomas Jefferson believed of meteorites that they did not exist"[1] or transforms of them, such as "Henry Adams's claim that the first postmaster general did not invent the lightning rod is about Benjamin Franklin."

Thus in order to identify vocabulary in alien languages that means what 'of' and 'about' do when used to ascribe intentionality and describe its content—or to introduce such vocabulary into expressively impoverished languages—one must be able to recognize expressions of *de re* ascriptions of propositional attitudes. The explanatory strategy pursued in this chapter is to address the issue of what is expressed by representational vocabulary, by showing how expressions must be used in order to be functioning as *de re* ascriptions of propositional attitudes. This provides an account of what is necessary to introduce, into an idiom that previously lacked them, representationally explicitating idioms with the expressive capacity provided by terms like 'of' and 'about'.

The tradition distinguishes two readings of, or senses that can be associated with, propositional attitude ascriptions. The usual way of introducing the vocabulary is to say that ascriptions *de dicto* attribute belief in (commitment to) a *dictum* or saying, while ascriptions *de re* attribute belief about some *res* or thing.[2] The distinction is not specific to sentential operators such as 'believes', which express ascriptions of propositional attitude. Consider, to begin with, the assertion, in 1994, of the sentence:

> The president of the United States will be black by the
> year 2000.

Read *de dicto*, this means that the *dictum* or sentence

> The president of the United States is black

will be true by the year 2000. Read *de re*, it means that the *res* or thing, the president of the United States in 1994, namely Bill Clinton, will be black by the year 2000. The difference between the two readings is naturally understood as concerning the order of application of the two operations involved:

tracking through time, as demanded by the tense operator, and picking out an individual, as demanded by the definite description operator. The more plausible reading corresponds to first shifting the context of evaluation of the sentence forward from 1994 to 2000, and then considering what expressions are in that context intersubstitutable with the definite description 'the president of the United States'. The less plausible reading corresponds to determining those intersubstitutabilities first, then seeing what sentences involving those terms will be true in the year 2000.[3]

The present concern, however, is with how this distinction applies to ascriptions of propositional attitude. There are two ways to read:

> Voltaire believed the man from whom Napoleon learned the most about the relations between war and diplomacy was a philosopher-prince.

One reading makes this a false claim, the other true. It is false that Voltaire believed the dictum

> The man from whom Napoleon learned the most about the relations between war and diplomacy was a philosopher-prince.

If you had asked Voltaire, he would have denied that he endorsed this claim;[4] after all, when Voltaire died in 1778, Bonaparte was only nine. In contrast, Voltaire would have endorsed the claim:

> Frederick the Great was a philosopher-prince.

This is a belief *about* the same *res*, since Frederick the Great was in fact the man from whom Napoleon learned the most about the relations between war and diplomacy. That is, Voltaire really did believe *of* or *about* the man from whom Napoleon learned the most about the relations between war and diplomacy (namely Frederick the Great) that *he* was a philosopher-prince.

In ordinary parlance the distinction between *de dicto* and *de re* readings is the source of systematic ambiguity. Sometimes, as in the case above, one of the readings involves a sufficiently implausible claim that it is easy to disambiguate. It is best, however, to regiment usage slightly in order to mark the distinction explicitly. This can be done with little strain to our ears by using 'that' and 'of' in a systematic way. Consider:

> Henry Adams believed the inventor of the lightning rod did not invent the lightning rod.

Thus expressed, the ascription is ambiguous in the same way as the example above. It is quite unlikely that what is intended is the *de dicto*

> Henry Adams believed *that* the inventor of the lightning rod did not invent the lightning rod.

Adams would presumably not have endorsed the *dictum* that follows the 'that'. It is entirely possible, however, that the *de re* claim

> Henry Adams believed *of* the inventor of the lightning rod that *he* did not invent the lightning rod

is true. For since the inventor of the lightning rod is the inventor of bifocals (namely Benjamin Franklin), this latter claim could be true if Henry Adams had the belief that would be ascribed *de dicto* as

> Henry Adams believed *that* the inventor of bifocals (or Benjamin Franklin) did not invent the lightning rod.

In the rest of this chapter the regimentation suggested by these expressions is followed: 'of' marks ascription *de re*, 'that' without 'of' marks ascription *de dicto*, and the absence of 'that' and 'of' marks an undifferentiated, potentially ambiguous ascription.

Quine emphasizes that the key grammatical difference between these two sorts of ascriptions concerns the propriety of *substitution* for singular terms occurring in them.[5] Expressions occurring in the *de re* portion of an ascription—within the scope of the 'of' operator in the regimented versions—are in his terminology *referentially transparent*—that is, coreferential terms can be intersubstituted salva veritate, that is, without changing the truth-value of the whole ascription. By contrast, such substitution in the *de dicto* portion of an ascription—within the scope of the 'that' operator in the regimented versions—may well change the truth-value of the whole ascription. (These formulations carry over directly into the present scheme in the by now familiar way—in terms of substitutions that do and do not preserve doxastic commitment.) Syntactically, *de re* ascriptions may be thought of as formed from *de dicto* ones by *exporting* a singular term from within the 'that' clause, prefacing it with 'of', and putting a pronoun (or other anaphoric dependent) in the original position. Thus the *de dicto* form

> S believes that $\Phi(t)$

becomes the *de re*

> S believes *of* t that $\Phi(it)$.

The important point is that, as the regimentation reminds us, it is *de re* propositional-attitude-ascribing locutions that we use in everyday life to express what we are talking and thinking *of* or *about*. This expressive function dictates the proprieties of substitution that govern expressions occurring in the scope of the 'of'. Since those expressions serve to specify *what* is represented by a belief rather than *how* it is represented, any singular term that picks out the right object is all right; one specification is as good as another. By contrast, expressions that occur within the scope of the 'that'

(either in a pure *de dicto* ascription or in the *de dicto* portion of a *de re* ascription) serve to specify *how* things are represented by the one to whom the belief is ascribed. Thus only expressions the believer acknowledges being committed to intersubstitute, and so treats as ways of recognizing the same object again, are intersubstitutable in these contexts, salva veritate (that is, preserving doxastic commitment).

In perspicuous, regimented-but-recognizable English, what marks this perspectival distinction within the content-specifications of ascriptions is the occurrence of the words 'of' and 'that'. Thought of in this way, the distinction between *de dicto* and *de re* should not be understood to distinguish two kinds of *belief* or even belief-contents, but two kinds of *ascription*[6]—in particular two different *styles* in which the *content* of the commitment ascribed can be *specified*. (To be more precise, this is how one ought to think of what will be distinguished as 'epistemically *weak de re* ascriptions'. Epistemically *strong de re* ascriptions, discussed in Section V, do ascribe a special sort of attitude.)

One way of trying to understand the representational dimension of propositional content is accordingly to ask what is expressed by this fundamental sort of representational locution. What are we *doing* when we make claims about what someone is talking or thinking *about?* How must vocabulary be used in order for it to deserve to count as expressing such *de re* ascriptions? Answering that question in a way that does not itself employ representational vocabulary in specifying that use is then a way of coming to understand representational relations in nonrepresentational terms.

The basic strategy, as indicated above, is to show how the use of this paradigmatic representational locution expresses differences in *social* perspective. The raw materials needed for a deontic scorekeeping account of *de re* ascriptions are already onboard: the attitudes of acknowledging and attributing doxastic commitments, substitutional commitments, and anaphoric commitments. Ascribing a doxastic commitment is explicitly attributing it—that is, attributing it by acknowledging another doxastic commitment. The substitutional commitments that govern the expressions used to specify the content of the commitment ascribed can similarly either be attributed to the one to whom the doxastic commitment is ascribed or be undertaken by the one ascribing it; that social difference of deontic attitude turns out to be what determines whether the ascription is *de dicto* or *de re*. Communication requires that scorekeepers be able to move back and forth between the significance tokenings have as governed by the commitments they themselves acknowledge, on the one hand, and by the commitments they take the speaker to acknowledge, on the other. That is why the expressive power of interpersonal anaphoric recurrence chains is important for securing the possibility of communication. It also turns out that what is expressed by *de re* ascriptions is just the understanding of what is represented by a speaker's claims or beliefs—what they are *about*.

3. *Deontic Scorekeeping Account of What Is Expressed by* De Dicto *and* De Re *Ascriptions*

These points emerge most clearly from looking at the expressive role of ascriptions, for that makes it possible to see the abstract need for two sorts of ways in which expressions could function in specifying the content of an ascribed commitment. Ascriptions are propositionally explicit attributions. Ascriptional locutions such as ". . . claims that . . ." or ". . . believes that . . ." permit one to *say* that one is attributing a commitment with a specified content to a specified individual, something that otherwise one could only *do.* In virtue of playing this expressive role, ascribings essentially involve two different deontic attitudes, to commitments with two different contents. Making an ascription involves doing two different things. Ascribing is *attributing* one commitment (to another), while *undertaking* (acknowledging) a different commitment (oneself).

It follows that expressions occurring in the ascription must somehow specify both the content of the commitment attributed by the ascriber and (thereby) the content of the commitment undertaken by the ascriber. That a particular expression occurs in the specification of the content of the attributed commitment accordingly can have a significance (affect the deontic score) for two different commitments. This makes possible a special sort of ambiguity. It may on occasion be important to be able to resolve such ambiguity; natural languages provide mechanisms for doing so by making explicit the sort of significance the occurrence of a particular expression ought to be understood to have. The expressions occurring in an ascription must specify three sorts of information, which differ in how their occurrence bears on the commitments being attributed and undertaken by the assertion of the ascription. The individual who is the target of the attribution, to whom a commitment is being ascribed, must somehow be indicated. The sort of commitment being attributed—for instance 'claims', 'believes', 'intends', or 'prefers'—must also be indicated. Finally, the content of the ascribed commitment must be specified.

The content-specifying expressions themselves can play two quite different roles, however, and that is where the difference between *de dicto* and *de re* ascriptions arises. Propositional contents, such as those characterizing both the ascribing and the ascribed commitment, are essentially inferentially articulated. But what else a commitment with a particular content commits one to, the committive-inferential consequences of adopting a commitment with that particular content, depends on its deontic context, on what concomitant commitments are available as auxiliary hypotheses or collateral premises. So the occurrence of an expression in the specification of the content of an ascribed commitment might have one inferential significance if evaluated with respect to the collateral commitments of the one to whom the commitment is attributed, and might have quite another if evaluated

with respect to the collateral commitments of the one undertaking the ascriptional commitment. A perspicuous regimentation of ascriptional locutions must notationally encode and display the distinction between expressions (or features of expressions) playing these various roles. Thus the content-specifying expressions should be marked somehow to indicate which sort of significance they are to be understood to have.

In specifying the content of the claim that is attributed by an ascription, a question can arise as to whom the ascriber takes to be responsible for this being a way of *saying* (that is, making explicit) what is believed, the content of the commitment. Consider the sly prosecutor, who characterizes his opponent's claim by saying:

> The defense attorney believes a pathological liar is a trustworthy witness.

The defense attorney may hotly contest this characterization:

> Not so; what I believe is *that* the man who just testified is a trustworthy witness.

To which the prosecutor might reply:

> Exactly, and I have presented evidence that ought to convince anyone *that* the man who just testified is a pathological liar.

If the prosecutor were being fastidious in characterizing the other's claim, he would make it clear who is responsible for what: the defense attorney claims that a certain man is a trustworthy witness, and the prosecutor claims that that man is a pathological liar. The disagreement is about whether this man is a liar, not about whether liars make trustworthy witnesses.

Using the regimentation suggested above, the way to make this explicit is with a *de re* specification of the content of the belief ascribed. What the prosecutor *ought* to say (matters of courtroom strategy aside) is:

> The defense attorney claims *of* a pathological liar that he is a trustworthy witness.

This way of putting things makes explicit the division of responsibility involved in the ascription. That someone is a trustworthy witness is part of the commitment that is *attributed* by the ascriber; that that individual is in fact a pathological liar is part of the commitment that is *undertaken* by the ascriber.[7]

Ascription always involves attributing one *doxastic* commitment and, since ascriptions are themselves claims or judgments, undertaking another. The suggestion is that the expressive function of *de re* ascriptions of propositional attitude is to make explicit which aspects of what is said express substitutional commitments that are being *attributed* and which express substitutional commitments that are *undertaken*. The part of the content

specification that appears within the *de dicto* 'that' clause is limited to what, according to the ascriber, the one to whom the commitment is ascribed would (or in a strong sense should) *acknowledge* as an expression of what that individual is committed to. The part of the content specification that appears within the scope of the *de re* 'of' includes what, according to the *ascriber* of the commitment, but not necessarily according to the one to whom it is ascribed, is acknowledged as an expression of what the target of the ascription is committed to. (This is what the target should, according to the ascriber, acknowledge only in a much weaker sense of 'should'.) What else someone is committed to by a claim that would be expressed in one way is a matter of the substitution-inferential commitments taken to govern the expressions that occur in the sentences the one whose acknowledgments are in question is disposed to assert or otherwise endorse in practice. Thus the marking of portions of the content-specification of a propositional attitude ascription into *de dicto* and *de re* portions makes explicit the essential deontic scorekeeping distinction of *social* perspective between inferences that are underwritten by substitutional commitments attributed and inferences underwritten by substitutional commitments that are undertaken.

So the difference expressed by segregating the content specification of a propositional attitude ascription into distinct *de re* and *de dicto* regions (marked in the regimentation by 'of' and 'that') can be thought of in terms of *inferential* and *substitutional* commitments. Propositional (= assertible) contents are inferentially articulated. Grasping such a content is being able to distinguish in practice what should follow from endorsing it and what such endorsement should follow from. But, as was pointed out above, the consequences of endorsing a given claim depend on what other commitments are available to be employed as auxiliary hypotheses in the inference. The ascriber of a doxastic commitment has got two different perspectives available from which to draw those auxiliary hypotheses in specifying the content of the commitment being ascribed: that of the one to whom it is *ascribed* and that of the one *ascribing* it. Where the specification of the content depends only on auxiliary premises that, according to the ascriber, the target of the ascription *acknowledges* being committed to, it is put in *de dicto* position, within the 'that' clause. Where the specification of the content depends on auxiliary premises that the *ascriber* endorses, but the target of the ascription may not, it is put in *de re* position.

More particularly, the use of expressions as singular terms is governed by *substitution*-inferential commitments. The rule for determining the scorekeeping significance and so the expressive function of *de re* ascriptions being proposed is then the following. Suppose that according to A's scorekeeping on commitments, B acknowledges commitment to the claim $\Phi(t)$. Then A can make this attribution of commitment explicit in the form of a claim by saying

B claims *that* $\Phi(t)$.

If in addition A acknowledges commitment to the identity $t = t'$, then whether or not A takes it that B would acknowledge that commitment, A can also characterize the content of the commitment ascribed to B by saying

B claims *of* t' that $\Phi(it)$.[8]

Again, the question is whose substitutional commitments one is permitted to appeal to in specifying the consequences someone is committed to by acknowledging a particular doxastic commitment. Where in characterizing the commitment the ascriber has drawn out those consequences employing only commitments the ascriptional target would acknowledge, the content specification is *de dicto*. Where the ascriber has employed substitutional commitments the ascriber, but perhaps not the target, endorses, the content specification is *de re*.

One counts as having *undertaken* a commitment wherever it is appropriate for others to *attribute* it. One may be committed to some claims by default, but there are two ways in which what commitments one undertakes can depend on what one *does* or is disposed to do. First, one may *acknowledge* the commitment, paradigmatically by being disposed to *avow* it by an overt assertion. Or one may acknowledge it by employing it as a premise in one's theoretical or practical reasoning. This last includes being disposed to *act* on it *practically*—taking account of it as a premise in the practical reasoning that stands behind one's intentional actions. Second, one may undertake the commitment *consequentially*, that is, as a conclusion one is committed to as an inferential consequence entailed by what one *does* acknowledge. Commitments one acknowledges and commitments one acquires consequentially in virtue of those acknowledgments correspond to two ways of talking about beliefs: one according to which one believes only what one takes oneself to believe, and the other according to which one believes willy nilly what one's beliefs commit one to.

According to one usage, I believe only what I think I believe, what I take myself to believe. I do not believe things behind my back; my sincere avowals are authoritative with respect to what I believe. According to another usage, however, I believe the consequences of my beliefs, whether I think I do or not. For my acknowledged beliefs can commit me to more than I acknowledge; so I can end up with beliefs I do not know I have. Also, my actions, perhaps together with avowed preferences, may commit me to certain claims.

There can be tensions between these two ways of talking about beliefs: a narrow one tied to empirical dispositions to avow, the other more normative and expansive, closing beliefs under a consequence relation not limited by the believer's acknowledgment of it. Indeed, the fact that people often move

back and forth between talk of belief in the empirical sense, which does not involve inferential closure, and talk of belief in the logical or ideal sense, which does, is one of the reasons to prefer talk of commitment to talk of belief—or simply not to believe in beliefs. *De dicto* and *de re* ascriptions correspond to two socially distinct perspectives from which the consequences of doxastic commitments can be extracted. Since the propositional content of a commitment depends on what it is a consequence of and what is a consequence of it, they accordingly also correspond to two different ways of specifying such contents.

II. INTERPRETATION, COMMUNICATION, AND *DE RE* ASCRIPTIONS

1. *Interpretation in Wittgenstein's Sense*

The representational dimension of propositional content (and of conceptual content more broadly) is made explicit—that is, expressed in the form of propositional content—by the use of *de re* ascriptions of discursive commitments. The use of such ascriptions turns essentially on *social* distinctions of doxastic perspective between the ascriber and the one to whom a commitment is ascribed—that is, between the deontic repertoire associated with the one by whom the ascriptional propositional commitment is *undertaken* and the one to whom a propositional commitment is thereby explicitly *attributed*. The social distinction between the fundamental deontic attitudes of undertaking and attributing is essential to the institution of deontic statuses and the conferral of propositional contents. This is, as was pointed out already in Chapter 1, an *I-thou* sociality rather than an *I-we* sociality. Its basic building block is the relation between an audience that is attributing commitments and thereby keeping score and a speaker who is undertaking commitments, on whom score is being kept. The notion of a discursive *community*—a we—is to be built up out of these *communicating* components.

Deontic scorekeeping is the form of understanding involved in communication. It is a kind of interpreting. But it is implicit, practical interpretation, not explicit theoretical hypothesis formation. It is presupposed by the capacity so much as to entertain the claims that would express a hypothesis, evidence for a hypothesis, or conclusions from a hypothesis. Looking at the sort of scorekeeping (and so interpretation) that is made explicit in the use of *de re* ascriptions of propositional attitude—the kind of understanding communication consists in—will make clear how the representational dimension of conceptual content that those ascriptions express is rooted in the social context of communication.

Thinking of ordinary intralinguistic understanding as essentially involving interpretation is objectionable if interpretation is thought of on the model of explicit hypothesis formation.[9] Linguistic understanding depends on inter-

pretation in this sense only in extraordinary situations—where different languages are involved, or where ordinary communication has broken down. Recall in this connection the rehearsal in Chapter 1 of Wittgenstein's argument that norms that are expressed explicitly as rules must be based on norms that are embodied implicitly as proprieties of practice. His pragmatist point is that the distinction between following a rule correctly and following it incorrectly must not be understood in every case as determined by the application of another rule—which could be thought of as an explicit theory or hypothesis concerning the first rule. The normative cannot be understood as rules all the way down. Since he is determined to "restrict the term 'interpretation' [Deutung] to the substitution of one expression of the rule for another,"[10] in his terminology it is accordingly a radical mistake to think of our ground-level practical mastery of linguistic proprieties as consisting exclusively in a capacity to interpret.

This argument is, as was indicated in Chapter 1, one of the fundamental insights from which the present approach proceeds. Yet the deontic scorekeeping account that has been developed here identifies the inferentially articulated sort of understanding characteristic of specifically discursive practice as a kind of interpretation. This sort of interpretation includes, but is not exhausted by, Wittgensteinian interpretation—substituting one expression of a rule for another. Since what matters for the regress-of-rules argument is the propositional explicitness of rules, substitution of one expression of a *claim* for another may be thought of as interpretation in the relevant sense, regardless of whether the claims in question are functioning as rules.[11] The question to be addressed here concerns how much of ordinary skillful intralinguistic practice depends on the capacity to substitute one explicit expression of a claim for another. The answer is that a great deal of ordinary communication between individuals who share a language requires interpretation in the broadly Wittgensteinian sense just indicated. This point will be made by considering four linguistic phenomena, all of which involve cases where we can often just understand what others are saying (without theorizing about it), but where just understanding them crucially depends on being able to substitute one expression of a remark for another.

The general structure that this argument depends on is the fact that a sentence in one person's mouth does not typically have the same significance as that same sentence emerging from another person's mouth, even where there is as much sharing of the language and as much mutual understanding as one likes. The fundamental reason is the kind of things claims and concepts are. As was emphasized above, the inferential articulation of conceptual contents is such that what someone becomes committed to by uttering a certain expression can be assessed only against a background of collateral commitments available as auxiliary hypotheses that can be brought in as other premises in drawing the inference. Even where people share a language (and so their concepts), which is the standard case of communication, there

will still be some disagreements, some differences in the commitments that people have undertaken. We each embody different perceptual and practical perspectives and so will never have exactly the same doxastic and practical commitments. In any case, if two interlocutors did (per impossibile) have exactly the same beliefs and desires, communication would be superfluous; so that case need not be considered here. As long as there are differences in the collateral set of commitments with respect to which the content of the claim expressed by a sentence needs to be assessed, the sentence in one mouth means something different from what that same sentence means in another mouth. So even in the smooth untroubled cases of communication, if you want to understand what I say, you have to be able to associate with it a sentence that in your mouth expresses the claim that the sentence I uttered expresses in mine. For your understanding it (your knowing what I have committed myself to) involves your being able to trace out the inferences that claim is involved in, the evidential significance of what I have said, in order to know what I am committing myself to. This means knowing how it could function as evidence for you, as well as for me, what claims its endorsement would preclude you, as well as me, from being entitled to, and so on. Apart from that capacity, you cannot extract information from what I say and cannot be said to understand it.

2. Four Linguistic Phenomena That Involve Interpretation

There are four phenomena in linguistic practice where it is clear this sort of capacity to interpret is of the essence of smooth conversational coping and practical grasp of the meanings expressed by the utterances of others. The first is the use of *personal pronouns*. When someone else says, "I'm talking," for a member of the audience to draw inferences from that remark and check out what would be evidence for it and what it would be evidence for, that auditor must be able (no doubt without thinking about it and smoothly, as a matter of unconscious skill) to substitute 'you' or 'he' or 'John' for 'I'. For that last expression in the mouth of the auditor does not mean what it means in the mouth of the speaker. Those tokenings are not corecurrent and so are governed (at least counterfactually) by different substitutional commitments. Understanding the expression requires being able to make that substitution of one expression of the claim for the other. The use of personal pronouns is a fundamental part of our shared language. So here is a simple case where sharing a language, being able to understand someone, consists in part in being able to make this substitution—that is to *interpret*, in Wittgenstein's sense. For the words in the one mouth do not have the same significance that they would in the other mouth.

A similar case is that of *demonstratives*. If someone says, "This is blue," one to whom the remark is addressed *might* (if in a favorable position) be able to use another token of the same type to make the same claim. But

typically, members of the audience will not be in a position to do that. If the speaker uses a demonstrative, the audience will have to use a pronoun, a description, or some other demonstrative (even if it is only 'that'). There is no requirement that in order to understand the original remark the audience must be in a position to make demonstrative reference to the same object at all. If the speaker used some demonstratives yesterday and someone in the audience wants today to rehearse the inferences in the vicinity of the commitments undertaken thereby, exercising an understanding of what was said, that auditor is not typically going to be in a position to redemonstrate those things. Other expressions must be available to be used to pick them out.

As was pointed out in Chapter 7, it has to be possible to pick up the demonstrative reference with a pronoun. When someone says, *"This* is blue," the audience must be able to say something like, "No *it* isn't either"—where 'it' is a pronoun that is anaphorically dependent on the speaker's tokening. Such anaphoric dependents continue a recurrence tree that gives everybody a repeatable way of expressing the content the speaker expressed, even those who are not in a position to make the same demonstration. There cannot be demonstratives without anaphora; such 'bare' or inaccessible demonstratives would be of no use in communication (nor, indeed, even for the thought of an isolated individual). The audience, in order to make use of the speaker's demonstratives, has to be able to pick them up with a pronoun. That is the only form in which they are repeatable. The current point is that the one understanding the speaker's remark must be able to *substitute* a *different* expression for the one the speaker used. Interlocutors must be able to perform this sort of interpretation—in the Wittgensteinian substitutional sense—in order to understand even in the smoothest intralinguistic conversation.

Those two examples, personal pronouns and demonstratives, are very straightforward and ordinary, even trivial cases where intralinguistic interpersonal understanding requires the ability to interpret in Wittgenstein's substitutional sense. The next two cases are more substantive. The important thing about the first two is that for both personal pronouns and demonstratives, understanding them *consists*, at least in part, in being able to substitute other expressions for them. The other two cases to be considered are *speaker's reference* and being able to give *de re* specifications of the content of someone's claim (of the expression of the belief expressed).

These are intimately linked. Both of them arise precisely because there are differences in repertoires of collateral commitments or background assumptions in any communicational situation. Conversation has the significance of communication only where the commitments of speakers and audience differ. It is because of this difference in background commitments that it is possible for one interlocutor to become entitled to a new commitment by listening to what someone else says—or to challenge the speaker's entitlement to the commitment that speech act expresses. This is the communica-

tive point of the assertional practice in virtue of which we are able to entertain propositional contents at all. But often in assigning or even just hearing the content of what someone says, it is necessary to take account of those differences.

The previous chapter discussed the classic sort of speaker's reference case, which arises when the one overhearing the remark "The man in the corner drinking champagne is a good husband" both *attributes* to the speaker a collateral commitment that would be undertaken by asserting "John is in the corner and is drinking champagne" and *undertakes* a collateral commitment that would be expressed by a sentence such as "John is in the corner and is drinking ginger ale." In order to understand the commitment the speaker intended to be understood as undertaking—and so potentially to extract information from what was said—the audience must be able to substitute for the defective description 'the man in the corner drinking champagne' some expression that, according to the substitutional commitments endorsed by the audience, succeeds in picking out the same individual. That individual is the one the speaker is talking about. (Though of course there is also a sense in which willy-nilly the speaker may have made a claim about someone else, to whom his expression ^ssemantically^s refers.) The expression the audience substitutes could be an explicit, anaphorically indirect definite description, such as 'the man the speaker is talking about'; or the same effect could be achieved implicitly by using a pronoun such as 'he', if one can arrange for it to be understood as having as its antecedent some other tokening the speaker *could* have used to indicate the same man, one that is not subject to the same conflict of doxastic perspective, such as 'the man in the corner', 'John', or 'the man with bubbly liquid in his glass'; or one might use one of those expressions directly. An interpreter's capacity to understand the remark and extract information across the gap between the commitments undertaken and those attributed depends on that scorekeeper being able smoothly, and perhaps completely unconsciously,[12] to substitute an expression that in the scorekeeper's mouth would undertake the commitment being attributed to the speaker.

The substitutional commitments that are *implicit* in the sort of interpretive competence required to be able to take someone's remark as having its significance determined in the manner characteristic of speaker's reference are made *explicit* in *de re* ascriptions of commitment. Indeed, the expressive capacity provided by ascriptions of propositional attitude that specify the content of the attributed commitment in the *de re* way is precisely what is needed to *say* what the difference is between interpreting a remark according to its speaker's reference and interpreting it according to its semantic reference. Thus, suppose there is someone, Bob, of whom the speaker is unaware, who is also in the corner and actually is drinking champagne, unlike ginger-ale-guzzling John, whom the speaker takes to be drinking champagne. The two interpretations then correspond to taking the speaker to have said (with-

out meaning to) *of* Bob that *he* is a good husband and taking the speaker to have said (using a false description) *of* John that *he* is a good husband. The difference between these two *de re* specifications of the content expressed is a matter of which substitutional commitments undertaken by the ascriber are taken to govern the speaker's tokening 'the man in the corner drinking champagne'.

In the first case, that tokening is treated as corecurrent with lexically cotypical tokenings in the ascriber's mouth. In the second, as indicated earlier (see 7.5.6 above), it is treated as anaphorically dependent on other tokenings, perhaps demonstrative ones, that the speaker could (according to the interpreting ascriber) have used equally well to express the commitment being acknowledged. This latter case involves substituting expressions of other types for the expression used, and it is these expressions that occur in the *de re* specification of the content. Indeed, (weak) *de re* ascriptions are formed by substituting, for locutions the target of the ascription might use in expressing the content of the commitment, locutions that the *ascriber* is committed to being intersubstitutable with them. Thus being able to offer *de re* characterizations of the contents of the commitments of other interlocutors requires interpretation in Wittgenstein's substitutional sense.

3. Communication, Truth Conditions, and De Re Specifications of Propositional Content

Unless one has this substitutional interpretive capacity, which is expressed explicitly in *de re* specifications of the contents of ascribed commitments, one would not be able to understand what others were saying—even in languages that lacked the personal pronouns, demonstratives, and the other officially token-reflexive structures considered above. The cognitive uptake or grasp of content that is the measure of the success of communication consists in auditors being able to move between their own doxastic perspective and that of the speaker in just the way expressed by *de re* ascriptions—that is, by inferentially exploring the significance the *propositional* commitment *attributed* to the speaker has in the context of the *substitutional* commitments *undertaken* by the audience who attributes it. There is no communication apart from this sort of interpretation. Being able to understand what others are saying, in the sense that makes their remarks available for use as premises in one's own inferences, depends precisely on being able to specify those contents in *de re*, and not merely *de dicto*, terms.

If the only way I can specify the content of the shaman's belief is by a *de dicto* ascription

> He believes *that* drinking the liquor distilled from the bark of that kind of tree will prevent malaria,

I may not be in a position to assess the truth of his claim. It is otherwise if I can specify that content in the *de re* ascription

> He believes *of* quinine that malaria can be prevented by drinking *it*,

for 'quinine' is a term with rich inferential connections to others I know how to employ. If he says,

> The seventh god graces us with his presence,

I may not know what to make of his remark. Clearly he will take it to have consequences that I could not endorse; so nothing in my mouth could *mean* just what his remark does. But if I come to believe (perhaps by being told) that the seventh god is the sun, and that his grace is sunshine, then I can specify the content of his report in a more useful form:

> He claims *of* the sun that it is shining,

or both *of* the sun and *of* shining, that one is doing the other, or indeed as saying *of* the claim that the sun is shining *that* it is true. These are forms I can extract *information* from, that is, can use as premises that I can reason with. Again, suppose a student claims that

> The largest number that is not the sum of the squares of distinct primes is the sum of at most 27 primes.

The student may have no idea what number that might be, or may falsely believe it to be extremely large. But if I know that

> 17,163 is the largest number that is not the sum of the squares of distinct primes,

then I can characterize the content of his claim in *de re* form as:

> The student claims *of* 17,163 that it is the sum of at most 27 primes,

and I can go on to draw inferences from that claim, to assess its plausibility in the light of the rest of my beliefs. (It is true, but only because *all* integers are the sum of at most 27 primes.) Identifying what is being talked *about*, what is represented by it, permits me to extract information across a doxastic gap. Thus, being able to substitute *de re* specifications of the content of a claim for *de dicto* ones is something anyone must be able to do, at least implicitly, in order to understand people whose beliefs are different (not massively or radically different, but just different on particular points).

De re content specifications not only indicate what a claim represents or is about, they are the form in which the *truth conditions* of claims are expressed. The sense in which they present a claim's representational content (the information it makes available to another interlocutor) is that they

express—from the point of view of the interpreter attributing the commit-ment—what would be required for the claim in question to be *true*. Recall the account (in Chapters 3, 4, and 5) of the scorekeeping practices that underlie truth assessments. The central context in which such assessments classically arise is attributions of *knowledge,* a normative status integral to the social practice of making claims and judgments. According to the tradi-tional account, knowledge is justified true belief. Transposed into a specifica-tion of a normative status something could be taken to have by interlocutors who are keeping score of each other's commitments and entitlements, this account requires that in order for it to be *knowledge* that a scorekeeper takes another to have, that scorekeeper must adopt three sorts of practical attitude:

> First, the scorekeeper must *attribute* an inferentially articu-lated, hence propositionally contentful, *commitment.*

This corresponds to the *belief* condition on knowledge.

> Second, the scorekeeper must *attribute* a (perhaps, but not nec-essarily inferentially inherited, but necessarily heritable) *entitle-ment* to that commitment.

This corresponds to the *justification* condition on knowledge. What is it that then corresponds to the *truth* condition on knowledge? For the scorekeeper to take the attributed claim to be true is just for the scorekeeper to endorse that claim. That is:

> Third, the scorekeeper must *undertake* the same propositional commitment attributed to the candidate knower.

Undertaking a commitment is adopting a certain *normative stance* with respect to a claim; it is not attributing a property to it. The classical meta-physics of truth properties misconstrues what one is doing in endorsing the claim as *describing* it in a special way. It confuses *attributing* and *undertak-ing* or *acknowledging* commitments, the two fundamental social flavors of deontic practical attitudes that institute normative statuses. It does so by assimilating the third condition on treating someone as having knowledge to the first two. Properly understanding truth talk in fact requires understanding just this difference of social perspective: between *attributing* a normative status to another and *undertaking* or adopting it oneself.

This is just the distinction that underlies the use of *de re* ascriptions. As the regimented form considered here emphasizes, they mark overtly the distinction between the *doxastic* commitment that is *attributed* and the *substitutional* commitments that are *undertaken* by the attributor. More specifically, they indicate whether responsibility for the words used being an appropriate expression of the content of the claim ascribed is attributed or undertaken by the ascriber. The connection between the account of ". . . is true" as an anaphoric prosentence-forming operator and the scorekeeping

account of truth assessments becomes evident in *de re* ascriptions in which the entire content-specifying clause is exported to *de re* position:

> Senator McCarthy believed *of* the first sentence of the Communist Manifesto that *it was true.*

The first sentence of the Communist Manifesto says that the specter of Communism is haunting Europe. That is undoubtedly something McCarthy believed, but it is quite unlikely that he had ever read the Communist Manifesto, and also quite likely that if asked whether he believed any of the claims of that manifesto, he would have denied that he did. Responsibility for specifying the content of his belief in this way is accordingly properly undertaken by the ascriber of the belief, rather than attributed along with the belief. The regimented *de re* form requires then that the whole content-specifying expression be exported to the scope of the 'of', and that an ascription-structural anaphoric dependent of it be left in its place inside the scope of the 'that'. Since what is exported is a sentence nominalization, such an anaphoric dependent is a prosentence. No truth *assessment* is being offered by such a use of ". . . is true," for that prosentence appears embedded here, in a context in which it functions only to express the undertaking of a *substitutional* commitment; the doxastic commitment is attributed, not undertaken. Nonetheless, in this expressive role the prosentence-forming operator is used to specify what must be the case for the attributed belief to have been true (according to the attributor). It would have been true just in case the first sentence of the Communist Manifesto were true; that is, according to a well-informed ascriber, just in case the specter of Communism *were* haunting Europe. The ascription-structural anaphoric connection between the exported expression (in *de re* or undertaking position) and its anaphoric dependent (in *de dicto* or attributing position) mirrors within an ascription in one mouth the interpersonal anaphoric connections among expressions in different mouths that make possible the uptake of speaker's reference and the extraction of representational, truth-conditional content across doxastic boundaries (as discussed in Chapter 7).

The role of *de re* specifications of propositional content in expressing the truth conditions of ascribed commitments is not restricted to cases where the entire content-specification is exported to *de re* position. For ordinary *de re* ascriptions, in which singular terms are exported, specify what a belief (or the sentence expressing it) *represents* in the sense of what it must be *true of* if it is to be true at all. A scorekeeper who takes it that the inventor of bifocals is the inventor of the lightning rod will take it that believing that the inventor of bifocals spoke French is believing *of* the inventor of the lightning rod that *he* spoke French. If it is true that the inventor of bifocals spoke French, then, according to that scorekeeper, it is *true of* the inventor of the lightning rod that *he* spoke French. In saying which individual a belief represents or is about, *de re* content specifications indicate the individual

whose properties must be consulted in order to determine the *truth* of the belief (all from the point of view of the ascriber). Thus a scorekeeper who in addition endorsed the substitutional commitment made explicit by the (provincially tempting and occasionally enunciated) claim that French is the only language in which words appear in exactly the order in which the corresponding ideas are arranged in thought could specify the truth conditions of the belief that the inventor of bifocals spoke French by saying that it is true just in case it is true *of* the inventor of the lightning rod and *of* the only language in which words appear in exactly the order in which the corresponding ideas are arranged in thought that the first spoke the second. The point is that scorekeepers must use the auxiliary hypotheses provided by their own commitments in assessing the truth of the beliefs they attribute or entertain, for taking-true a claim is just endorsing it. *De re* specifications of their contents are just those that employ the substitutional commitments undertaken by the scorekeeper or interpreter, rather than those attributed to someone else, and so are the specifications that present propositional contents in a form apt for assessments of *truth*. Thus the *information* a claim potentially communicates, its *representational* content in the sense captured by its truth conditions, is what is expressed by *de re* specifications of propositional content.

Extracting information from the remarks of others requires being able to do the sort of substitutional interpreting that is expressed explicitly in offering *de re* characterizations of the contents of their beliefs—that is, to be able to tell what their beliefs would be true *of* if they were true. It is to grasp the *representational* content of their claims. The most important lesson of the deontic scorekeeping account of the use of *de re* ascriptions is that doing this is just mastering the *social* dimension of their inferential articulation—the way in which commitments undertaken against one doxastic background of further commitments available for use as auxiliary hypotheses can be taken up and made available as premises against a different doxastic background. Having the sort of representational significance that is analyzed here in terms of relations between the socially distinct perspectives of an interlocutor who undertakes a commitment and a scorekeeper who attributes it is an essential component of propositional content, and hence conceptual content generally. It is where the notion of propositional contents as truth conditions—as depending for their truth on the facts about the objects they represent—gets its grip. This is the sense of 'represents' that is fundamental to the intentionality of thought and talk. The analysis of the uptake of this sort of representational content in terms of the kind of social substitution-inferential interpretation deontic scorekeeping consists in is what stands behind the claim that the discursive practice that embodies such intentionality must be understood as essentially a *social* practice. Propositional, and so conceptual, *understanding* is rooted in the *interpretation* that *communication* requires, rooted in the *social* practice of deontic *scorekeeping*.

4. Two Other Senses of 'Represents'

The intentional aboutness of thought and talk—the representational dimension that arises out of the interaction between the social and the inferential articulation of discursive practice—should be distinguished from two other ways in which things can be treated in practice as having a broadly representational content. These are two ways in which some things or states of affairs can be treated as representing others. The first does not presuppose grasp of the specifically propositional contents conferred by distinctively discursive practice; the second does.

The preconceptual variety can be implicit in the practical capacity to navigate by a map. It is a matter of reliable differential responsive dispositions—altering one's behavior with respect to the *terrain* in response to features of the structure being employed as a *map*. In its simplest form, this sort of taking or treating as a representation need be no more complicated than seeking shelter from impending rain upon sighting a certain sort of cloud. It is based on treating one thing as a sign of another in the sense of acting in a way appropriate to the latter in response to the former. Navigating by a map need be no more than a systematic constellation of such differential responsive dispositions. Creatures evidently need not be able to talk in order to treat things as representations in this sense. This is a primitive ability, without which it would not be possible to engage in the more sophisticated practices linking inferentially and socially articulated propositionally contentful perception and action that provide the indispensable framework of the empirical and practical aboutness of our judgments. There is no reason to quibble about whether the implicit practical preconceptual attitudes of treating something as a representation by suitably systematic dispositions to differential responses embody some sort of intentionality, so long as the distinction between such attitudes and propositionally contentful ones is kept firmly in mind. The account of socially and inferentially articulated deontic scorekeeping practices, including the discussion of the substitutional and anaphoric substructures of that articulation, have been aimed precisely at saying what must be added to the primitive preconceptual representational ability in order to arrive at the full-blooded conceptual one.

Within the context of genuinely discursive practice, which institutes intentionality properly so-called, a propositionally explicit form of the primitive, practically implicit taking or treating of something as a representation becomes possible. At this level it is possible to endorse inferences from claims about a representing thing to claims about a represented thing. This is how we concept users typically use maps: *inferring* from the noninferentially acquired claim that there is a wavy blue line between the two black dots on the map *that* there is a river between the two cities in the terrain. This capacity explicitly to use one thing or state of affairs as a representation of another is parasitic on our capacity to make claims or judgments *about*

maps and pictures, on the one hand, and *about* what is mapped or pictured, on the other. On pain of an infinite regress we cannot understand what it is for our thought and talk to represent objects and states of affairs as a matter of our taking or treating it *as* representing them in this sense. One way of understanding the strategy of agent semantics, discussed in Chapter 3, is as assimilating the representational content of *talk* to this model of *explicit* treating as a representation, while appealing to the prior representational content of *thought*, which must be otherwise explained. A prominent candidate for the account that is then required of the more fundamental intentionality of thought is then to assimilate it to the sort of primitive preconceptual taking or treating as a representation by noninferential differential responsive dispositions considered in the previous paragraph. The deontic scorekeeping strategy, by contrast, interposes an intermediate level of socially and inferentially articulated practices that provide the context for the propositional and so conceptual contentfulness of both thought and talk—practices that presuppose the lower level of implicit, merely differentially responsive taking or treating as a representation, and are in turn presupposed by the upper level of explicit inferential taking or treating as a representation.

In the previous section it was claimed that the primary representational locution in ordinary language—the one we use to talk about the representational dimension of our thought and talk, to specify what we are thinking and talking *about*—is *de re* ascriptions of propositional attitude. It is the role they play in such ascriptions that gives their meanings to the 'of' or 'about' we use to express intentional directedness. It has also been claimed that the expressive role of these locutions is to make explicit a distinction of social perspective involved in keeping our books straight on who is committed to what. The social dimension of inference involved in the communication to others of claims that must be available as reasons both to the speaker and to the audience (in spite of differences in collateral commitments) is what underlies the representational dimension of discourse. Beliefs and claims that are *propositionally* contentful are necessarily *representationally* contentful because their inferential articulation essentially involves a *social* dimension. That social dimension is unavoidable because the inferential significance of a claim (the appropriate antecedents and consequences of a doxastic commitment) depends on the background of collateral commitments available for service as auxiliary hypotheses. Thus any specification of a propositional content must be made from the perspective of some such set of commitments.

One wants to say that the *correct* inferential role is determined by the collateral claims that are *true*. Just so; that is what *each* interlocutor wants to say—each has an at least slightly different perspective from which to evaluate inferential proprieties. Representational locutions make explicit the sorting of commitments into those attributed and those undertaken, without which communication would be impossible given those differences of per-

spective. The *representational* dimension of propositional contents reflects the *social* structure of the *inferential* articulation conferred by their role in the game of giving and asking for reasons.

III. *DE RE* ASCRIPTIONS AND THE INTENTIONAL EXPLANATION OF ACTION

1. From Theoretical to Practical Reasoning

The intentional states characteristic of sapience are distinguished by their propositional contentfulness. When attributions of such states are made propositionally explicit, those contents must be specified as part of the ascription. There are two different forms in which the propositional contents of ascribed intentional states can be expressed: *de dicto* and *de re*.[13] The different expressive roles played by the vocabulary employed in these two sorts of content-specifications reflect differences in social perspective that are already implicit in the underlying practical attitudes of attributing and acknowledging propositionally contentful deontic statuses. To understand intentional states or deontic statuses as having specifically *propositional* semantic contents is to understand the norms governing their pragmatic significance as *inferentially* articulated. The *de dicto* and *de re* styles of specifying such contents are the explicit expression of the *social* dimension of that inferential articulation.

The two ways of expressing those contents arise because the inferential role of a claimable content appears differently from the various points of view provided by the background beliefs on the basis of which the one acknowledging a commitment and those attributing it assess its significance by projecting consequences, antecedents, and incompatibilities. This social articulation of inferential role—that endorsing a semantic content has systematically different pragmatic significances for different parties to a conversation—is essential to the discursively fundamental communicative practice of giving and asking for reasons. Since *de re* ascription of propositional attitude is the primary locution serving to make representational relations explicit in natural languages—providing the use in virtue of which words like 'of' and 'about' come to express intentional directedness—this social articulation is also what underlies the representational dimension of the conceptual content of intentional states. *Representational* content can be understood in terms of *inferential* content, provided the *social* dimension of inferential practice is properly taken into account.

Doxastic commitments, deontic statuses of the sort that corresponds to the intentional state of belief, are caught up in two kinds of reasoning. The discussion of what is expressed by *de re* ascriptions of doxastic commitments has so far been directed at *cognitive* or theoretical reasoning: reasoning whose premises and conclusions are doxastic commitments. Presystematically, this

is reasoning that leads from beliefs to further beliefs. But there is also *practical* reasoning: reasoning whose conclusions are practical, rather than doxastic, commitments. Practical commitments are commitments to *act*, to make-true a certain propositional content. Chapter 4 discussed how to make sense of such practical discursive commitments and the practical inferences they are involved in, in the context of a prior story about the practices that institute doxastic discursive commitments, and how these notions can then be used to explain the concepts of *action* and *intentional explanation.* Practical discursive commitments are deontic statuses of the sort that corresponds to the intentional state of intention, rather than belief. Presystematically, basic practical reasoning accordingly leads from beliefs to intentions. Inferential commitments that are endorsements of the propriety of patterns of practical reasoning articulate the significance both of individual preferences or desires and of norms of various sorts. These can appear as further premises provided the language in question has sufficient expressive power to make explicit the endorsement of those patterns of practical inference.

To treat a performance as an action is either to take it to be the result of exercising a reliable differential disposition on the part of the agent to respond to acknowledgment of a practical discursive commitment with a certain content by producing a performance with a certain character or to take it to be itself the acknowledgment of such a commitment. As Anscombe and Davidson have taught us, for a performance to be an action is for it to be intentional under some description; the specification under which the performance is intentional is determined by the practical commitment the performance acknowledges. To offer an intentional explanation of an action is then to attribute a sample piece of practical reasoning, whose conclusion is the practical commitment or intention with which the action is taken to be performed. Action depends on noninferential capacities reliably to *make-true* claims of various kinds as, or in response to, the acknowledgment of practical commitments, as perception depends on noninferential capacities reliably to *take-true* claims of various kinds, to respond to environing situations by acknowledging doxastic commitments.

Reliability in the practical sphere of action, as in the cognitive sphere of perception, is assessed by a scorekeeper on the basis of the match between the contents of noninferentially eliciting or elicited commitments *attributed* to the agent or perceiver, on the one hand, and the contents of the commitments regarding what the agent is making- or taking-true that are *acknowledged* by the scorekeeper who is assessing such reliability, on the other. Attributing cognitive reliability involves assessments of the *truth* of judgments that are noninferentially elicited; attributing practical reliability involves assessments of the *success* of performances that are noninferentially elicited (the truth of the claim one sought to make true). Assessments of both of these sorts involve hybrid deontic attitudes on the part of the score-

keeper—*attributing* one commitment and *undertaking* another. Such attitudes are articulated according to the fundamental *I-thou* social structure of deontic scorekeeping practices. *De re* ascriptions of propositionally contentful intentional states (or deontic statuses) express just such hybrid scorekeeping attitudes; they incorporate the undertaking of substitutional commitments as well as the attribution of doxastic ones. That is why the *truth* conditions of ascribed beliefs are to be determined, according to the ascriber, by the substitutional commitments that govern expressions occurring transparently (that is, inside the scope of the 'of' when the ascriptions are suitably regimented) in *de re* specifications of the content of the ascribed belief. For these are the substitutional commitments *undertaken* by the scorekeeper, who is ascribing, and so *attributing*, a doxastic commitment.

It might accordingly be expected that *de re* ascriptions of commitments would serve an expressive role in making explicit the representational dimension of assessments of the *success* of intentions and actions that corresponds to their role in making explicit the representational dimension of assessments of the *truth* of beliefs and claims. So they do. There are two kinds of intentional explanation, roughly those that explain what an agent *tried* to do, and those that explain what an agent *succeeded* in doing. In each case the performance is made intelligible by attributing to the agent a constellation of commitments structured as a piece of practical reasoning. That structure is provided by the attribution of an implicit, practical-inferential commitment linking premises the scorekeeper takes the agent to be committed to with a conclusion consisting of a practical commitment whose acknowledgment is taken to have been, or resulted noninferentially in, the performance in question. The difference between the two kinds of intentional explanation lies in how the action being explained is specified, and so in the content of the governing practical commitment.

2. De Dicto *and* De Re *Ascriptions Underwrite Two Different Kinds of Intentional Explanation*

The first form of intentional explanation is the kind discussed in Chapter 4, which yields a specification under which the performance being explained is taken by the scorekeeper to be intentional. Why, it might be asked, did Nicole discharge her rifle in the direction of the animal in front of her? She believed that the animal in front of her was a deer, and she desired to shoot a deer, and she believed that if she discharged her rifle then and in that direction, she would shoot the animal in front of her, and so shoot a deer. Those commitments, in the absence of competing ones, provided sufficient reasons for the formation of an intention that she shoot the animal in front of her by discharging the rifle then and in that direction. This last specification is one under which her performance was intentional, and hence an action. For this is a specification under which she *acknowledges* the

practical commitment, and it is that attitude that noninferentially elicits the performance.

But Nicole's performance admits of other descriptions. As Davidson reminds us under the heading of "the accordion effect," any consequence of the performance can be appealed to in order to specify what Nicole did—although most of these will be descriptions under which it was not intentional. Suppose that another such description of Nicole's performance is that she killed a cow. That is something she succeeded in accomplishing, something she did, not merely something that happened, though it is not (under that description) something she intended to do. Nonetheless, her action can still be made intelligible under that description by offering an intentional explanation of it. The animal in front of her was in fact a cow, not a deer. So Nicole believed *of* a cow (namely the animal in front of her) that it was a deer, and she desired to shoot a deer, and believed that if she discharged her rifle then and in that direction, she would shoot the animal in front of her, and so shoot a deer. Those commitments, in the absence of competing ones, provided sufficient reasons for the formation of an intention of, about, or directed at a cow (the animal in front of her) that she shoot a deer by shooting it (that animal) by discharging the rifle then and in that direction.

As the awkward expression of the intention appealed to by this intentional explanation is supposed to suggest (via the ascriptional regimentation put in play earlier), the difference between the two sorts of intentional explanation is just whether the contents of the attributed commitments it appeals to are specified *de dicto* or *de re*. *De dicto* ascriptions specify the contents of attributed commitments in terms that, according to the ascriber, the one to whom they are ascribed would acknowledge as specifications of the contents of commitments undertaken. From Nicole's point of view (according to the ascriber), her intention is that she shoot a deer and her belief is that the animal in front of her is a deer. She does not acknowledge these as states that are of, about, or directed at a cow. The claim that they are cow-representing states, that she represents a cow as a deer, is part of the commitment the ascriber *undertakes*, not part of what the ascriber *attributes*. Thus an intentional explanation whose point is to show things from Nicole's point of view—to show what she was *trying* to do, under what specifications she acknowledged the performance as her own, under what descriptions it was intentional—employs ascriptions that express the contents of the attributed commitments in the *de dicto* style.

In contrast, an intentional explanation whose point is to show why the action turned out as it did—why what Nicole intentionally did had the *consequences* it did, accomplished what it did, succeeded or failed as it did—is provided from the point of view of the ascriber. Intentional explanations of this sort employ ascriptions that express the contents of the attributed commitments in the *de re* style. That Nicole's belief that the animal in front of her was a deer was a belief *of* a cow that *it* was a deer in front of her

is just what is needed to explain why Nicole shot a cow. Corresponding to these two specifications of the content of Nicole's belief are two specifications of the content of Nicole's intention. She intended that she shoot a deer, but in the actual circumstances, where she intended to do that *by* shooting the animal in front of her and the animal in front of her was in fact a cow, she thereby (without realizing it) intended *of* a cow that she shoot *it*.

The basic form of intentional explanation is, on this account, the one that employs *de dicto* specifications. For according to this way of understanding the distinction, behind every *de re* ascription there must be a *de dicto* ascription (though the ascriber need not always be in a position to produce it in order to be entitled to the *de re* ascription). For to be committed to the claim that S believes *of* t that $\Phi(\text{it})$ is to be committed to there being some t' such that $t = t'$ and S believes that $\Phi(t')$. Since the distinction has to do with how the *content* of the ascribed commitment is specified, rather than with what sort of a commitment it is, the story about intentions is no different from that about belief.[14] In this case it is because the animal in front of Nicole was a cow and Nicole intended that she shoot the animal in front of her that she intended *of* a cow that she shoot *it*. Shooting a cow is something she *succeeded* in doing, although it is not something she was *trying* to do—that is, this *de re* way of specifying the content of her practical commitment is not one under which she acknowledges it. Shooting a deer is something she was *trying* to do but *failed*.

Shooting the animal in front of her is something that she both tried to do and succeeded in doing. Of course she might have failed even in this, if her aim had been bad. But then she would still have succeeded in firing the gun. If she failed even in this because of wet powder, she would have succeeded in pulling the trigger, or if it had jammed, in moving her finger. Davidson takes it that in each case of intentional action, there will be *something* the agent both intended to do and succeeded in doing. That is, he insists that some description under which a performance is intentional be *true* of that performance. The characterization of intentional action offered here does not involve such a commitment. As long as the performance was produced by exercising a reliable noninferential capacity to respond differentially to the acknowledgment of a practical commitment with a given propositional content by producing performances of which that content is true,[15] it does not matter if that particular case is one in which the usually reliable capacity misfired. If intending to pick up the bread I spill the wine, there is on this line no theoretical commitment to insisting that there must be something I intended to do that I succeeded in doing, for instance moving my arm. In this case my performance is intentional under a description, 'picking up the bread', that is not true of it.[16] Such a characterization makes no sense for intentions-in-action, where the performance *is* the acknowledgment of the practical commitment. But even an *unsuccessful* performance that is the result of exercising a usually reliable differential disposition to respond to the

acknowledgment of a prior commitment by producing a suitable performance can, just because of its provenance, qualify as intentional, according to this non-Davidsonian line.

Picking up the bread entails moving my arm, so a practical commitment to doing the one entails a practical commitment to doing the other. Undertaking commitment to picking up the bread accordingly involves undertaking a commitment to move my arm. But what corresponds to intentions is practical commitments not just *undertaken* but *acknowledged,* just as the beliefs that are efficacious in deliberation are not all those consequentially undertaken but only those acknowledged. In each case the attitude of acknowledgment need not be preserved by committively good inferences. Picking up the bread also entails contracting a variety of muscles in complicated ways, sending electrochemical signals through various synapses, and so on, but doing so need not in any ordinary sense be part of my intention. Even though in one sense I can move my arm only *by* doing these things, I can *intentionally* contract those muscles in just that way and so on only *by* reaching for the bread. Since in these circumstances picking up the bread *is,* in fact, contracting my muscles in those particular ways, in intending that I pick up the bread, I am intending *of* something I can do only by contracting my muscles in those particular ways that I do *that.* The identity is not one that I need to acknowledge in order to have the intention. A theorist who does acknowledge it, however, is in a position to offer a *de re* specification of the content of my intention. Just such specifications may need to be appealed to in explaining why my attempt to pick up the bread succeeded—or failed because of the unnoticed interposition of the wine bottle. Underlying such *de re,* or attributor-centered, ascriptions, however, must be some *de dicto,* or agent-centered, ascription. These specify the content of the intention in terms the agent would (according to the ascriber) acknowledge, and so cannot rely on inferential commitments undertaken by the ascriber but not attributed to the agent.

For individuals can be trained to have reliable noninferential dispositions to respond differentially only to their own *attitudes* of acknowledging commitments, not to their *statuses* of being committed, insofar as these outrun those they acknowledge. In this sense only such acknowledgments are causally efficacious in an agent's behavioral economy. In another sense, of course (a sense that the analysis of *de re* ascriptions as expressing hybrid deontic attitudes based on *de dicto* ascriptions seeks to explain), actions are causally responsive to features of the contents of intentions that are expressed by *de re* specifications, which outrun the agent's dispositions to acknowledge commitments. For what Nicole does, shooting a cow, depends on the fact, unacknowledged by her, that the animal in front of her is a cow. The two sorts of intentional explanation correspond to reconstructing the context of *deliberation,* in which only what the agent is aware of or acknowledges is relevant, and reconstructing the context of *assessment* of action, in which various

factors that are not accessible to the agent can be invoked in explaining how things turned out. Beliefs ascribed *de dicto* answer to the believer's dispositions to acknowledge, avow, or act according to the belief. Beliefs ascribed *de re* answer to what would make the belief true—which is to say, to what it in fact (according to the ascriber) represents.

It is tempting to formulate this difference in terms of a distinction between *narrow* and *wide* aspects of intentional content, the first of which depend only on the individual whose intentional states are in question, the latter of which depend also on the environment in which the individual with those states is embedded.[17] This is misleading if the narrow sense is identified with a methodologically solipsistic or individualistic point of view.[18] For the expressions that are used to specify the contents of both ascriptions *de dicto* and ascriptions *de re* are drawn from the same public vocabulary. There is no reason to think that any of this vocabulary can be semantically funded on the basis of the activities of a single individual. The content specifications of *de dicto* ascriptions, owing allegiance as they do to the acknowledgments of the one to whom they are attributed, do express (according to the ascriber) the agent's *subjective* perspective on intention and action. The content specifications of *de re* ascriptions, owing allegiance as they do to the facts bearing on actual outcomes or consequences of performances, do express (according to the ascriber) the *objective* aspect of intention and action. But these are complementary points of view that arise within the essentially social (in an *I-thou* sense) practice of deontic scorekeeping that alone is capable of conferring genuinely conceptual content.

What is expressed by *de dicto* ascriptions is conceptually prior to what is expressed by *de re* ascriptions *locally*, but not *globally*. That is, each true *de re* ascription is on this account true in virtue of some underlying true *de dicto* ascription: S believes of t that $\Phi(it)$ just in case for some t', $t = t'$ and S believes that $\Phi(t')$. It is also true (as will be seen in the next section) that *de dicto* ascriptions can be introduced into a language that previously had no way of making attributions of intentional states explicit, in advance of introducing on that basis also *de re* ascriptions; and the converse is not true. But what is expressed (where suitable explicitating vocabulary is available) by *de dicto* specifications of the content of intentional states cannot be made sense of apart from what is expressed (where suitable explicitating vocabulary is available) by *de re* specifications of the content of intentional states. Propositional contents (and so conceptual contents more generally) are essentially the sort of thing that can appear from both perspectives. Just as one cannot understand the content of a claim or judgment unless one has a practical mastery of the social distinction of attitude involved in assessments of it as *true* or *false*, so one cannot understand the content of an intention or an action unless one has a practical mastery of the social distinction of attitude involved in assessments of it as *successful* or *unsuccessful*.

3. Successful Action, Truth, and Representation

The success or failure of an action concerns the relation between the intentional *antecedents* of a performance and its actual *consequences*. It is assessed (as the truth of a belief is assessed) by comparing the commitments the assessor takes the agent to acknowledge with those the assessor acknowledges. Where suitable expressive resources exist, such assessments are made explicit by *de re* ascriptions. Understanding action accordingly requires that scorekeepers have binocular vision—that they be able to move back and forth between what is expressed by the *de dicto* ascriptions that figure in intentional explanations of one kind and the *de re* ascriptions that figure in intentional explanations of another kind. Without the capacity to produce and understand the first kind of intentional explanation, one cannot understand the difference between action and mere behavior: between shooting the gun and the gun's just going off. Without the capacity to produce and understand the second kind of intentional explanation, one cannot understand how shooting the cow can be something genuinely *done:* an action and not merely behavior, even though it was not done intentionally. (Turning on the light, which one intends, may be alerting the burglar, which one does not intend.) That is, one cannot understand the success and failure of intentions. To do that one must grasp the practical scorekeeping relation between specifications of the performance that are privileged by being the ones under which the agent is prepared to acknowledge or take responsibility for it, on the one hand, and the specifications (often in terms of consequences) under which it is available to anyone interpreting (keeping score on, attributing commitments to) the agent, on the other.

Explanations of the success or failure of an action—of why one's trying to do what one did had the consequences it did—appeal to what is true (according to the one doing the explaining), in addition to what (according to the one doing the explaining) the agent believed and so intended. To explain the success of an action, one must initially be able to ascribe beliefs and intentions *de dicto*—in order to say what the agent was *trying* to do, which sets the criteria of adequacy for *success. De re* recharacterization of those beliefs then puts them in a form in which their truth can be assessed. Typically the *success* of an action is to be explained by citing the *truth* of various of the agent's motivating beliefs. Nicole would have been successful if her belief that the animal in front of her was a deer had been true, that is if that belief had been a belief of a deer—rather than of a cow, that it was a deer. This is not the invariable form of an explanation of success, for I might succeed in my intention to put (that I put) oregano in the stew because of my false belief that the bottle you just carried to the counter contains oregano, in combination with my false belief that it is the bottle on the left, in case the oregano happens to be in the bottle on the left. Explanations of an agent's *reliable*

capacity to act successfully in some domain, however, cannot depend on such accidents and must accordingly appeal to true beliefs: Nicole believes of the animal in front of her that it is a deer, and it is true of the animal in front of her that it is a deer (that is, the animal in front of her is a deer). So the *success* of actions is to be assessed and explained by appealing to what is *true of* what the various intentional states leading to those actions represent or are *about.*

The analysis of what is expressed by *de re* specifications of the content of intentional states shows what this sort of representational contact with reality consists in—what is being *done,* in scorekeeping terms, when an assessor attributes such contact. Representational content is an essentially *social* affair—a consequence of the role claims play in interpersonal communication. *De re* ascriptions express the potential one individual's remark has as a premise from which *others* can draw conclusions. That is why they are the ones that figure in assessments of *truth,* which involve what the assessor endorses (not just attributes).

In diachronic communication and assessment of truth—for instance when explaining the (limited) practical successes of earlier theories and assessing the truth of their claims—the same structure is evident. Cognitive progress consists in saying more and more *true* things *about* the objects we talk about. Only such progress can explain increases in practical reliability—improvements in the long-term likelihood of success of actions based on those theories. The deep connection between our tautological *subjective* interest in the *success* of our practical endeavors (in getting what we want) and the *objective* *truth* and representational adequacy of our theoretical commitments (how we take things to be) is accordingly rooted in the social structure of hybrid deontic scorekeeping attitudes that is expressed in *de re* specifications of the contents of ascribed intentional states.

Though thoroughly social and linguistic, these representational contents are not *merely* linguistic—for they are not entirely up to us in the way in which what noises we use is entirely up to us. Assessing the success of action, for instance, involves, not only consideration of noninferential *exit* moves (from acknowledgments of commitments to performances) on the part of agents, but also noninferential *entrance* moves (from performances to acknowledgments of commitments) on the part of perceivers of those actions. The discursive practice that confers representational content is empirically and practically solid. It is a matter not just of hollow words and other linguistic expressions but also of mastery of their concrete circumstances of appropriate application and appropriate consequences of application, not all of which are purely inferential. Indeed, assessments of the success of actions often involve perceptual responses to intentional performances: Nicole and her scorekeepers can tell whether what she has shot is a deer by looking at it. This is one way in which agents themselves can come to be in a position to offer *de re* characterizations of their own previous beliefs, which can then

be used inferentially to alter the plans and intentions adopted in aid of some goal. Acknowledgments of doxastic commitments adopted noninferentially in response to performances produced noninferentially in response to acknowledgments of practical commitments can accordingly provide the feedback needed to complete the loop of a Test-Operate-Test-Exit cycle controlling goal-seeking behavior. The process by which commitments and concepts develop in a community over time through this sort of collaboration of the practical and empirical aspects of discursive practice is the expressing (in an inferentially articulated and hence conceptually explicit form) of the concrete constraints supplied by the fact that we think and act in an objective world.

The aim of the chapter so far has been to delineate the essential *social* structure that underlies the representational dimension of discursive practice, by offering an account of the hybrid deontic attitudes that are expressed by *de re* specifications of conceptual content—the primary representational idiom in natural languages. The objectivity of conceptual content—the way in which its proper applicability is determined by how things are in such a way that anybody and everybody might be wrong in taking such a content to apply in particular circumstances—is not by these means to be explained away. Rather the way in which the social and inferential dimensions of discursive practice combine in conferring, securing, and deploying such contents is to be investigated.

IV. FROM IMPLICIT ATTRIBUTION TO EXPLICIT ASCRIPTION

1. Prospectus

The fundamental claim made in this chapter is that the *representational* dimension of discourse is rooted in *social* differences of inferential perspective, which correspond to different repertoires of discursive commitments. Understanding that representational dimension is the key to understanding the objectivity of conceptual norms—how assessments of the objective correctness or incorrectness of the application of concepts (and so of the making of claims) can be explained in terms of social practices of deontic scorekeeping. As indicated in the opening chapters, it is a critical criterion of adequacy of any account of sapience or discursive practice that it be able to make intelligible this objectivity of the norms that govern the use of concepts. In the first three sections of this chapter, the crucial underlying representational dimension of the semantic contents conferred on states, attitudes, and performances by the role they play in discursive practice has been picked out as what is expressed by (weak) *de re* ascriptions of propositional attitude, and a deontic scorekeeping account of the use of such ascriptions has been sketched.[19]

This section examines more closely the explicitating expressive function

of ascriptions employing *de dicto* and *de re* specifications of propositional content. It focuses not, as heretofore, on *what* they make explicit but rather on *how* they do so. This account of ascriptional locutions takes the form of a structured recipe for introducing such locutions into an idiom that previously lacked their characteristic expressive power. The next section then moves from concern with epistemically *weak de re* ascriptions—which have not traditionally received much attention, though they are philosophically central according to the line of thought pursued here—to discuss epistemically *strong de re* ascriptions (and indeed, beliefs), which have standardly been taken to be of the greatest theoretical interest. With these accounts in hand, the final section returns to the issue of objectivity—of what it means for claims and the application of concepts to answer for their correctness to what is *true of* what they in fact represent or are about, rather than to what anyone or everyone *takes* to be true of whatever they *take* them to represent or be about.

The leading idea of the treatment of ascriptions presented here is that propositional-attitude-ascribing locutions are a species of *logical* vocabulary. Logical vocabulary has been distinguished in this work by its expressive role in making *explicit*, as something that can be *said*, some constitutive feature of discursive practice that, before the introduction of that vocabulary, remained *implicit* in what is *done*. In the paradigmatic case of introducing conditional locutions, where before one could treat inferences only as good or bad in practice, one comes by their means to be able to claim explicitly that one claim follows from another—to *say that* if *p* then *q*. Thus interlocutors can express the implicit inferential commitments in virtue of which (along with noninferential empirical and practical circumstances and consequences of application) sentences have the contents they do as the contents of explicit assertional commitments. Conditional, negating, identity, and quantificational locutions make explicit the inferential and substitutional commitments that articulate the *semantic contents* of states, attitudes, and performances. Ascriptional locutions, by contrast, make explicit *pragmatic attitudes*—paradigmatically attributions of commitment—which articulate the scorekeeping practices that confer those contents. In this sense they are pragmatic, rather than semantic, explicitating vocabulary—a different species within the genus of locutions qualifying as logical in virtue of their explicitating role.[20]

2. Introducing 'Says': Direct Quotation

The simplest form of ascription makes the attribution of an assertional speech act explicit in the form of a claim:

> Wittgenstein says, "Wovon man nicht sprechen kann, darüber muß man schweigen."

Since what is being considered is assertionally explicit ascription of spe-cifically *assertional* speech acts, 'says' here has the sense of 'claims', or 'says claimingly'—that is, it means *utters with assertional force,* rather than sim-ply *utters.* (So in this sense in saying a conditional, one does not say the antecedent.) How must an expression be used in order to play the ascriptional role of 'says' here in making explicit the attribution of assertional speech acts?

One of the implicit practical capacities interlocutors must be able to exercise in order to keep score on each other's commitments (and so to engage in discursive practice at all) is to treat each other as producing perfor-mances that have the pragmatic significance of claimings or assertions, which express explicit acknowledgment of doxastic commitments.[21] So part of being a competent discursive scorekeeper is being a reliable perceiver (under appropriate circumstances) of claimings, in that one is able differen-tially to respond noninferentially to the assertional performances of others by attributing to them those speech acts and the commitments they have thereby acknowledged. The response that is reliably elicited is then just the adoption of a practical scorekeeping attitude.

One essential expressive job of the assertion-ascriptional 'says' is to add an assertional response that is elicited by the same range of performances interlocutors are trained to respond to noninferentially by adopting such scorekeeping attitudes. 'Says' is used to form *reports* of assertional speech acts. Part of what it is for those reports to make explicit in the form of assertions what is implicit in attributions of doxastic commitments is that the assertional ascriptional commitments they express be *undertaken* or acknowledged noninferentially as responses to the same performances that elicit *attributions* of the commitments they ascribe. This expressive condi-tion provides these ascriptions with noninferential circumstances of applica-tion. Like any perceptual reports, however, once these claims are available in the language, they can be relied upon as *testimony.*[22] Thus someone who has never read Wittgenstein can come to be committed (and entitled) to the ascriptional claim made above by overhearing someone else making that claim. And that individual may have acquired entitlement to the claim either directly, as a response to the speech act itself (in this case, an inscriptional one), or by testimony. Thus it follows from the general account of assertional practice that the circumstances of application corresponding to these ascrip-tions include not only assertional speech acts of the sort ascribed but also type-similar ascriptional assertions. Since the latter class of testimonial cir-cumstances of application is common to all assertions as such, it is only the former class that is distinctive of specifically ascriptional assertions.

One essential *consequence* of application of these ascriptions has already been mentioned: asserting an ascription involves attributing a commitment as well as undertaking one. The distinctive consequences of application of these ascriptions also include their use in making explicit *deferrals* of the

responsibility to redeem one's *entitlement* to a claim. Recall that the basic model of assertional practice (introduced in Chapter 3) is structured by a constellation of complementary assertional *authority* and *responsibility*. The testimonial authority of assertional speech acts (their pragmatic significance as licensing others to undertake the same doxastic *commitments* they express) is balanced by a responsibility to exhibit, if appropriately challenged, one's *entitlement* to those commitments. This may be done in any of three ways: offering an inferential justification, exhibiting the commitment (and so the corresponding entitlement) as the result of exercising a reliable non-inferential perceptual ability, or invoking the testimony of an informant, to whom one defers the responsibility to redeem the entitlement in question.

Inferential justification is always a matter of making further explicit claims, which are thereby implicitly put forward as inferentially related to the claim in question, as premises to conclusion. The other two sorts of entitlement vindication are conceived of, to begin with, as nonassertional performances—of any kind whatever that might be invested with the appropriate practical scorekeeping significance. Thus in the most primitive cases, to invoke perceptual authority one might point to one's eyes or ears. To invoke testimonial authority one might point to one's informant, implicitly deferring responsibility the way someone in a theater line does by passing the ticket taker with a thumb jerked over the shoulder to indicate that the one next in line has tickets for both. But tropes may be introduced that permit these sorts of authority to be claimed explicitly. In the case of perception, these could range from Wittgenstein's suggestion "I speak English" to "I am a reliable reporter of red things under circumstances like these." Explicit deferrals are ascriptions.

In the case of the basic ascriptions formed using 'says', inferential justifications are minimal. As with any assertions, in languages with sufficient Boolean logical vocabulary, ascriptions follow from conjunctions in which they figure as conjuncts, and so on. Since as just indicated they can play the role of noninferential reports, they will share this mode of entitlement vindication with all other reportables. Their distinctive consequence of application, however, is that they can function to make deferrals explicit. That is, they are assertions that can have the pragmatic scorekeeping significance of deferrals. Instead of implicitly attributing an assertion to an informant by a wordless, (or at any rate assertionally inarticulate) deferral, one can *say* that so-and-so made the claim in question. Thus someone might attempt to vindicate entitlement to an assertive tokening of "Wovon man nicht sprechen kann, darüber muß man schweigen" by explicitly ascribing it to Wittgenstein. The assertion of the ascription offered as an example above then would function as a petition to inherit his authority for that view: anyone interested in the speaker's credentials for the commitment should investigate the reasons Wittgenstein has to offer for his. This role in making deferrals explicit is distinctive of ascriptions and is not inherited simply from their role as assertions, nor from their role as possible noninferential reports.

These ascriptions can accordingly be picked out by three fundamental features of the assertional commitments they express:

—the undertaking of those commitments involves attributing further assertional commitments acknowledged by others,

—their acknowledgment can be noninferentially elicited by speech acts by others that acknowledge further assertional commitments, and

—speech acts that acknowledge them can have the significance of deferring responsibility to vindicate entitlement to further assertional commitments to other interlocutors who have acknowledged commitment to those further assertional commitments.

To play this role, ascriptions must specify who these others are—the others to whom commitments are attributed, whose speech acts are being reported, to whom the ascriber can defer (should the ascriber not only attribute, but also undertake commitment to the claim ascribed). Even before locutions with the expressive power of ascriptions are available, interlocutors must already be able to keep score on the deontic statuses of different individuals. This implicit practical capacity to distinguish interlocutors is then exercised by the explicit use in ascriptions of singular terms that play the role that 'Wittgenstein' does in the example above. The use of singular terms (in 'referentially transparent' contexts such as this) has already been described.

More important, however, each of these distinctive features of the use of ascriptions of assertions appeals to a relation between an ascribing assertional commitment or its acknowledgment and the acknowledgment of some further assertional commitment, possibly on the part of another.[23] Ascriptional expressions must accordingly somehow specify those further assertional commitments. They do this by containing an embedded declarative sentence, one whose freestanding use would be the undertaking of an assertional commitment. The interesting differences between the various kinds of ascriptions concern the relation between this embedded tokening, on the lips or the pen of the ascriber, and various possible tokenings on the lips or the pen of the one to whom a commitment or its acknowledgment is ascribed. The basic ascriptions of assertional speech acts (overt expressions of acknowledgments of assertional commitments) whose expressive role has been considered so far are *direct* discourse or *quotational* ascriptions. This means that the relation between the ascribed sentence-tokening and the embedded sentence-tokening that specifies what is ascribed is one of lexical *cotypicality.* Acknowledgment of the commitment expressed by the ascription "Joanne says ⌜p⌝"[24] involves attributing to Joanne the assertional commitment she would be acknowledging by assertively uttering a tokening of that very same sentence p—that is, one of the same lexical-syntactic type as that appearing in the ascription. Similarly, acknowledgment of the commitment expressed by that ascription is to be appropriately elicited noninferentially by a tokening by Joanne that is of the same type. And that ascription

is in order as an explicit deferral (though not necessarily successful in re-deeming entitlement, for that depends on Joanne's entitlement) just in case Joanne has acknowledged an assertional commitment by uttering a sentence of that type. Any linguistic device for segregating the embedded sentence that has the significance of indicating this relation of cotypicality between reported and reporting tokening in ascriptions is playing the expressive role of quotation marks.

3. Introducing 'Says That': Indirect Discourse

Ascriptions accordingly can be thought of as having three parts: one specifying the *target* of the ascription by a singular term ('Wittgenstein', 'Joanne'); one specifying what sort of state, attitude, or performance is being ascribed ('says'); and one specifying its propositional content ("Wovon man nicht sprechen kann, darüber muß man schweigen," $\lceil p \rceil$). It has been shown how to understand in deontic scorekeeping terms the role played by each of these, in the case of the direct discourse or quotational ascriptions considered so far. The next step is to move from direct to indirect discourse, from *oratio recta to oratio obliqua.* Examples are such sentences as:

> Kant says that the essence of Enlightenment is responsibility.
> Hector says that he is hungry.
> John says that I am confused.

The circumstances and consequences of application of this sort of ascription are structurally much the same as for ascriptions of sayings, with one important exception. Direct discourse ascriptions can serve (as well as indirect discourse ascriptions) as inferential grounds capable of justifying indirect discourse ascriptions.

For example, the ascription

> Wittgenstein says, "Wovon man nicht sprechen kann, darüber muß man schweigen"

can be used to justify the ascription

> Wittgenstein says that what one cannot speak about one must pass over in silence.

Similarly, Wittgenstein's inscriptional tokening of

> Wovon man nicht sprechen kann, darüber muß man schweigen

could be responded to noninferentially by that same indirect ascription and could be the basis for a deferral achieved by asserting it. What is special about these indirect discourse ascriptions—and what must be understood if the asymmetric inferential relation between direct and indirect varieties of as-

cription is to be intelligible—is the use of the content-specifying clause, now marked in the regimentation by 'that' rather than quotation marks.

The fundamental considerations needed to understand these indirect discourse ascriptions have been laid out by Davidson, in his seminal essay "On Saying That," and have been clarified and developed further by McDowell in his commentary "Quotation and Saying That."[25] The approach they elaborate, which Davidson calls "the *paratactic* theory," conjoins three basic ideas. First, it focuses on tokenings, rather than on types.[26] Second, the speech act being ascribed is specified by *displaying* a sentence tokening related to the one being ascribed. The essence of the account accordingly consists in its understanding of the relation between the *reporting* tokening and the *reported* tokening. It is worth elaborating this point briefly.

Of course, direct discourse ascriptions also specify the ascribed saying by displaying a tokening, namely the one that is quoted, but all that matters about that tokening is its type. Indirect discourse loosens the constraint of cotypicality. Davidson emphasizes that the reported and reporting tokenings need not be in the same language, as in the specification in English above of what Wittgenstein said (or wrote) in German. McDowell acknowledges the significance of *translation* in indirect ascription, but he is concerned also to emphasize that *indexicals* behave very differently in 'that' clauses than they do when sealed inside quotation marks. Quite different remarks are reported by the ascriptions:

> John says, "I am confused"

and

> John says that I am confused.

The first has John characterizing himself, the one to whom the claim is ascribed, as confused, while the second has John characterizing the ascriber as confused.

They differ as to whether the indexical is to be evaluated according to the context of the reported tokening or according to the context of the reporting tokening. The demonstration and context relevant to understanding the indirect

> John says that *that man* is confused

is that of the ascriber, while the demonstration and context relevant to understanding the quotational

> John says, "That man is confused"

is that of the one to whom the speech act is ascribed. Although McDowell does not make the point, expressions that are not canonically indexical but whose recurrence classes are also not determined simply by cotypicality—

paradigmatically *anaphoric* proforms—also behave differently in the two sorts of ascriptions.

>Hector says, "He is hungry"

and

>Hector says that he is hungry

admit quite different readings. The interpretation of the first ascription turns on the anaphoric antecedents that are made available by the discursive context of the reported tokening, while the interpretation of the second ascription turns on the anaphoric antecedents made available by the discursive context of the reporting tokening.

The third idea underlying the paratactic approach is that the relation between the reported and the reporting tokenings is what Davidson calls "samesaying." Davidson is content to leave this as the rough-and-ready, more-or-less notion of sameness of content as what is preserved by good translations. This characterization of samesaying, however, is not of much help in thinking about the use of indexical and anaphoric tokenings in indirect ascriptions. A more useful way to put the condition is that the displayed tokening specifies the content of the ascribed uttering by being a samesaying with it, in the sense that its assertional utterance *would be* an undertaking of commitment to the very same claim that the ascriber undertook commitment to by the assertional utterance of the ascribed saying.[27]

In the context of the model of discursive practice developed here, an important feature of the expressive role of ascriptional contents is that their assertion can have the pragmatic significance of an assertionally explicit *deferral* of justificatory responsibility for a claim to the one to whom it ascribes a claim. The notion of sameness of claim appealed to by this formulation can be picked out in terms of the pragmatic notion of deferral potential. For a scorekeeper to treat a displayed sentence-tokening as standing in the appropriate samesaying relation to some other possible ascribed tokening is to take it that if the displayed tokening *were* uttered assertionally, deferral to the ascribed tokening would be in order.[28] Thus understanding what ascriptions make explicit offers a way of clarifying the notion of samesaying that governs the relation between reported and reporting tokenings in indirect ascription. For understanding the relation between reported and reporting tokening in terms of what claim *would* be made by the reporting tokening were it used assertively entails that the content of indexical expressions be evaluated according to the circumstances surrounding the ascriber's utterance, not the ascribed one.

This way of doing the work Davidson invokes samesaying for also deals with the difficulty that leads McDowell, in a later article,[29] to retreat from the samesaying formulation in favor of a notion of "corresponding senses"—a

relation stronger than coreference but weaker than identity of sense. The difficulty, discussed in more detail in the next section, arises from the fact that the special role that expressions such as 'I' and 'now' play in practical reasoning and the formation of intentions, particularly intentions in action, determines a sense in which no one else can have a thought with just the same content that John expresses by saying, "I am confused." "You are confused," in my mouth, at best expresses a corresponding thought, not the very same one. For present purposes it is enough to notice that understanding the relation between reported and reporting tokenings in terms of scorekeeping proprieties of deferral of justificatory responsibility permits these cases to be taken in stride; I can justify my assertion of "You are confused," addressed to John, by deferring to his assertion of "I am confused." The issue of whether the significance of the special features of the use of 'I' and 'now' is best addressed by individuating claim contents more finely accordingly need not be dealt with here.

Thus the paratactic approach to indirect discourse explains the expressive role of 'that' clauses (which specify the contents of ascribed speech acts and commitments) by focusing on the relation between a possible tokening in the ascriber's mouth and a possible tokening in the mouth of the target of the ascription. What primarily distinguishes this style of content specification from the directly quotational is that quotation marks seal off content-specifying expressions from the context of utterance of the ascriber. By contrast, in indirect discourse the content expressed by words appearing in the scope of the 'that' depends on the context of evaluation of the ascriber who actually utters them: the personal pronouns the ascriber uses, the demonstrations the ascriber makes, the anaphoric antecedents available to the ascriber, even the language the ascriber speaks. Davidson sets out his account as a response to a particular problematic; its task is explaining why substitution within 'that' clauses of sentences with the same truth-value does not preserve the truth of the ascriptions containing them. His official resolution of the problem is that the sentence tokening that occurs within the 'that' clause is not strictly part of the ascription at all. Properly understood, the ascription has the form "S said that." Another sentence tokening has been appended—without benefit of punctuation, which is why Davidson calls the theory "paratactic"—simply to display the referent of 'that', which is used demonstratively.[30] Thus substituting another sentence tokening for the one occurring in the scope of the 'that', whether or not the substituend has the same truth-value, alters the *referent* of 'that', and so should not be expected to preserve the truth-value of the whole ascription. McDowell is right to point out the affinities such an analysis has to an assimilation of indirect discourse to quotational discourse (in spite of the sharp distinction Davidson wants to make between his approach and quotational ones)—though the quotational model needs just the sort of commentary McDowell provides about differential semantic permeability to pragmatic context. For it is just because the

corresponding substitution alters the reference of quote-names that it is not expected to be truth preserving in direct discourse.

The lesson of the presentation of the fundamental insight of the paratactic analysis presented in the last few paragraphs, however, is that *all* the work here is being done by the emphasis on the relation between the *tokening* displayed in the ascription and a possible tokening by the one to whom a speech act or commitment is ascribed. So long as this point is insisted upon, it is otiose—a matter of being cute at the cost of being misleading—to treat the 'that' in indirect quotations as a demonstrative referring to (or an anaphoric dependent whose antecedent is) the sentence tokening that follows it. Although this is the most striking feature of Davidson's approach, it is just as well that this feature can be shown to be dispensable. For treating the role of 'that' as subordinating conjunction in ascriptions as derivative from its role as a demonstrative is both an etymological howler in English and depends on a coincidence that is not repeated in many other languages. It can simply be jettisoned, however, and Davidson's genuine insights put in the form of a direct description of the expressive role of the sentence tokenings displayed in the 'that' clauses of indirect discourse ascriptions.

The account offered here consists of two suggestions for filling in Davidson's analysis of saying-that. The first would expand his formula, along the following lines.

> Galileo said that the earth moves

is to be understood as having the sense of:

> Galileo said (something that *in his mouth then* committed him to what an assertional utterance of this *in my mouth now* would commit me to): The earth moves.

This way of putting it entails that the content of indexical (and anaphoric) expressions should be evaluated according to the circumstances surrounding the ascriber's utterance, not the ascribed one, which is what is distinctive of indirect quotation. The second suggestion is that the notion of identity of commitment undertaken that does duty here for Davidson's samesaying be understood in terms of the (ceteris paribus) scorekeeping propriety of *deference* to the first performance in vindicating entitlement to the second (merely possible) one. Each of these ways of developing Davidson's idea expresses an understanding of the paraphrase as explicating the rules by which the report*ing* tokening, occurring in the content-specifying clause, relates to the report*ed* tokening.

A common way of botanizing approaches to what is expressed by propositional-attitude-ascribing locutions distinguishes *sentential* and *propositional* approaches. Sentential approaches understand belief and similar attitudes as essentially involving relations to sentencelike entities,[31] where propositional approaches understand them as essentially involving relations

to more abstract objects. Sentential approaches have the advantage of relating the possession of propositionally contentful states to something intentional agents can unproblematically *do,* namely use sentences, but they are widely believed to have difficulty accommodating the sort of dependence of content on context exhibited by indexicals. Propositional approaches have the advantage that they can easily accommodate such contextual dependences, but they face a greater challenge in explaining how what believers *do* connects them with the abstract objects in virtue of which their beliefs are content-ful.[32] The present approach, by focusing on tokenings, displays some features in common with each of these approaches but avoids their difficulties. On the one hand, it shows how indexicality and anaphora can be accommodated in an approach to intentional states that is centered on relations to linguistic items. On the other hand, its semantics is rooted directly in its pragmatics: the notions that do duty for abstract propositional contents are explained in terms of linguistic scorekeeping practices.

Davidson's invocation of samesaying suggests that the relation between the tokenings that matter for reports of intentional states and the linguistic performances that express them is to be explicated by appeal to an antecedent notion of expressing the same content. The idea here is rather that the theoretical explanatory work usually done by appeal to notions of semantic content is done instead by the notion of constellations of tokenings, structured by the commitments (inferential, substitutional, and anaphoric) that link those tokenings. What it is for the significance of the occurrence of a tokening to be governed by a commitment is in turn explicated in terms of scorekeeping practices. For some purposes it is then acceptable to speak loosely about tokenings linked in such ways as "expressing the same semantic content." For other purposes, for instance where the asymmetric structure of anaphoric linkages is to the fore, assimilating these constellations of tokenings to equivalence classes by talk of each being related to one and the same content, the one it expresses, is not appropriate. Structures of tokenings governed by the same commitments bear a certain resemblance to classical notions of sentence and term, and they bear a certain resemblance to classical notions of the contents expressed by the use of sentences and terms, but the differences must be borne in mind. It is just these differences that distinguish the account offered here of propositional-attitude-ascribing locutions from both the purely sentential and the purely propositional approaches.[33]

4. Introducing 'Claims That'

The indirect discourse locutions considered so far are ascriptions that make propositionally explicit the attribution of speech acts that have the pragmatic significance of acknowledging propositionally contentful commitments. It is a small step to indirect discourse locutions that ascribe the

propositionally contentful commitments acknowledged by those speech acts. Thus a third form of regimented ascription can be introduced:

> Ruskin claims that medieval Gothic architecture is aesthetically and morally superior to Renaissance architecture.

'Claims' here should be understood as short for 'is committed to the claim'. Just as ascriptions formed using the regimented 'says' (claimingly) attribute a pragmatic force or significance that distinguishes sentences said from those merely uttered, so one need not actually say (claimingly) that p in order nonetheless in the regimented sense to claim that p. For one can adopt the attitude of acknowledging a commitment without having occasion overtly to avow it. (For this reason the status being ascribed here by the technical 'claims' would be expressed more naturally in English by something like 'holds'.)

The most straightforward version of such an ascription makes it appropriate just in case the individual to whom the commitment is ascribed is *disposed* to acknowledge it assertionally if suitably queried or prompted. The sentence tokening following the 'that' then purports to specify the content expressed by a tokening with assertional significance that the ascriptional target *would* produce under appropriate conditions, rather than the content expressed by an assertional tokening the target actually *has* produced. One way this expressive role shows up as a feature of the regimented use of 'claims' is in the propriety of inferences of the form

> S says that p,

therefore

> S claims that p,

while the converse form of inference is not in general valid (that is, commitment preserving). 'Claims' used in this way ascribes doxastic commitments that are, according to the ascriber, acknowledged by the one who is committed. Thus it corresponds to one way in which 'believes' is used (in accord with the overall explanatory strategy of having deontic status and attitude do much of the theoretical work generally done by appeal to the notions of intentional state and propositional attitude): namely, the use in which attribution of belief is tied most closely to dispositions to avow.

The difference between indirect discourse ascriptions regimented with 'says' and those regimented with 'claims' consists in a difference in how the reported tokening is to be understood as related to the target of the ascription. In the former case, what is attributed is an actual avowal of a doxastic commitment—the overt expression in an assertional speech act of the deontic attitude of acknowledging such a commitment. In the latter case, what is attributed is merely the disposition to such an avowal, and the tokening in question may be merely virtual—one that *would* be produced if the disposi-

tion *were* actualized. But the relation between these two tokenings—the reporting tokening in the scope of the 'that' clause uttered by the ascriber and the reported tokening whose assertional utterance, on the part of the one to whom the commitment is attributed, *would* be the avowal of the ascribed commitment—is not different in the case of 'claims' than of 'says'. That is, the expressive role of terms occurring in the scope of the 'that' clause is not different in these two cases.

The major difference in the circumstances and consequences of application of ascriptions of commitment that follows from relaxing the requirement from actual to merely dispositional avowal of the ascribed commitment concerns inferential relations among such ascriptions. Inferential relations among the propositional contents of claims are reflected in inferential relations among the ascriptions of commitments. Thus one may offer as justification for the ascriptional claim that

> S claims that $\Phi(t')$

the further ascriptional claims that

> S claims that $\Phi(t)$

and

> S claims that $t = t'$,

where these claims might be backed up in turn by ascriptions (regimented using 'says') of overt avowals of those commitments. To say this is not to say that interlocutors always acknowledge commitment to what follows from commitments they acknowledge, even where the inference involved appeals only to auxiliary premises to which they acknowledge commitment. They *ought* to acknowledge commitment to such consequences of their acknowledged commitments, and that means that in the absence of incompatible commitments, entitlement to

> S claims that $\Phi(t)$

and

> S claims that $t = t'$

entitles one as well to

> S claims that $\Phi(t')$.

These are *entitlement*-preserving inferences, not *commitment*-preserving inferences, and are defeasible as such. Thus the claim above, together with the claim that

> Ruskin claims that medieval Gothic architecture is the purest
> outward expression of admiration for the concrete, natural, vari-

> able, and changeful, and Renaissance architecture is the purest
> outward expression of admiration for the abstract, artificial, uni-
> form, and unchanging

give good reason for endorsing also the ascription

> Ruskin claims that the purest outward expression of admiration
> for the concrete, natural, variable, and changeful is aesthetically
> and morally superior to the purest outward expression of admi-
> ration for the abstract, artificial, uniform, and unchanging.

If, based on things Ruskin actually says, the ascriber were to find equally
good reasons to deny this latter ascription (without thereby also finding good
reason to withdraw the two previous ascriptions on which it is based), the
conclusion of the ascriptional inference should indeed be withdrawn—but
the fault would lie with Ruskin, rather than with the ascriber.

5. Introducing De Re Ascriptions

Three forms of regimented ascriptions of doxastic commitments
and the performances that express them have been introduced in this section:

1. Direct discourse ascriptions of assertional speech acts, of the form:

> S says ⌜p⌝,

as in

> Wittgenstein says "Wovon man nicht sprechen kann, darüber
> muß man schweigen."

2. Indirect discourse ascriptions of assertional speech acts, of the form:

> S says that p,

as in

> Kant says that the essence of Enlightenment is responsibility.

3. Indirect discourse ascriptions of acknowledged doxastic commit-
ments, which are avowed overtly by assertional speech acts, of the
form:

> S claims that p,

as in

> Ruskin claims that medieval Gothic architecture is aestheti-
> cally and morally superior to Renaissance architecture.

There are two distinctions in play in these sorts of ascriptions. There is
the distinction between attributions of a speech act and attributions of deon-
tic status or attitude expressed by such acts, which is marked in the regimen-

tation by the use of 'says' or 'claims'. There is also the difference between specifying the content of what is ascribed in the mode of direct or of indirect discourse, which is marked in the regimentation by the use of quotation marks, or of 'that' segregating a reporting sentence tokening. These are independent distinctions, so there is no reason not to recognize as well a fourth category:

4. Direct discourse ascriptions of acknowledged doxastic commitments, of the form:

    S claims $\ulcorner p \urcorner$,

as in

> Mill claims, "The internal sanction of duty is . . . a feeling in our own mind: a pain, more or less intense, attendant on violation of duty."

Claims of the form "S claims that p" express in propositionally explicit form the attribution of endorsement of what is expressed by a propositionally contentful *dictum*. They correspond to *de dicto* ascriptions of propositional attitude. The next step in reconstructing the way *implicit* scorekeeping attitudes of *attribution* of performances and statuses can be made *explicit* as *ascriptions* is to look at how *de re* content specifications are related to these *de dicto* ones. These can be formed from either ascriptions of speech acts or of commitments, yielding two further sorts of regimented ascription:

5. *De re* ascriptions of assertional speech acts, of the form:

    S says of t that $\Phi(it)$,

as in

> Johnson says of a man now otherwise forgotten that he was obscurely wise and coarsely kind.

6. *De re* ascriptions of doxastic commitment, of the form:

    S claims of t that $\Phi(it)$,

as in

> Gibbon claims of his most famous work that its English text is chaste, since all its licentious passages are left in the decent obscurity of a learned language.

This use of 'of' has already been discussed. Expressions occurring within its scope are those for whose use in specifying the content of the ascribed commitment or speech act the ascriber *undertakes,* rather than attributes, responsibility. For a scorekeeper to endorse S_2's *de re* ascription:

S_1 claims of t that $\Phi(it)$,

the scorekeeper must take it that there is some appropriate *de dicto* ascription

> S_1 claims that $\Phi(t')$,

and that it is proper to attribute to S_2 a commitment that could be ascribed by endorsing

> S_2 claims that $t = t'$.

The point is that in specifying the content of the doxastic commitment S_2 takes it that S_1 acknowledges, its inferential consequences have been extracted by conjoining it with substitutional commitments acknowledged by S_2, not necessarily S_1. The principle behind the regimentation is that one moves from the *de dicto*

> S claims that $\Phi(t)$

to the *de re*

> S claims of t' that $\Phi(it)$,

by adding the nonascriptional premise

> $t = t'$,

while one moves to the *de dicto*

> S claims that $\Phi(t')$

by adding instead the ascriptional premise

> S claims that $t = t'$.

Ascriptions that specify the content of the doxastic commitment (or the speech act that is the overt acknowledgment of such a commitment) in the *de re* way accordingly let the ascriber express explicitly the *attribution* of a doxastic commitment and also the *undertaking* of the substitutional commitments that license specifying its content by using the particular expressions that occur in *de re* position. The ascriber is not claiming that the one to whom the claim is ascribed would *acknowledge* commitment to the content as specified *de re*—only that the one to whom the claim is ascribed has in fact *undertaken* commitment to that claim in virtue of what that individual *would* acknowledge. One undertakes, not only those commitments one acknowledges and so is disposed to avow, but also their consequences. The expressions used to specify the contents of commitments ascribed *de dicto* determine (perhaps requiring suitable shifts of language, indexicals, and anaphoric chains) tokenings the one whose commitments are in question would (according to the ascriber) recognize as expressing commitments that individual acknowledges. The expressions used to specify the content of commitments ascribed *de re* determine tokenings the ascriber recognizes as expressing commitments that the one whose commitments are

in question would (according to the ascriber) acknowledge as expressed by some other (unspecified) tokenings. The difference is just whether the *expressive commitment*, to a particular form of words being a way of expressing the commitment in question, is *attributed* along with the doxastic commitment being ascribed or *undertaken* along with the ascriptional commitment itself. These expressive commitments can be thought of as hybrid deontic statuses compounded out of attributed doxastic (or practical) commitments and substitutional commitments, which can be either attributed (in the *de dicto* case) or undertaken (in the *de re* case). The appendix to this chapter shows how the regimentation presented here can be extended to include *iterated* ascriptions that mix *de dicto* and *de re* styles of content specification, as in

> This textbook claims that Russell claims of one of my favorite philosophers that he claimed of the planet Herschel discovered that it did not exist,

and shows that the ascriptional idiom that results is *expressively complete* with respect to possible constellations of scorekeeping attitudes toward doxastic and substitutional commitments.

It is worth noticing in passing that the hybrid deontic attitude made explicit by *de re* ascriptions has a dual, and that locutions have been introduced to make this attitude explicit as well. *De re* ascriptions *attribute* a doxastic commitment, while *undertaking* or taking responsibility for the expressive commitment involved in using a certain form of words to specify its content. One can also *undertake* a doxastic commitment, while *attributing* rather than undertaking the expressive commitment involved in using a certain form of words to specify its content. The locution used to make this attitude explicit is *scare quotes*. Consider this dialogue:

A: The freedom fighters succeeded in liberating the village.
B: Those ˢfreedom fightersˢ butchered two-thirds of its inhabitants.

The second interlocutor is undertaking responsibility for a claim but is disavowing responsibility for using the term 'freedom fighters' to express that claim. (Much the same effect could be achieved by prefixing 'so-called' to the offending expression.) In this context, the scare-quoted expression (regimented by superscripted 's's) functions *anaphorically*, picking up the previous interlocutor's use of the expression (just as the tokening /it/$_B$ is anaphorically dependent on /the village/$_A$). The point of scare quotes is to stipulate that one is talking about the same thing that someone else was (using an expression intersubstitutable with it, in the extended sense that includes expressions that are only repeatable as elements in anaphoric structures), while disavowing responsibility for the propriety of using the expression that individual used to pick out that common topic of conversation. The sense in which scare quotes are performing an expressive function comple-

mentary to that performed by *de re* locutions is evident in the fact that in these circumstances *A* should characterize *B*'s claim by an ascription such as

> *B* claims *of* the freedom fighters [that is, of the ones who according to me, *A*, are freedom fighters] that they butchered two-thirds of the village's inhabitants.[34]

It may also be remarked that there is no necessity to choose between conveying the information that is the specialty of *de dicto* specifications of ascriptional contents and conveying the information that is the specialty of *de re* specifications of ascriptional contents. *De dicto* specifications convey terms that the one to whom the doxastic commitment is attributed would (according to the ascriber) acknowledge as expressing that doxastic commitment. *De re* specifications convey terms that the ascriber takes responsibility for as expressing that doxastic commitment. Both sorts of information can be conveyed by using an ascription such as

> *A* claims *of* the ones who butchered two-thirds of the village's inhabitants, *as* freedom fighters, that they succeeded in liberating the village.

Here the effect of scare quotes is implicit in the significance of the 'as' clause: the ascriber is attributing, not undertaking, responsibility for the use of this expression. This 'of' . . . 'as' . . . 'that' . . . regimentation segregates the various words used to specify the content of the ascribed commitment, so as to mark clearly the expressive role being played by each—whose commitments determine the admissible substitutions that settle the significance of the occurrence of each of those expressions.

The regimentation can accordingly be of use in construing and disambiguating ascriptions in which a single expression is called on to perform both *de dicto* and *de re* expressive roles, where in "Alfred believes the man in the corner is a spy," 'the man in the corner' may "be doing double duty at the surface level—both characterizing Alfred's conception and picking out the relevant *res*."[35] That dual role can be indicated by regimenting this ascription as

> Alfred believes *of* the man in the corner *as* the man in the corner that he is a spy.

It is significant that *B* can characterize *A*'s claim, not only by an ascription such as

> *A* claims of the ones who butchered two-thirds of the village's inhabitants that they succeeded in liberating the village,

but also by one such as

> *A* claims of the ones that he calls 'freedom fighters' that they succeeded in liberating the village.

Here B achieves the same effect as scare quotes: securing a common topic without undertaking expressive responsibility for using the words 'freedom fighters' to characterize the ones who are nonetheless being talked about. In this case that expressive power is achieved by using the *explicitly* anaphoric indirect description 'the ones that he calls (refers to as) "freedom fighters"'. Recall that Chapter 5 described the use of these anaphorically indirect descriptions and argued that what makes an expression mean *refers* is its use in forming them. This sort of definite description makes explicit what is implicit in the use of scare quotes.

Descriptions can also be formed from occurrences of singular terms whose expressive role is complementary to that marked in the regimentation by occurrence within scare quotes, namely terms occurring within the scope of the 'of' in *de re* ascriptions.[36] Thus both A and B can use the *ascriptionally indirect description*

> the ones *of* whom A claims (believes) that they succeeded in liberating the village,

and both can endorse its intersubstitutability with

> the ones *of* whom B claims (believes) that they butchered two-thirds of the inhabitants of the village.

They can then use these terms to establish the common communicational ground necessary for a debate about the warrant for a commitment to intersubstituting either of these expressions with 'the freedom fighters'. Since the use of the terms that occur in *de re* position in ascriptions is governed by the ascriber's substitutional commitments in the same way that the use of terms having primary occurrence in the expression of nonascriptional claims is, the commitments involved in using ascriptionally indirect descriptions are just those outlined in Section II of Chapter 7, in the original discussion of definite descriptions. Putting this discussion together with that of *de re* ascriptions yields conditions under which

$$t = !x(S \text{ bel } \Phi(x))$$

that is, t is the one *of* whom S believes (claims) that Φ(it).

V. EPISTEMICALLY STRONG *DE RE* ATTITUDES: INDEXICALS, QUASI-INDEXICALS, AND PROPER NAMES

1. *From Weak* De Re *Ascriptions to Strong* De Re *Beliefs*

In one fundamental regard the account offered in this chapter is located firmly within the mainstream of contemporary philosophical thought. The last three or four decades have seen the development of a consensus that *de re* ascriptions of propositional attitude express something of fundamental importance for understanding the intentionality or repre-

sentational dimension of thought and talk—what it is for them to be *of*, *about*, or *directed at* objects and states of affairs that are not (in general) themselves in the same sense of, about, or directed at anything else. In other respects, however, the discussion so far is bound to appear simply to have missed the point of much of the discussion that has taken place within the scope of that consensus. Quine initiated the modern phase of interest in *de re* ascriptions by distinguishing, in terms of the substitutional behavior of singular terms, *notional* from *relational* senses of 'believes'—attributed by the use of *de dicto* and *de re* ascriptions respectively.[37] The distinction that has been reconstructed here concerns rather two ways of specifying the content of ascriptions employing one unitary sense of 'believes'.[38] One of the root intuitions that Quine begins with is that one does not put oneself in a position to have genuinely relational beliefs about an object simply by mastering the use of some singular term that in fact refers to that object. Someone who believes only that the shortest spy is a spy has not yet managed to have a belief that is *of* or *about* a particular object in the sense Quine is after with his notion of relational belief. Yet according to the reconstruction offered here, if Rosa Kleb is the shortest spy, then anyone who believes that the shortest spy is a spy thereby counts as believing *of* Rosa Kleb that *she* is a spy. Quine thinks of the latter ascription as appropriately used only to report someone who is in possession of information that might be of interest to the FBI—believing of someone in particular that that individual is a spy—a distinction that someone who believes only that the shortest spy is a spy surely does not achieve.

In short, the notion of *de re* belief that has exercised philosophers since Quine has more epistemic oomph to it than the thin notion reconstructed here. Indeed Kaplan, in "On Quantifying In"—which represents the first great milestone after Quine in the development of thought on this topic—formulates what is in its technical form essentially the reading of *de re* ascriptions offered here, only to discard it immediately as inadequate on the basis of the considerations Quine had raised about the shortest spy. The phenomenon of belief *de re* taken to be of interest is one that involves being en rapport with a particular object in a stronger sense than merely having some way or other of denoting it. Still, Kaplan does formulate what Dennett aptly calls this "denotational" sense of 'about' as a general strategy for moving from a notional sense of belief, conceived in terms of relations to linguistic items, to a relational sense of belief, conceived in terms of relations to the nonlinguistic items they denote. The idea is in essence that S believes$_{Rel}$ *of* t that $\Phi(it)$ just in case there is some expression α such that S believes$_{Not}$ $\ulcorner \Phi(\alpha) \urcorner$ and α denotes t.[39]

The present account derives from this by three developments. First, a more sophisticated account is offered of the underlying notional or *de dicto* belief ascription—an account that appeals to linguistic tokenings rather than types and is not directly quotational. Second, the notion of substitutional

commitments is used in place of an explicit denotation relation. Third, the whole account is set in a social-perspectival framework that explicates the shifts of perspective from one repertoire of commitments to another that are required to deal properly with the *iteration* of mixed *de dicto* and *de re* ascriptions. Having introduced the basic strategy that underlies a generically denotational reading of 'about', however, Kaplan rejects it, and the subsequent tradition has followed him in this regard.

Kaplan suggests that what is required to strengthen the flabby denotational sense of 'about' into a genuinely relational form is what he calls "a frankly inegalitarian attitude toward various ways of specifying" objects.[40] He explores the epistemically stronger varieties of rapport with an object involved not just in being able to use an expression that in fact refers to it but in knowing what object one is talking about. Not all terms that can be used to refer to an object are equally serviceable from the point of view of providing this sort of epistemic access. Even though the use of a particular term, say "the shortest spy," counts in some sense as picking out a particular object, knowing what one is talking or thinking about by using that term requires being able to pick out the object by the use of one's own resources in some narrower sense. In this connection Kaplan points to numerals and quote-names as privileged designators ("standard names") of numbers and linguistic expressions respectively.[41] Though the expressions

> the number of spheres that can touch any one sphere in a
> 24-dimensional Leech lattice

and

> the eleventh through thirteenth lines of the first of Rilke's
> *Duino Elegies*

pick out a number and a linguistic expression respectively, and so can be used in the denotational sense to talk and think *about* them, in a relatively clear sense one nevertheless knows what one is thereby talking and thinking about only if one understands them as intersubstitutable with the expressions

> 196,560

and

> und die findigen Tiere merken es schon / daß wir nicht sehr verläßlich zu Haus sind / in der gedeuteten Welt.

respectively.

Although Kaplan tries to develop a specific notion of an expression α being a "vivid name" of object t for believer S that accords with broadly Quinean intuitions about relational belief, the general lesson that should be drawn from his efforts is that different relational senses of belief, different sorts of epistemic rapport with the objects of belief, correspond to different restric-

tions on the expressions believers have available to them to pick those objects out. Just which privileged class of designators corresponds to "knowing who" or "knowing which" turns out to vary with particular situations in a way that defies antecedent systematization.

Recall the discussion (in Section I of Chapter 7) of the *Grundlagen* criterion that for an expression in the language to be properly understood as playing the role of a genuine singular term, and so as picking out a particular object, it must be understood as intersubstitutable with some other term. It was then argued (in Section II of that chapter) that to take the expression to pick out an object that exists in a particular sense (for example numerically, physically, or fictionally) is to take it that it is intersubstitutable with some term that is privileged as canonical with respect to that sort of existence. One can adapt the *Grundlagen* criterion applied there to the language as a whole and apply it to individual believers. The result is a minimal criterion of being able to use a singular term to put oneself cognitively en rapport with a particular object. To do so the believer must endorse *some* nontrivial identifying belief about it. That is, for some singular term $\ulcorner t \urcorner$ to figure *for a particular believer* as a name *of* an object in a sense stronger than the merely denotational, the believer must endorse *some* (true) nontrivial identity claim relating it as intersubstitutable with another term.

This minimal inegalitarian attitude would rule out ascriptions of strong *de re* belief regarding the shortest spy that are based only on the belief that the shortest spy is a spy. According to this criterion, however, if Holmes believes the murderer left the footprints, and so has *two* ways of picking out that individual, he will count as having beliefs that are *of* that individual in a stronger sense than the inspector, who can refer to him only as "the murderer"; yet such a situation surely falls short of Holmes knowing who the murderer is. The claim is that beyond this minimal criterion there are many stopping places, many sorts of demands that can be made on the other termini of nontrivial identities that have as one of their termini tokenings that the individual whose attitudes are being reported is willing to use. In an important article Sosa follows out the idea that matching intuitions about relational belief requires that the substituend denoting the object be a "distinguished term." His persuasive and influential conclusion is that in order to serve this function, the notion of distinguished term will have to be a "wholly pragmatic matter, which can change radically from one occasion to the next."[42] One stronger demand has been of particular interest: requiring that one of the terms the believer can use to pick out the object—one of the expressions occurring in the nontrivial identity claim that identifies it for the believer—be used (or be anaphorically traceable back to one that is used) in the demonstrative or indexical way.

In the fifties, then, Quine distinguishes the attribution of merely notional belief from the attribution of genuinely relational belief. In the sixties, Kaplan shows how various sorts of relational belief can be understood in

terms of notional belief, by adding to the thin denotational model of aboutness restrictions to privileged classes of expressions available for use in specifying the content of the underlying notional belief. Sosa then showed how radically the appropriate restrictions must vary with circumstances of belief attribution and the interests of the attributors. In the seventies Burge, Perry, Lewis, and Kaplan himself, among others, focus attention on the kind of beliefs that are *de re* or relational in the epistemically strong sense that results from insisting that the believer be in a position to pick out the object of belief by the use of *demonstratives* or, more generally, *indexical* expressions.[43] Since that time much effort has been invested in the notion that ˢdirectly referentialˢ expressions, paradigmatically indexical ones, make possible a fundamental sort of cognitive contact with the objects of thought, a kind of relational belief that is not conceptually mediated—in which objects are directly present to the mind, rather than being presented by the use of concepts.[44] According to this line of thought, what is of prime philosophical interest is this variety of epistemically strong *de re* beliefs. There are of course dissenters. Davidson, for instance, finds no use for anything but ordinary beliefs ascribed *de dicto*, and Dennett thinks (along lines somewhat similar to those pursued in the last two sections) that *de dicto* and (epistemically weak) *de re* ascriptions are essentially notational variants of one another, different styles of ascription of one kind of belief, and that insofar as the notion is intelligible at all, strong *de re* beliefs are of little significance.[45]

The view put forward here belongs to neither camp—neither to that of the boosters of nor to that of the scoffers at strong *de re* beliefs. It is distinguished from the scoffers in taking seriously the notion of epistemically strong *de re* beliefs. On the one hand, one should distinguish *de dicto* and *de re* styles of specifying the contents of ascribed propositional attitudes, as discussed thus far in this chapter. On the other hand, one should also recognize the existence of a distinct kind of epistemically strong *de re* beliefs—which deserve to be so called because of the way their distinctive contents derive from the sort of acquaintance with their objects that is expressed by the use of demonstratives and other indexicals. While sharing with them an acknowledgment of the existence of a distinct class of essentially indexical beliefs, the point of view adopted here nonetheless differs from that typical of the boosters of strong *de re* beliefs in its assessment of their theoretical significance. First, these beliefs do not form an autonomously intelligible sort or stratum of beliefs; one cannot coherently describe a situation in which this is the first or only kind of belief that is in play.[46] Second, although essentially indexical beliefs have a special sort of object-involving content that other beliefs do not, that object-involvingness should not be thought of as a *nonconceptual* element in their content; rather, the special sort of access to the objects their contents are about that the use of indexicals makes possible is a special kind of *conceptually* articulated access (though it is not correct to think of the role of the conceptual as *mediating* mind and its objects).[47] Finally, important

as essentially indexical beliefs are for our empirical knowledge and practical activity, what is of primary significance for understanding the representational dimension of thought and talk—its intentionality or aboutness in general—is the combination of doxastic perspectives expressed by (weak) *de re* specifications of the conceptual contents of beliefs of any sort, *not* the special sort of content possessed by strong *de re* beliefs, nor the special sort of rapport with objects they embody.

2. Essentially Indexical Beliefs: The Use of 'I'

The content of a belief or doxastic commitment may be termed "essentially indexical" if it cannot be expressed in the form of a claim without using (or appealing to the use of) some locution as an indexical. Following Castañeda, Perry has argued persuasively that there are things that can be said using expressions such as 'I', 'now', and 'here' that cannot be said by using nonindexical vocabulary.[48] I can of course refer to or describe myself in many ways, but no other term can do the expressive job done by 'I'. For any other term $\ulcorner t \urcorner$ I might use to pick myself out, there would always be some possible circumstances in which I could believe that t had a property—was about to be eaten by a bear, was standing in the Stanford library, was leaving a trail of sugar on the floor of the supermarket, and so on—without thereby believing that *I* have that property. This difference can be manifest in my actions, since a belief I would express using 'I' can be relevant to my practical reasoning and action in a way that no belief that can be otherwise expressed can match across the whole range of counterfactual situations. There will always be some of those situations in which I fail to realize that *I* am t. In those cases, I will fail to form intentions and to be motivated to act on the basis of beliefs about what t should do in the same way I would on the basis of beliefs about what *I* should do. Similarly, the belief that the meeting is starting *now* plays a different role in my behavioral economy than the belief that the meeting is starting at ten o'clock (or at any other time specified without the use of token-reflexives), since conditions can always be described under which I would fail to realize that the time so specified was *now*, as evidenced by my failure to form an intention-in-action.

In the idiom developed here, the key feature of the use of 'I' that is not reproduced by other coreferential expressions (even those that are *necessarily* coreferential, since one could always fail to realize that they were) is its use in *expressing the acknowledgment of a commitment*. What 'I' expresses is a potentially *motivating* acknowledgment of a commitment. The only expression that cannot counterfactually be separated from this motivational role is 'I'. Acknowledgment of a practical commitment is the deontic attitude that corresponds to forming an intention—what is expressed by a sentence of the form "I shall . . ."

One of the central sorts of practical capacity involved in rational agency

consists of reliable dispositions to respond noninferentially to acknow-
ledgments of some range of practical commitments by producing suitable
performances.[49] But what one is trained to respond to in this way is what is
expressed by sentences of the first-person form "I shall . . . ," not what is
expressed by sentences of the third-person form "*t* should . . ." The practical
reasoning that leads to such efficacious acknowledgments of commitments
to act also makes essential use of first-person pronouns:

> *I* am in danger of being eaten by a bear,
> so I *shall* run away.

The identity of the terms appearing in the premise and the conclusion is
essential to the commitment-preservingness of this form of practical substi-
tution inference.

> The bearded man (or N.N.) is in danger of being eaten by a bear,
> so I *shall* run away

is a good inference only insofar as it implicitly involves a commitment to
the auxiliary premise

> I am (=) the bearded man (or N.N.).

The fact that Perry points to—that no other expression can in general be
substituted for 'I' while preserving the role of the resulting sentences in
practical reasoning and rational agency—is a consequence of the possibility
of describing circumstances that ensure the agent will either lack the com-
mitment expressed by this identifying auxiliary hypothesis or have commit-
ments incompatible with it.

'I' finds its home language-game in acknowledgments of commitments to
act, and secondarily in the expression of beliefs and desires that are directly
relevant, as premises, to bits of practical reasoning that have formations of
intention as conclusions. The central defining uses of 'I' are not its uses in
such sentences as "I can run the mile in five minutes" but its uses in "I shall
open the door," as expressing the conclusion of practical deliberation, and
therefore as used in the expression of the premises. Its primary use is as
acknowledging a commitment. The lesson of Perry's examples is that one
cannot replace 'I' in its role in practical reasoning by any other designator,
while preserving its motivating role in deliberative practical reasoning, under
counterfactual circumstances. That is, although in particular circumstances
'I' may be replaceable by other expressions, that is always defeasible by
altering collateral beliefs and desires, in a way that is not paralleled for 'I'.
This is because what it expresses, the *undertaking* by *acknowledging* of a
commitment, is something that no one but I can do. Locutions that do not
function in practical reasoning in the way characteristic of first-person locu-
tions (such as names and descriptions) can be used to *attribute* commitments
to someone who may or may not (in different possible situations) turn out

to be, or to be known by the agent to be, that very agent. But *acknowledging* commitments is the basic way of *undertaking* them, and undertaking commitments cannot be reduced to *attributing* them, even to oneself. In the deontic scorekeeping social practices that institute commitments, these are two different kinds of practical doing. The phenomena that Perry displays in arguing for the ineliminability in favor of other thoughts of those expressed using 'I' is a manifestation of the irreducibility of the deontic attitude of acknowledging commitments to that of attributing them.

These same considerations can be seen to underlie the special features of the expressive role of 'I' that lead Anscombe, in her important meditation "The First Person,"[50] to deny that 'I' is a referring expression at all. From the point of view developed here, that dark claim appears to be an overstatement; the important thing is to specify how the use of 'I' differs from that of other referring expressions.[51] The line of thought that leads to the stronger conclusion could be put into the terms of this book something like this. If 'I' *were* a referring expression, it would refer to someone such that *attributing* a commitment (for instance, an intentional action or a belief) to that individual by its use would have the same practical significance as *undertaking* that commitment. One's *self* (= df. the object referred to by 'I') would just be the one, such that this sort of attributions to it *are* acknowledgments. Since one can show that no expression that could figure in attributions in this way could achieve the effect of acknowledgment, the conclusion is drawn that 'I' does not function as a referring expression at all. A less dramatic conclusion would be that it does not play a role analogous to that played by expressions that can be used merely to attribute commitments. The considerations Anscombe advances to distinguish the use of 'I' from that of other singular terms are somewhat different from, though related to, those Perry invokes.

The features she focuses on are:

1. One can have a special sort of nonobservational knowledge of one's own doings, of what one would express using 'I', and
2. The central sorts of claim one expresses using 'I' involve immunity from the possibility of misidentification.

These are closely related. Thus, the first point is that though I can be fooled in exceptional circumstances, in the usual run of things I do not need to *perceive* (visually, kinesthetically, or otherwise) my intentional bodily movements in order to know what they are. The second point is that when I acquire in this way a belief about what I am doing (about "actions, postures, movements, intentions," as Anscombe says),[52] there is no room (as their would always be were I merely observing) for a question as to *whose* it is I know about—no room for questions of the form "*Someone* is talking, sitting, perceiving a red triangle, trying to open the door, but is it *I*?" (Of course there are *also* cases in which questions of these kinds *are* in order, but they are precisely the ones where I have only observational knowledge of what I am

doing and how it is with me.) I cannot in these cases *misidentify* myself, accidentally pick out the wrong object, using 'I', as I can with *any* other way of referring to myself. (It should be clear that contact is made at this point with the considerations and counterfactuals that Perry points to.) It is easy to become mystified, thinking about what sort of peculiar object selves must be, and what sort of peculiarly intimate contact we must have with them, in order to be able to use first-person expressions that have these special properties. And it can be tempting to think that these features of 'I' show something deep about the difference between selves or subjects and about the objects that we can merely observe and refer to. But it is not obligatory to take this Cartesian, ontological road. The question is how best to understand the expressive role of 'I'.

Anscombe teaches us that the possibility of nonobservational knowledge expressed using 'I' that precludes the possibility of misidentification shows the inadequacy of an account of 'I' as the "word each of us uses to speak of [refer to] himself."[53] Anscombe exhibits the difference by considering a language (the 'A' language) in which an expression ('A') *is* used by each interlocutor just to refer to that very interlocutor. In the 'A' language, as opposed to our 'I' language, one must observe in order to know who ('A', or someone else) is doing something, and one is for that reason always liable to be mistaken about who really is doing it. One way to understand the particular sort of epistemically strong *de re* belief that can be expressed by using 'I' is to add practices to the 'A' language in stages, until 'A' has the fundamental features characteristic of 'I'. Anscombe describes the starting point: "Imagine a society in which everyone is labelled with two names. One appears on their backs and at the top of their chests, and these names, which their bearers cannot see, are various: 'B' to 'Z' let us say. The other, 'A', is stamped on the inside of their wrists, and is the same for everyone."[54] Individuals use chest-and-back names for each other and learn to respond to them as their own, but observations of one's own doings are reported by using the wrist-label 'A'. By hypothesis 'A' is used only on the basis of observation: tracing a bodily connection between the limbs whose motions are being reported and a wrist with an 'A' stamped on it. Thus misidentification is possible, for sometimes it will be someone else's token of 'A' that one glimpses. (It may be supposed that where one is in this way entitled to use 'A' in a report, one thereby becomes *inferentially* entitled to make the report that differs from the original in having the chest-and-back label one has learned to respond to as one's own substituted for 'A'.)

These labels can be used in keeping deontic score. So suppose that Anscombe's 'A' speakers engage in the sort of linguistic practices described in Chapter 3 and 4. They make claims, and they can make some of them noninferentially, as expressions of, or reports on, perceptions. These they can sometimes take-true by exercising a reliable differential responsive disposition, responding to a fact or state of affairs by a performance that has the

pragmatic scorekeeping significance in the community of acknowledging a doxastic commitment. A scorekeeper treats another interlocutor as a reliable observer (of a particular range of states of affairs under certain circumstances) by attributing entitlement to reports with those contents issued by that interlocutor under those circumstances. But at this stage there need be no way explicitly to *claim* reliability, for instance as a justification in response to a challenge by someone who is committed (noninferentially or not) to an incompatible claim.

One of the things they can observe each other doing is making observation reports. Suppose then that explicit ascriptions of claims are added to their vocabulary. Here it will be helpful to consider a special variety: ascriptions of reports of observations. These ascriptions will have inferential uses, but the uses of primary concern are noninferential. Thus T may report: "S claims-perceptually that p." A deontically hybrid ascribing locution could then be introduced so that if T also is disposed, noninferentially or otherwise, to endorse the claim that p, T could say, "S perceives (sees) that p." T could now respond to a challenge by another to S's observation report that p by saying, "S perceives that p." This response is also open to S, who in the basic case will phrase the same response as "A perceives that p" (for by hypothesis S's self-application of 'S' is always inferential, based on a noninferential disposition to use 'A').

At this stage, scorekeepers learn to treat the claim "A perceives that p" as in order *whenever* a perceptual claim that p is in order—commitment or entitlement to the 'A'-ascription of perception that p is attributed in just those cases in which commitment or entitlement to the underlying perceptual report that p is attributed. (Note that a scorekeeper can take a reporter to be *entitled* to the claim that p, and equivalently to the self-ascription of a perception that p, even if the scorekeeper does not endorse the claim that p—that is, does not take it to be true; attribution of entitlement depends on taking the reporter to be reliable, not to be infallible.) "A perceives that p" commitment-entails "Someone perceives that p" and is commitment-entailed by "Everyone perceives that p," in just the same way that "S perceives that p" does.

For the next stage, suppose that this functional equivalence of perceptual reports that p and self-ascriptions of perceptions of the form "A perceives that p" is extended so that "A perceives that p" becomes available as an alternate *noninferential* response to the same situation that would previously have prompted only a report to the effect that p. By hypothesis, the folk in question can already respond differentially to environing situations by producing the claim that p, so all that is being imagined is that they add as another possible response the claim "A perceives that p." Although clearly possible, this addition to the grammar of 'A' makes a substantial difference. For the first time, 'A' is given a use that does *not* depend on observation of a labeled wrist. In its initial use, "A perceives that p" still required tracing

in observation a connection between the mouth issuing the perceptual report that p and a wrist marked with an 'A'. It was accordingly liable to failures of identification on those occasions when reporters could see someone else's wrist. That possibility does not exist for the noninferential *responsive* use of 'A' in ascribing perception. Here 'A' is applied nonobservationally—in the sense that while one must observe that p in order to be entitled to the claim "*A* perceives that p," one need not *in addition* observe a wrist inscribed with an 'A'. Perceptual reports of the form "*S* perceives that p" always leave room for misidentification, but perceptual reports of p in the form "*A* perceives that p" do not. They do not simply because there is no question about whether the dispositions one is exercising or evincing are one's own—any more than when an environing situation noninferentially elicits the report that p there can be a question as to whether claiming that p is the performance one is disposed to produce in response to that situation.

What has been done for reliable differential dispositions to respond noninferentially to a situation by acknowledging a *doxastic* commitment, in *perception*, can also be done for reliable dispositions to respond noninferentially to the acknowledgment of a *practical* commitment by producing a performance, in *action*. To be agents, the 'A' creatures must learn to respond to the acknowledgment of some kinds of commitments by making-true some claims—just as to be perceivers, they must learn to respond to some facts by taking-true some claims, that is, by acknowledging some kinds of commitments. Interlocutors can observe each other's nonlinguistic intentional performances just as they can their perceptual reports. Suppose then that locutions are introduced making it possible to *ascribe* performances to individuals as intentional. These are *agentive* locutions, which may be represented for present purposes as having the form "*S* does q." For *S* to do q in this sense is for its being true that q to be brought about by the exercise of a reliable disposition to respond differentially to *S*'s acknowledgment of a practical commitment. As in the case of perception, these locutions are first used with chest-and-back names and then extended so as to give a use to "*A* does q," in the case where the one making that claim can observe the doing and trace the movements it involves to a wrist on which 'A' is inscribed. Thus at this stage the use of 'A' is still purely observational and liable to misidentification. At this stage, scorekeepers learn to treat *S*'s claim "*A* does q" as in order whenever the scorekeeper would endorse the ascription "*S* does q"—those are the circumstances in which *S* is, according to the scorekeeper, entitled to the commitment expressed by "*A* does q." "*A* does p" commitment-entails "Someone does p" and is commitment-entailed by "Everyone does p," in just the same way that "*S* does p" does.

For the next stage, then, suppose that limited scorekeeping equivalence of attributions of doings and of self-ascriptions of doings is extended in practice by interlocutors learning reliably to respond to practical commitments to do q, not only by doing q, but also by undertaking doxastic commitment to the

self-ascription "*A* does *q.*" Again, this is just adding to an already-established differential responsive disposition an alternate response that the creatures are already capable of producing in other circumstances. Here again, a substantial difference in the use of 'A' is being envisaged: one need not see a wrist label in order to use 'A' this way. The result is that although it is possible for *S* noninferentially to claim "*A* does *q*" without in fact doing *q*—for instance when the agentive ascription in response to acknowledgment of the practical commitment is not prevented, but doing *q* as a response is prevented—in the cases where reliability prevails and both responses are elicited by that acknowledgment, *S* will have *nonobservational* knowledge of those doings and will express such knowledge by the 'A' ascription. Furthermore, uses of "*A* does *q*" that are elicited in this way are immune from errors of misidentification, as those based on observation of a labeled wrist were not.

At this point the initially impoverished and unselfconscious speakers of the 'A' language have a rudimentary but recognizable first-person pronoun. They have, in favored cases, immediate, nonobservational knowledge of their own doings and perceivings. They are not incorrigible about them, for they cannot guarantee the truth of what they take themselves to perceive or the success of what they take themselves to do, but when they are right, they *know* they are right: scorekeepers within the community will attribute to them a hybrid deontic status that corresponds to justified true belief. The claims about their perceivings and doings that, if true, express this sort of immediate, nonobservational knowledge are immune from errors of misidentification. Challenges to *S* of the form "Were you right to say '*A* perceives . . .' or '*A* does . . .'? Perhaps it was really *T*?" will not be in order, except as queries about whether the warrant for the claim really was observational— that is, elicited by the inspection of wrist labels rather than by the exercise of nonobservational responsive dispositions. The *immediacy* of these acknowledgment-expressing uses of 'A'—as either noninferentially elicited concomitantly with the acknowledgment of doxastic commitments ('perceives') or noninferential responsive consequences of acknowledgments of practical commitments ('does') elicited concomitantly with actions—ensures that the question of whether one should not look more closely and see whether it is someone else's commitments that are at issue cannot arise.

Finally, this fable promises to make intelligible not only Anscombe's but Perry's observations about the use of 'I', though only the briefest sketch can be offered here of how this story would go. What one expresses using 'I' plays a special motivational role. 'Does' as described here in effect attributes the intentions-in-action into which prior intentions mature. It would be possible to extend this discussion by introducing next ascriptions of *prior intentions*, which would lead to self-ascriptions of the form "*A* shall do *q.*" When acknowledgments of these ascriptions come to be among the noninferential responses to acknowledgments of practical commitments regarding future doings, they express nonobservational knowledge of one's intentions, rather

than just of one's actions. (Discussion of the motivational role of 'now' that Perry points to would require going further into the expressive role it plays in connection with these ascriptional locutions in making explicit the practical reasoning implicit in the maturation of prior intentions into intentions-in-action.) The final stage is reached when training secures it that when one says 'I shall' or uses 'I' in substitution inferences that lead to that conclusion, then one is disposed to respond to (the acknowledgments that are made explicit in) such concomitant expressions by producing a performance. There will be no other expressions, S or T, which de jure have this significance. For each of them, it is only if something that means "I am (=) S" is endorsed that the motivational significance is guaranteed to be in play, just as in Perry's examples. The fable shows just how "A perceives a bear" or "A is making a mess" can have a significance in practical reasoning across various possible situations that is not matched by any of the chest-and-back labels, or indeed any other designators that do not play the expressive role with respect to perception and action that has been built into the use of 'A' by the end of the story.[55]

3. Anaphoric Quasi-Indexicals and the Ascription of Essentially Indexical Beliefs

In this way, then, it is possible to understand the basic outlines of the expressive role played by the first-person pronoun 'I'. There is no necessity that every idiom include an expression playing that role; there need be nothing incoherent in descriptions of communities of judging and perceiving agents, attributing and undertaking propositionally contentful commitments, giving and asking for reasons, who do not yet have available the expressive resources 'I' provides. Yet when this bit of vocabulary is available, it makes explicit something important about such a set of discursive practices. 'I' is a logical locution. While those who lack it can be *conscious* in the sense of *sapient*—can be explicitly aware of things by making judgments about them—they are in an important sense not yet *self*-conscious. One of the normative social statuses instituted by any scorekeeping practices that qualify as discursive is that of being an individual *self*: a subject of perception and action, one who both can *be* committed and can *take* others to be committed, a deontic scorekeeper on whom score is kept. Selves correspond to coresponsibility classes or bundles of deontic states and attitudes—an indispensable individuating aspect of the structure of scorekeeping practices that institutes and articulates discursive commitments.

The notion of *one* performer who is responsible for *two* different claims is implicit in the practical acknowledgment of relations of inferential consequence among claims. One interlocutor is not responsible for the inferential consequences of commitments undertaken by another; such consequence relations govern only the commitments of a single interlocutor. The notion

of incompatible propositional contents similarly presupposes the assignment of responsibility for commitment to those contents to reidentifiable interlocutors. There is nothing wrong with one interlocutor being entitled to commitment to a content incompatible with a content to which another is committed; what is forbidden is that one and the same interlocutor should count as entitled to a commitment incompatible with another claim to which that same interlocutor is committed. Asserting, challenging, deferring, justifying—all these significances that performances can have according to basic discursive scorekeeping practices depend on sorting commitments and performances into the concomitancy classes corresponding to individuals.

That two commitments are to be assigned to the same individual is accordingly a fundamental social status instituted by those scorekeeping practices. It is not arbitrary how such commitments are sorted into equivalence classes, for the reliability of the responsive dispositions that make possible perception and action, and hence the empirical and practical aspects of conceptual contents, are tied to the careers of ongoing organisms. The sorting depends on the bodily relations between the retina that is irradiated and the mouth that can be trained to make reliable reports as responses to what was seen, between what that mouth acknowledges and the limbs that can be trained reliably to move in response to that acknowledgment, and so on. The empirical and practical abilities presupposed by our conceptual contents essentially involve locating beliefs—connecting for instance the public coordinate addresses appealed to by commitments to physical existence and the egocentric space induced by the use of demonstratives[56]—that depend in turn on the fact that our bodies trace out continuous paths in space.

'I' plays the substitutional role of a singular term; it *is* a pronoun and so should be counted, when correctly used, as giving us access to an object. Selves are precisely the kind of thing it gives us access to. But so would chest-and-back names, provided they played the proper role in scorekeeping practices. The difference lies in the sort of privileged access that 'I' gives us to those objects. The sort of self-consciousness that the use of first-person pronouns makes possible is not available through the use of any other sort of expression. Thus what can be expressed only by the use of 'I' is an example of a special kind of conceptual content—one that is *essentially indexical.* Beliefs with this sort of content embody a special sort of acquaintance with their objects; they are *de re* in an epistemically strong sense. Doxastic and practical commitments that are explicitly expressed (in the sense of overtly acknowledged by the production of propositionally contentful speech acts) by the use of other indexical expressions such as 'now' and 'here', including demonstrative ones such as 'this' and 'that', can also be strongly *de re* in this sense. The sort of cognitive access to their objects that they incorporate is expressible *only* by the use of indexicals; the role they play in the conceptual economy of the believer cannot be duplicated by beliefs expressible entirely in nonindexical terms.

Beliefs that are about their objects in the epistemically strong sense characteristic of this kind of *de re* beliefs are essentially indexical; being able to use an indexical expression to pick out a particular object depends on being in the right circumstances or context (in the case of 'I', it depends on *being the right object*). Thus one must be in the right circumstances or context in order to have strong *de re* beliefs (or intentions) about an object. Scorekeepers cannot pick each other out by using 'I', and it is only rarely that in assessing and attributing the commitments of another a scorekeeper will be in a position to pick out demonstratively all the objects with which the other was demonstratively acquainted. What constraints do these facts impose on the *communication* of the strong *de re* beliefs expressed by the use of indexicals? Does the fact that demonstrative references are only *accidentally* shareable across interlocutors, and that first-person references *never* are so shareable, mean that essentially demonstrative thoughts are only accidentally shareable, and first-person thoughts never are? Can the audience keeping score on another's commitments understand the essentially indexical commitments that the interlocutor undertakes, in cases where the audience is not (whether de facto or de jure) in a position to pick out the same objects indexically? Are strong *de re* beliefs strictly or largely incommunicable?

Concluding that they are is adopting what Perry calls the "doctrine of limited accessibility." He takes the collateral costs of such a theoretical commitment to be sufficiently high to motivate resisting this doctrine and suggests a strategy for doing so. His way out is to distinguish two components of belief: a belief *state* exhibited by the believer and a belief *content* exhibited by that state. These correspond to Kaplan's notions of *character* and *content*. Two different people who each believe something they could express by asserting the sentence "I am threatened by a bear" share a belief state, though the contents of their beliefs (who it is they take to be so threatened) differ. Michele, who believes "I am threatened by a bear," and Nicole, who believes "Michele is threatened by a bear," are in different belief states, which share a content. Thus belief states are shareable, and belief contents are shareable; what is not shareable is their conjunction: only Michele can have a belief with that particular content by being in the belief state she could express by using the sentence "*I* am threatened by a bear." The question then is what constraints the fact that sharing this conjunction of state and content may be impossible places on the theoretical understanding of *communication* of the state one is in and the content that it has. Clearly in some cases one understands the commitments of another only if one understands both the belief *state* the other is in and the *content* of that belief. But both of these are shareable. What needs to be worked out is an account of the sense in which being able to share these individually suffices to make intelligible (hence communicable) their conjunction, which is not shareable.

McDowell is led to a different attitude by these same considerations

concerning essentially indexical beliefs: he endorses the doctrine of limited accessibility toward strong *de re* beliefs but denies that doing so incurs the cost of making it unintelligible how they can be understood by and communicated to others. No one else can have the thought that Michele would express by saying, "I am threatened by a bear," but that does not mean that no-one else can understand what thought she would be expressing by that claim. He says: "Frege's troubles about 'I' cannot be blamed simply on the idea of special and primitive senses; they result, rather, from the assumption—which is what denies the special and primitive senses any role in communication—that communication must involve a *sharing* of thoughts between communicator and audience. That assumption is quite natural, and Frege seems to take it for granted. But there is no obvious reason why he could not have held, instead, that in linguistic interchange of the appropriate kind, mutual understanding—which is what successful communication achieves—requires not shared thoughts but different thoughts which, however, stand and are mutually known to stand in a suitable relation of correspondence."[57] So McDowell proposes that different *de re* senses can be sorted in two ways: one corresponding to Kaplan's characters (according to which you and I express senses of the same sort by saying 'I'), and one corresponding to my use of 'I' and your use of 'you'. Successful communication requires only that the audience be able to associate an appropriately corresponding *de re* sense with the speaker's utterance. It is on this basis that he retreats from his earlier endorsement of *samesaying* as the relation between reported and reporting tokenings in (*de dicto*) ascriptions of such beliefs and their expressions. That the senses expressed by 'I' in my mouth and 'you' in yours, or by 'today' uttered today and 'yesterday' uttered tomorrow, correspond in the right way is conceived as a requirement stronger than mere coreference but weaker than identity of sense or content.

Clearly the facts about essentially indexical or epistemically strong *de re* beliefs can be rendered in either way. Employing the machinery of the previous section makes it possible to describe more precisely the phenomenon that each of these approaches renders in general terms. For the sort of understanding that is the cognitive or semantic uptake that must be appealed to in defining the notion of successful communication can be expressed in scorekeeping terms by the requirement that the scorekeeping audience interpreting the performances of the speaker or agent be able to *attribute* the very same commitments that the performer *undertakes* or *acknowledges* by those performances. If one wishes to construe such understanding on the model of shared contents, then it is this pragmatic scorekeeping condition that determines what counts as grasping the *same* conceptual content; if one wishes to construe it instead on the model of grasping (in the cases at issue strongly *de re*) senses of the same *sort* ("corresponding senses"), then it is this pragmatic scorekeeping condition that determines what counts as those contents being of the *same* sort. So from the point of view of deontic scorekeeping,

the phenomenon to focus on is that of *attributing* essentially indexical commitments. This is an implicit practical capacity; as always, the best route to theoretical understanding of something *implicit* in a practice is by consideration of the use of the expressions by which it is made *explicit.* Since attributions are made propositionally explicit in the form of *ascriptions*, this methodological precept dictates a look at *ascriptions* of essentially indexical commitments.

Transposed into these terms, the problem becomes that of understanding how an ascriber can specify the content of an ascribed commitment, even though no tokening in the mouth of the ascriber could have just the same significance as a tokening of, say, "I am threatened by a bear," in the mouth of the one to whom the commitment it expresses is to be ascribed. So put, the problem has a straightforward solution. Castañeda is concerned that the ordinary use of a *de dicto* ascription such as "Michele believes that she is threatened by a bear" does not distinguish between its use to report the belief Michele could express as "I am threatened by a bear" and other beliefs she might express by making such claims as "Michele is threatened by a bear" or "The woman under the pine tree is threatened by a bear."[58] He suggests that the genuinely first-person belief might be reported by some such idiom as "Michele believes that *she herself* is threatened by a bear." Anscombe points out that Greek, among other languages, contains special *indirect reflexive* constructions, which are used precisely to ascribe the use of reflexives to others.[59] Castañeda suggests regimenting the ascription of essentially indexical propositional attitudes by the use of what he calls "quasi-indicators," which are stipulated to perform this expressive function. In his notation, what is expressed in the informal example above by 'she herself' is expressed by the use of the quasi-indicator 'she*', so that the ascription "Michele believes that *she** is threatened by a bear" is correctly used only to ascribe the first-person belief Michele could express by using "*I* am threatened by a bear."

An analogous difficulty arises for the reporting of other essentially indexical beliefs, notably those the believer would express using *demonstratives.* As was pointed out in the discussion of the expressive role of 'that' clauses in specifying the contents of *de dicto* ascriptions, one of the characteristics that distinguishes indirect discourse from direct discourse is that indexical and demonstrative expressions appearing in them are used by the ascriber, and so evaluated according to the context of the ascriber's speech act, not that of the speech act ascribed. Thus ascriptions such as "Danielle believes that *this* is an interesting painting" do not settle whether the belief ascribed is one Danielle is in a position to express by claiming "*This* is an interesting painting," rather than by using some nondemonstrative singular term, perhaps by claiming "Edward Hopper's *Excursion into Philosophy* is an interesting painting"—a belief one could have without knowing *which* painting one was talking about, in the strong sense of 'knowing which' that goes with

demonstrative acquaintance. Castañeda's quasi-indicators can mark this distinction as well: "Danielle believes that *this** is an interesting painting" ascribes the use of a demonstrative without using one.[60]

How do quasi-indicators perform this expressive function? Castañeda notes that tokenings functioning as quasi-indicators are not used demonstratively and that they have antecedents, but he does *not* claim that they are anaphoric dependents. The only reason he gives is a bad one: they are not replaceable by their antecedents. But it was pointed out already in Chapter 5 that anaphoric dependents are not in general replaceable by their antecedents; syntactically lazy anaphora is only one species. The analysis in Chapter 7 of asymmetric repetition structures of tokenings in terms of inheritance of governing substitutional commitments showed how anaphora can be understood more generally. In this context, the indirect indexical constructions Castañeda regiments as "quasi-indicators" can be understood to be functioning in a straightforwardly anaphoric fashion. They differ from ordinary anaphoric dependents in the extra information they carry about their antecedents—namely, the antecedent of a tokening of *this** is a demonstrative of the sort the ascriber would express using 'this'.[61]

There is no special difficulty in understanding how anaphoric proforms could convey such added information; pronouns that do so are more common in natural language than purely anaphoric forms. Thus the use of 'he' carries the added information that its antecedent picks out an animate organism, and a male one. In many languages, of course, the gender information conveyed by a pronoun concerns the antecedent itself, rather than what that antecedent refers to. Anaphorically indirect indexical expressions are to be understood as working the way these pronouns do, except that the additional information they convey concerns not gender and number but the status of the antecedent as an indexical. In this way they can be compared to anaphorically indirect definite descriptions such as "the one he referred to as 'that airhead'," which also convey information about the expression on which they are anaphorically dependent.

To add the expressive power of these expressions to the regimented ascriptions of the previous section, let $\lceil a^{\langle t \rangle} \rceil$ stand for an expression that is anaphorically dependent on a tokening of type $\langle t \rangle$, such as 'he$^{\langle I \rangle}$' (which expresses what Castañeda's "he himself" does).[62] This usage differs from Castañeda's in that where he simply indicates that there *is* an indexical or demonstrative antecedent by using a '*', here the type of that antecedent is explicitly specified in the superscript. 'Quasi-indicator' is Castañeda's technical term, and he does not use this convention or understand the functioning of the expressions to which it applies in anaphoric terms. Thus it seems well to adopt a different term. They might be called 'anaphorically indirect indexicals', but this would suggest that they are a kind of indexical, whereas they are really a kind of anaphoric dependent. So the term 'quasi-indexicals' will be used here, as a reminder of the origin of this category in Castañeda's quasi-indicators, with the difference being the anaphoric analysis of their expressive function.

The addition of these expressive resources makes it possible to ascribe essentially indexical beliefs as such. "Danielle believes that $it^{\langle this \rangle}$ is an interesting painting" is a *de dicto* ascription of such an epistemically strong *de re* belief. Since what a strong *de re* belief is a belief *of* or *about* is essential to its having the content that it does (see the discussion below of the "object-dependence" of this sort of belief), a full specification of its content requires also a *de re* specification.

> Danielle believes that $it^{\langle this \rangle}$ is an interesting painting, and the painting she believes that of or about is Edward Hopper's *Excursion into Philosophy*.

This is a *de dicto* ascription conjoined with an identity one side of which is an ascriptionally indirect *de re* definite description. It could be expressed in a single mixed *de dicto/de re* ascription by something like:

> Danielle believes of Edward Hopper's *Excursion into Philosophy* that it = $it^{\langle this \rangle}$ is an interesting painting.

The basic effect of this mixed ascription can be achieved without the use of quasi-indexicals by employing the full "... of ... as ... that ..." structure mentioned at the end of the previous section. Thus the ascription

> Danielle believes *of* Edward Hopper's *Excursion into Philosophy*, <u>as</u> *'this'*, <u>that</u> it is an interesting painting[63]

specifies the content of the ascribed belief in full. Understanding the anaphoric function of both quasi-indexicals and indirect definite descriptions formed using 'refers' makes it possible to appreciate that this is because that ascription has the sense of

> Danielle believes *of* Edward Hopper's *Excursion into Philosophy*, which she refers to by using a tokening of type $\langle this \rangle$, *that* it is an interesting painting.

What have been considered so far are *de dicto* ascriptions of strong *de re* beliefs. The use of quasi-indexicals also permits a straightforward analysis of *de re* ascriptions of beliefs that are strongly *de re* in the sense of being essentially indexical or demonstrative.

> S believes of$_{strong}$ t that $\Phi(it)$

is an ascription expressing the perspectivally hybrid deontic attitude that consists of the *attribution* that would be expressed by the *de dicto* ascription

> S believes that $\Phi(that^{\langle this \rangle})$,

and the *undertaking* of the substitutional commitment that would be expressed by the identity

> $that^{\langle this \rangle} = t$.[64]

This understanding just adapts the reading of weak *de re* ascriptions by combining it with a restriction on the terms the ascriber takes it the believer could use to express the belief. Such an account combines:

—The denotational reduction of weak *de re* to *de dicto* ascription;

—Kaplan's and Sosa's idea that stronger sorts of rapport with objects correspond to a "frankly inegalitarian attitude toward different ways of specifying objects";

—Perry's point that there are essentially indexical or demonstrative beliefs, and Burge's suggestion that it is just these that are strongly *de re*;

—Davidson's and McDowell's paratactic account of *de dicto* ascriptions in terms of relations between the report*ing* tokening, in the ascriber's mouth, and the (possible) report*ed* tokening, in the mouth of another, which is or expresses what is ascribed.

The final theoretical ingredient, which serves to combine all the rest, is then

—An anaphorically rendered notion of quasi-indexicals, whose expressive role in ascriptions is modeled on Castañeda's quasi-indicators and indirect reflexives.

With these expressive resources it is possible even for an interlocutor who cannot have that 'I' thought and that demonstrative thought to ascribe *exactly* the beliefs that Michele would express using the first-person pronoun "I am threatened by a bear"[65] and that Danielle would express using a demonstrative "This is an interesting painting." What is required is only that the reporting tokening occurring in the ascription determine in a systematic fashion what token is being reported—either the utterance being ascribed or the tokening that would express the state being ascribed. For that it is not necessary that the content-specifying token in the ascription be one that *would*, if uttered by the ascriber as a freestanding assertion, undertake the very same commitment as is being ascribed. That samesaying relation between the reported and the reporting tokening holds in the case of *de dicto* ascriptions of ordinary beliefs, but not of essentially indexical ones. This is the point McDowell makes in the passage cited above. The discussion of quasi-indexicals has shown how to make precise his notion of suitably related corresponding thoughts or *de re* senses. The relation in question is stronger than mere coreference in just the way that anaphoric dependency is stronger than coreference: a tokening used quasi-indexically and the token used indexically that is its anaphoric antecedent are members of one single token-repeatability structure.

The quasi-indexicals used in the *ascription* of essentially indexical beliefs are *logical* vocabulary. They make explicit the implicit practical understanding of another's indexical beliefs that is required for their *attribution*, and therefore for their *communication*. They show what is required to attrib-

ute *exactly* the belief that another undertakes, even in the cases where the essentially indexical character of the belief makes it impossible, in principle or in practice, for speaker and audience, believer and scorekeeper, to *share* that belief. In the least demanding extensional or representational sense, communication requires securing only coreference. Further grasp of what is expressed by the terms the speaker did (or the believer would) use is not needed. In this sense it is sufficient for an interpreter to understand another's remark, and so for communication to count as successful, if that scorekeeper adopts a deontic attitude of attributing a commitment that would be undertaken explicitly by asserting a (weak) *de re* ascription. Understanding in this sense is gathering the information conveyed by an utterance—knowing what is being talked about and how it is being classified.

The use of quasi-indexicals in ascriptions of strong *de re* beliefs preserves more than reference, however. The ascriber's access to the object the belief is strongly about is anchored, by anaphoric links, in the believer's demonstrative acquaintance with the object. The sort of attribution that is in this way made explicit is attribution not just of a belief *about* that object but of belief *strongly* about that object. This is a more demanding, broadly intensional sort of communication and understanding of another's claim. It is an extension—via anaphoric commitments to defer regarding what substitutional commitments govern the use of an expression—of the sort of sameness-of-claim content defined by deference of assertional responsibility to vindicate entitlement to a claim. A member of the audience of a speech act containing a demonstrative reference can pick it up anaphorically (perhaps, but not necessarily quasi-indexically) and so understand it, without thereby counting as able to entertain the very same (demonstrative) thought the speaker expressed.

4. Object-Dependent Singular Thought

The propositional contents of demonstrative beliefs and claims are paradigmatic of a wider class, sometimes called 'singular thoughts' in virtue of the peculiarly intimate cognitive relations to their objects that they embody. Russell thought that the best way to accommodate this sort of strong *de re* belief in semantic theory is to treat the propositional contents they express as having the objects themselves as constituents.[66] Less ontologically alarming contemporary descendants of these Russellian propositions are proposed by those who treat demonstratives and indexicals as devices of *direct reference*, as picking out their objects without employing any conceptual resources or mediating senses.[67] Each of these is motivated by the thought that the behavior of demonstratives and indexicals requires an anti-Fregean semantic theory. McDowell, following Evans, has argued forcefully to the contrary—that the underlying phenomenon is the *object-dependence* of the thoughts or contents expressed by the use of these locu-

tions, and that a thoroughly Fregean account of object-dependent thoughts and singular senses is possible.[68] There are some thoughts one can entertain regardless of whether the singular terms occurring in the sentence tokenings that express them succeed in picking out objects. Thus the claim that the woman who wrote *Sordello* was hopelessly sentimental and syntactically unsound expresses a genuine thought—even though the singular term fails to refer (since Robert Browning authored the work that one would by implication be criticizing by endorsing that claim). Object-dependent thoughts, by contrast, are those that can be entertained only if the singular terms occurring in the sentence tokenings that express them do succeed in picking out objects. Prime among them, the suggestion is, are those propositional contents whose expression requires the use of demonstratives or indexicals.

That essentially demonstrative thoughts are object-dependent is a claim about their accessibility—their accessibility not to different individuals but to the same individual in different possible circumstances. If there is indeed a cup in my hand, the thought I could express by asserting "This cup (in my hand) was on my desk yesterday" is one that might be true or false, depending on the previous career of that cup. But if there is no such cup, if I am hallucinating or dazzled by reflections into mistaking a stone in my hand for a cup, then that thought is not available for me to think.[69] That thought is strongly about a particular cup, and if it is not available to be demonstrated, I am not in a position to think a thought that is strongly about that cup. This need not mean that I am not thinking at all; I may be thinking (falsely) that there is a cup in my hand and that it was on my desk yesterday. That is a thought, but it is not the same one that I would express by the use of the demonstrative in the different case in which I am not mistaken about there being a cup in my hand. Calling singular thoughts and beliefs—those that are strongly about this very object in the way that is paradigmatically expressed by the use of demonstratives and indexicals—'object-dependent' is intended to be a neutral specification of the phenomenon addressed theoretically by Russell's claim that the object itself is a *constituent* in the proposition entertained and by the claim that such thoughts and beliefs refer *directly* to their objects, without the use of concepts or associated senses.

The phenomenon of object-dependent propositional contents is widely thought to be of particular significance for semantic theory. The reason for this is that it makes manifest a fundamental feature of the way in which thought and talk (insofar as they have an empirical and practical dimension) relate to what they are about—the objects that must be consulted in order to assess the *truth* of what is believed and claimed. The expressive role played in semantic assessment by the theoretical vocabulary whose paradigm is 'true' has been explained here in terms of the adoption of perspectivally hybrid deontic attitudes. Treating a claim as *true* is *attributing* one doxastic commitment while *undertaking* another which shares or anaphorically inherits its propositional content. Similarly, treating an action as *successful* is

attributing a practical commitment while undertaking a doxastic commitment with a corresponding propositional content.

The object-dependence of strong *de re* propositional contents can also be understood in terms of perspectivally hybrid deontic attitudes. Taking someone to have a strong *de re* propositional attitude is *attributing* a *doxastic* commitment the expression of which would involve the demonstrative or indexical use of a singular term (or more broadly, of an anaphoric dependent traceable back to such an antecedent)—an attitude that can be *ascribed de dicto* using a quasi-indexical. But it also requires that the one attributing that attitude *undertake* an *existential* commitment regarding the singular-term tokening in question. The object-dependence of a propositional content consists in the additional existential commitment *undertaken* by anyone who *attributes* a status or attitude exhibiting such a content.[70] (Recall that existential commitments were explained in Chapter 7 in substitutional terms, with respect to a special class of designators that are canonical for the sort of existence that pertains to the objects at issue—in this case, accessible spatiotemporal ones.)

For this reason, *de re* ascriptions of epistemically strong *de re* beliefs involve existential commitments. Recall that (weak) *de re* ascriptions as here conceived need not involve such commitments. If someone who has never heard the name 'Pegasus' believes that Bellerophon's horse has wings, I can specify the content of that belief in the weak *de re* way by saying that he believes *of* Pegasus that he has wings, without undertaking any existential commitment to the existence of such a horse. But if I say that he believes of$_{strong}$ Pegasus that he has wings, I am committed to his being able to pick out Pegasus by using a demonstrative, and hence to the spatiotemporal accessibility of that horse in the common environment he and I share, which is what a commitment to the physical existence of the horse comes to on the analysis presented in Chapter 7. That *de re* ascription of an epistemically strong *de re* belief commits me to there being some indexical expression *t* such that

> (a) he believes that $a^{\langle t \rangle}$ has wings, and
> (b) $a^{\langle t \rangle}$ = Pegasus.

Undertaking commitment to that identity involving the quasi-indexical $a^{\langle t \rangle}$ commits me to its antecedent tokening /t/ picking out Pegasus, and this is the source of the existential commitment I undertake in making the *de re* ascription of a belief strongly of Pegasus. The result is that *de re* ascriptions of essentially indexical beliefs provide a way of understanding what Quine was after with his notion of *relational* belief. For his conception of such beliefs embodied both the idea that they involve a special sort of epistemic access to or rapport with objects, and that *de re* ascriptions of them involve existential commitments.

Expressions making explicit the stronger sense of 'of' can then be under-

stood in terms of their role in this kind of ascription. If S attributes the strong *de re* doxastic commitment that would be ascribed *de dicto* by asserting "T believes that she$^{\text{(that woman)}}$ is a spy" and undertakes the doxastic commitment that would be acknowledged by asserting "The one T referred to as 'that woman' is (=) Rosa Kleb," "She$^{\text{(that woman)}}$ is Rosa Kleb," or in suitable circumstances just "She is Rosa Kleb," then T is also committed to the *de re* ascription of that commitment: "T believes of_{strong} Rosa Kleb that she is a spy." The subscript distinguishes the 'of' as expressing a strong *de re* belief by marking the requirement that the underlying *de dicto* ascription involves the use of a quasi-indexical (or a proper name functioning analogously to a quasi-indexical), and hence ascribes a strong *de re* belief.

A scorekeeper who attributes a demonstrative tokening but who does not undertake an existential commitment regarding it—and so does not take it to pick out an actual object—does not take the thought or belief in whose expression it occurs to be strongly of or about any object. What is expressed is taken to be only a mock-thought, one that could be mistaken for a thought with demonstrative content but that in fact has none (though it may be associated with genuinely contentful nondemonstrative thoughts).[71] In this case the scorekeeper must not undertake substitutional commitments with respect to the original demonstrative tokening—and so not with respect to its anaphoric dependents, including quasi-indexicals. It is for this reason that *de re* ascriptions of essentially indexical beliefs involve existential commitments and so count as attributions of successful object-dependent beliefs.

Besides the sort of error involved in taking it that someone (perhaps oneself) has a demonstrative thought when there is no such thought to be had, it is also possible to have two genuine demonstrative thoughts that are strongly about one and the same object, without realizing that they are. In the classic case, one could say of Venus in the morning, "This planet is the last one visible as the sun rises," and of Venus in the evening, "This planet is the first one visible as the sun sets," without realizing that the two claims were about one planet. Conversely, it is possible to have two genuine demonstrative thoughts that are strongly about different objects while believing that they are about one single one. I may at one moment think truly of the cup in my hand (as "this cup") that it is the one that was on my desk last night and, after a moment of confusion, think the same thing of the indistinguishable duplicate that replaces it, without realizing that the first thought is true and the second false, since they are about different cups.

In Fregean terms, what one grasps when one has practical mastery of the use of an expression—understands the thoughts expressed by that use—is its *sense*. The possibility of this sort of individuative mistake regarding the use of demonstratives shows that the notion of *de re* or object-involving senses or conceptual contents is incompatible with the doctrine that such senses are epistemically *transparent*. According to the doctrine of their transparence, grasping a sense entails individuative omniscience regarding it: one

cannot grasp two different senses without realizing they are different, nor the same one twice without realizing that. On this line, the fact that one may fail to realize that the objects picked out by two different senses are different or that they are the same—may endorse false identities and deny true ones—just shows that one may be ignorant about the world, not that by grasping the senses one does not thereby know everything about them. According to this ultimately Cartesian picture of conceptually articulated senses, they lie entirely within the mind, open to its survey in a way that rules out the possibility of error or ignorance. Only in collaboration with the world lying wholly without the mind do they determine referents, however, and it is the epistemic opacity of that second semantic component that introduces the possibility of error and ignorance—which accordingly pertains not to the representational *purport* of those senses but only to its specific *success* or failure.

It may be doubted whether Frege was ever committed to this picture.[72] In any case, it does not work well for the sort of strong aboutness exhibited by essentially demonstrative or indexical claims and the beliefs they express. Here the model of sense and reference that has traditionally been associated with Frege must be modified. If strong *de re* or object-involving senses are admitted, then the transparency principle regarding the individuation of senses must be relinquished; if the transparency principle is retained, then object-involving thoughts must be understood as picking out their objects without the intervention of senses. The latter option is the one pursued by theories that understand strong *de re* claims and beliefs as *directly referential*—namely, that the object picked out by singular terms such as demonstratives, indexicals, and indeed proper names is the *only* contribution their occurrence makes to the claims expressed by the sentence tokenings in which they appear.

Two related phenomena raise particular challenges for such a view. First is the version of Frege's puzzle about identities that arises for the special case of directly referential expressions. How is one to explain the difference in inferential significance between acquiring commitment to a claim of the form $t = t'$ and acquiring commitment to a claim of the form $t = t$? The former, as Frege points out, can be "rich in consequences,"[73] furthering one's knowledge by substitutionally entitling one to new claims; the latter is not something that could ever count as a discovery—it can license no new conclusions. Yet if the two terms are understood as directly referential, it would seem that claims of the two forms say just the same thing: they are strongly about an object, and they say of it just that it is identical to itself. The second phenomenon concerns how this difference is reflected in explicit *ascriptions* of belief or commitment. It seems that one can coherently ascribe the belief that t and t' are not identical even when those singular-term tokenings are expressions that secure an epistemically strong rapport with one and the same object and so according to this approach should be under-

stood as directly referential. But how is one to explain the content of the belief that is being ascribed in that case? The classical Fregean account of such ascriptions turns precisely on the *senses* of the expressions used to specify the content of the ascribed belief—and it is the essence of the direct reference approach to eschew appeal to any such senses.

To understand these phenomena one must look at the way singular reference can be passed on *anaphorically*. Anaphoric chains contribute to both of the theoretical tasks for which Frege postulated senses: they are ways in which objects can be given to us, and they determine the reference of the expressions occurring in them (whose senses they articulate). Anaphoric chains of tokenings—explained (as in Chapter 7) in terms of inheritance of substitutional commitments—provide a model for object-involving, *de re* senses. The access such anaphoric senses provide to the objects they pick out is *conceptual*. For the conceptual is understood in this work in terms of *inferential* articulation, and anaphora is a fundamental mechanism whereby unrepeatable tokenings are linked into repeatable classes subject to *substitutional*, and hence indirectly to *inferential*, commitments (see further discussion below in 9.1). The cognitive accessibility of empirical objects depends in no small part on the possibility of grasping object-involving anaphoric senses—even though such senses are not themselves epistemically transparent to those who grasp them.

Singular thought, which is strongly about particular objects,[74] extends beyond essentially demonstrative or indexical thought. The discussion of the use of quasi-indexicals in ascriptions shows how singular reference, originally secured by the use of demonstratives or indexicals, can be preserved and extended anaphorically. This should come as no surprise, in light of the argument in Chapter 7 to the effect that unrepeatable tokenings, paradigmatically those used demonstratively and indexically, can function as singular terms at all, and so as picking *out* objects, only in virtue of the possibility of picking *up* their references anaphorically, to construct token repeatables that can figure in substitutional commitments. If my claim "*This cup* was on my desk last night" succeeds in being strongly of or about a particular cup, then so does the anaphorically dependent claim (whether in my mouth or in that of another) "And *it* was also there the night before." The singular reference can likewise be extended by the use of explicitly anaphoric definite descriptions, such as "the cup he referred to on that occasion as 'this cup',"which give the sense (anaphoric, not demonstrative) of 'that cup' in sentences such as "That cup is *not* the one that was on your desk last night" when uttered by another, later, out of sight of the cup in question.

Indeed, without such mechanisms one often could not so much as state identities linking essentially demonstrative or indexical thoughts. One cannot, for instance, be in a position to express the strong *de re* version of the discovery that Frege was interested in by saying "This planet [Venus, in the morning] is (=) this planet [Venus, in the evening]." For at any time when one

could perform the one demonstration, one could not perform the other. One must rather say something such as "This planet [Venus in the morning] is (=) that planet, that is, the one I referred to this morning as 'this planet'." Even in cases where one is situated so fortunately as to be able to say "this is (=) this" or "this is that," (where 'that' is used demonstratively, rather than anaphorically), such an utterance has the significance of asserting an *identity* only insofar as it can be used to license *substitutions*. (The appearance of 'is' or '=' can mark that aspiration but cannot by itself make the remark able to perform that function.) That it can be so used depends on the possibility of picking up those demonstrative references anaphorically, using expressions such as 'it', 'that planet', 'the cup N.N. referred to yesterday as "this"', and so on. So much as stating the Frege puzzle about the difference between the claim that would be expressed by a sentence of the form $t = t'$ and one of the form $t = t$ requires the repeatability of t in order to formulate the trivial identity (or the overtly anaphoric reflexive construction in t = itself). It is only *as* initiators of anaphoric chains that demonstrative and indexical to-kenings provide ways of talking or thinking about objects at all, and hence potentially as *strongly* of or about them. These chains provide the *point* of using demonstratives or indexicals, and they articulate the *significance* of doing so.

5. The ^{s}Puzzle about Belief^s

Proper names can be used to express singular thoughts. The ob-ject-involvingness of these uses of proper names has been an important topic addressed by theories of direct reference. It was suggested already in Chapter 7 (Section IV) that the considerations that motivate causal-historical theories of proper names can be understood equally as motivating an *anaphoric* the-ory of proper-name usage. According to such an approach, the tokening-repeatables corresponding to proper names should be understood not as equivalence classes of lexically cotypical tokenings but as anaphoric chains. Such chains are anchored in a tokening that plays the role of anaphoric initiator, which corresponds to the introduction of the name. In favored cases, that initiating token might be a singular term used demonstratively or indexically: "I hereby christen *this ship* 'Beagle'." The role of the introducing token as anaphoric initiator, and hence as reference-fixer, means that even where it has the syntactic form of a definite description, it should be under-stood as used demonstratively—that is, as 'dthat'ed, and so rigidified. This is not a special feature of anaphoric chains that are proper names but a general feature of anaphora. Consider the following discourse:

> #The president of the organization is a brilliant woman. She would have been a brilliant woman even if she had never be-come president of the organization.#

The initiator 'the president of the organization' serves to pick out an individual; anaphoric dependents on it, such as the two tokenings of 'she', continue to refer to that same individual even in counterfactual situations in which one could not pick out that individual with that description or title. Indeed, as this example suggests, this feature of the use of anaphoric dependents is essential to the possibility of expressing such counterfactual circumstances.

This approach assimilates the way proper names express singular thoughts to the mechanism of overtly anaphoric preservation of object-involvingness just considered. On the anaphoric account of proper-name usage, one who uses a tokening of a proper name is committed thereby (in the eyes of those keeping deontic score) to the pragmatic significance of that tokening being determined by the same substitutional commitments that govern its anaphoric antecedents—whatever those are. The one using the proper name need not know what the antecedents are (never mind what the original initiator of the chain is) in order to undertake that anaphoric commitment. Nor, provided the one using the name is generally competent in the relevant practices, need that individual *intend* to use the tokening to pick out the same object as the one uttering some particular antecedent did. The anaphoric commitment to inherit substitutional commitments can be undertaken and attributed in the absence of any such intention, just as one can undertake an assertional commitment as an inferential consequence of acknowledging another one, without intending to or realizing that one has done so. Each scorekeeper who treats a tokening as a proper name, and hence as an anaphoric dependent, must as part of that interpretation of the utterance assign it to some anaphoric chain (or at least, take it that there is some such chain), for it is that chain that determines what commitments one undertakes by the utterance, and hence its pragmatic significance. Both intentions and conventions may be relevant to such scorekeeping and may be appealed to by other scorekeepers in assessing the correctness of an interpretation. But the essential thing to understand theoretically is what it is for a scorekeeper to treat one tokening as anaphorically dependent on another, and that is not a matter of intention or convention but of commitment.[75]

If anaphoric chains of tokenings are the semantically relevant structure that governs the behavior of proper names, then it is easy to understand how different tokenings of the type ⟨Aristotle⟩, for instance, can refer to different people. This happens when they owe allegiance to different anaphoric chains, anchored in different antecedents—in the same way that different tokens of the type ⟨it⟩ can refer to different objects. Individual speakers are not omniscient about the commitments they undertake by their use of various expressions; they need not know much about the anaphoric chains they participate in and appeal to in order nonetheless to have their commitments determined (according to scorekeepers) by those chains. As a result, they may not realize that two singular propositions they are entertaining are strongly about the

same individual, or they may mistakenly think that one individual is at issue when two are. These are the phenomena Kripke addresses in his essay about the use of object-involving expressions in ascriptions of propositional attitude, "A Puzzle about Belief."[76] Seeing through the difficulties Kripke raises requires bringing together the treatment of expressions of strong *de re* belief with the previous discussion of *de dicto* ascriptions.

The puzzle concerns an apparent incoherence in ascriptions of belief that arises when someone uses two different proper names without realizing they refer to the same object, or mistakenly takes different occurrences of one name to refer to different objects. The example that is usually discussed concerns Pierre, who, though generally bilingual in French and English, does not realize that 'Londres' and 'London' are names of the same city. Keeping separate substitutional accounts under the two headings, he endorses both "Londres est joli" and "London is not pretty." The question is whether or not he should be said to believe *that* London is pretty.

The puzzle arises from two seemingly harmless principles governing the specification of the content in indirect discourse ascriptions of belief, which Kripke calls the *disquotational* principle and the principle of *translation*. The translation principle governs the relation between reported tokening (which would express the belief ascribed) and the reporting tokening inside the scope of the 'that' (which specifies the content of the belief ascribed). It is just the principle considered in Section IV above in the discussion of Davidson's "On Saying That." In fact, however, only the disquotational principle is required. As Kripke says after discussing the bilingual case: "Even if we confine ourselves to a single language, say English, and to phonetically identical tokens of a single name, we can still generate the puzzle . . . Only the disquotational principle is necessary for our inference; no translation is required."[77] The case he considers is one where someone learns of the pianist Paderewski, and so it is true to say, based on the avowals he is willing to make:

> Peter believes that Paderewski had musical talent.

But then he learns of the nationalist leader Paderewski and concludes from general beliefs what can be reported, based on his avowals, as

> Peter believes that Paderewski had no musical talent.

From the vantage point of this example, it is clear that the example of Pierre involves a translation principle governing the ascription of beliefs only because it amounts to a case like that of Peter, but in which one of the lexical expressions under which he has segregated his beliefs (in spite of their actually being about one individual) is translated into a different language. Thus the translation principle really adds nothing essential to the case.

The puzzle or paradox comes in two forms. In the weaker form, the sincere avowals of Peter lead us to attribute to him inconsistent beliefs, namely the belief that Paderewski did have musical talent and the belief that Paderewski

did not have musical talent. Now one might think that there is nothing particularly paradoxical about having inconsistent beliefs. Certainly in the *de re* sense, one can believe *of* Benjamin Franklin, as the inventor of bifocals, that he did not invent the lightning rod, and also belief *of* Benjamin Franklin, as the inventor of the lightning rod, that he did invent the lightning rod. But what is at issue is *de dicto* ascriptions of belief. The relevant parallel would accordingly be ascribing both the belief *that* Benjamin Franklin did invent the lightning rod and *that* Benjamin Franklin did not invent the lightning rod. Kripke makes contradictions of the sort exemplified by the Paderewski case particularly telling by invoking in addition a principle that may be called the *transparency of inconsistency:*

> We may suppose that [the believer], in spite of the unfortunate situation in which he now finds himself, is a leading philosopher and logician. He would *never* let contradictory beliefs pass. And surely anyone, leading logician or no, is in principle in a position to notice and correct contradictory beliefs if he has them.[78]

> A straightforward application of the principles of disquotation and translation yields the result that Pierre holds inconsistent beliefs, that logic alone should teach him that one of his beliefs is false.[79]

The principle of the transparency of inconsistency is what underwrites the inference from holding inconsistent beliefs to the claim that logic alone should show that one must be discarded. Again, with respect to another analogous case: "Our Hebrew speaker both believes, and disbelieves, that Germany is pretty. Yet no amount of pure logic or semantic introspection suffices for him to discover his error."[80] In the stronger form of the paradox, the inconsistency is charged not against the one whose beliefs are being reported but against the ascribers. For those ascribers are charged with both taking it that Peter does believe that Paderewski has musical talent and taking it that Peter does *not* believe that Paderewski has musical talent.

It is worth looking closely at the disquotational principle that causes all the trouble. Here is the full passage that introduces and specifies it:

> Let us make explicit the *disquotational principle* presupposed here, connecting sincere assent and belief. It can be stated as follows, where 'p' is to be replaced, inside and outside all quotation marks, by any appropriate standard English sentence: "*If a normal English speaker, on reflection, sincerely assents to 'p', then he believes that p.*" The sentence replacing 'p' is to lack indexical or pronominal devices or ambiguities that would ruin the intuitive sense of the principle (e.g. if he assents to "you are wonderful," he need not believe that *you*—the reader—are wonderful) . . .
>
> A similar principle holds for sincere affirmation or assertion in place of assent.[81]

Kripke says of this principle: "Taken in its obvious intent . . . the principle appears to be a self-evident truth."[82]

There are two principles with quite different functions being endorsed here. One of them connects overt linguistic *avowals* of belief (assertions, or affirmations) with reports, attributions, or ascriptions of belief. It says, in effect, that the best evidence one could have, or very strong evidence, or prima facie evidence that is hard to override, for ascribing a belief with a certain content is a sincere avowal or affirmation with that content. 'Self-evident' is not a word whose use is encouraged or endorsed by the point of view pursued in this work (in view of its standard circumstances and consequences of application), but it can nonetheless be agreed that anyone who denied all principles along these lines would thereby simply have changed the subject and shown that he or she was not talking about our concept of belief at all. At any rate no philosophical ground is to be gained by denying that there is any connection of this sort between avowal and ascription of belief.

The second principle concerns the relation between the expression used to specify the content of the belief avowed, on the one hand, and the expression used to report or ascribe that belief, in indirect discourse, on the other. The relation asserted here is the one that gives the principle Kripke appeals to its name: the *disquotational* principle. Under this heading the claim is that the very same words used to avow the belief are to be used to report it. The relation between direct discourse quotation of what was said and indirect discourse reporting of what was claimed is that of lexical cotypicality. This principle has a very different status from the first. It should not count as "self-evident" in anyone's book. Indeed, as Kripke immediately acknowledges, it is evidently false for a number of different sorts of locution. The whole problem of indirect discourse consists in specifying the relation between the tokening actually uttered, or that the believer could utter to express the belief in question, on the one hand, and the tokening another can admissibly use in reporting what was asserted or believed, on the other. Specifying this relation involves a myriad of subtleties, which the 'disquotational' strategy simply ignores. These are just the subtleties that were discussed in Section IV in connection with McDowell's emendation of Davidson's paratactic theory of the expressive role of content-specifying 'that' clauses in *de dicto* ascriptions of propositional attitude. The upshot of that discussion is precisely that translation plus disquotation (= type repetition) will *not* do in general as an account of the relation between the reported and the reporting tokening in indirect discourse ascriptions.

In the passage above, Kripke explicitly formulates the disquotational principle, on which the paradox or puzzle rests, so as to acknowledge the existence of locutions to which it does not apply. He mentions both the occurrence of indexicals (presumably including demonstratives), and of pronominal or anaphoric constructions. Why should we not conclude from the

puzzle cases offered just that proper names, at least under some circum-
stances, are also locutions for which the disquotational principle is inade-
quate? Such a conclusion seems particularly apt in view of the fact that
demonstratives, indexicals, and tokenings that are anaphorically dependent
on them are the paradigm of object-involving singular-term usage, and that
Kripke sets up the problem about proper names by appeal to the very features
of their use that have made direct reference theories attractive (both for them
and for demonstratives, indexicals, and so on). Kripke nowhere addresses the
possibility that the disquotational principle does not apply to proper names
in the same way and for the same reasons that it does not apply to demon-
stratives, indexicals, and their anaphoric dependents.

The concessive qualification that follows the explicit statement of the
principle does not quite say that the principle is to be taken as holding except
in those cases where it does not. Three classes of exceptions are acknow-
ledged. Consider first "ambiguities that would ruin the intuitive sense of the
principle" (that is, presumably, render it false). What sort of ambiguities are
these? Kripke considers a case which he does *not* take to present the difficul-
ties characteristic of his "puzzle." Suppose that Arthur has in his idiom two
names that share the lexical type ⟨Cicero⟩. Some of his tokenings of that type
are intended to refer to a Roman orator, and others are intended to refer to a
famous spy from the Second World War. Then Arthur would affirm both
tokenings of the type ⟨Cicero was a spy⟩, when the latter is under discussion,
and tokenings of the type ⟨Cicero was not a spy⟩, when the former is under
discussion (the denouncer of Catiline has been accused of many sins, but so
far as Arthur knows, never of being a spy). In such a case Kripke does *not*
take it that Arthur has inconsistent beliefs, believing both that Cicero is and
is not a spy, nor that we have inconsistent beliefs in both taking him to
believe that Cicero is a spy and not to believe that Cicero is a spy. The forms
of words parallel the Paderewski case, but the dual use of the name 'Cicero'
prevents inconsistency, and therefore paradox. This seems to be the sort of
ambiguity that is intended to be excluded by the third phrase in the qualifica-
tion of the disquotational principle.

How does this case differ from that of Paderewski? Peter believes that he
is using different tokens of the type 'Paderewski' systematically as different
names of different individuals, in a way exactly parallel to Arthur's use of
'Cicero'. As it turns out, this is not the case. But to claim that the one
situation is unproblematic and the other is paradoxical evidently requires
more of a difference than simply that Peter's uses are in fact coreferential,
while Arthur's are not. After all, 'the inventor of bifocals' and 'the inventor
of the lightning rod' are coreferential, like the various uses of 'Paderewski',
and there is no suggestion that paradox lurks in the beliefs of someone who
attributes a property to Franklin by the use of the one expression that he
denies to him by the use of the other. Kripke's appeal to the transparency

principle in stating his puzzle commits him, however, to very strong criteria of adequacy on the difference he discerns between these cases. The principle of the transparency of inconsistency, cited above, makes explicit a commitment to Peter's being able to tell his case from Arthur's by the use of "pure logic and semantic introspection alone"—for only this can set up Kripke's dilemma. At the very least, such resources ought to suffice to tell to which cases the disquotational principle applies ("self evidently"!), as in Peter's case, and to which it does not, as in Arthur's. Yet clearly Peter and Arthur are in no position to tell their cases apart. Kripke offers no argument whatever for not assimilating Peter's case (and therefore Pierre's) to Arthur's ambiguity case.[83] Neither identity of lexical type nor identity of referent will do the job. Presumably the difference has something to do with the fact that 'Paderewski' is used as *one* name by the community that Peter and also we, who are reporting on his beliefs, belong to, while 'Cicero' is used by us as *two* names. But these are not facts that are accessible to Peter and Arthur, least of all by "logic and semantic introspection." Thus it is difficult to see how such facts can do the work Kripke needs to have done in formulating the paradox. Again, why is not the proper conclusion from the puzzle cases just that sometimes proper names are so used that the disquotation principle does not apply to them because of dual uses not only of the 'Cicero' type but also of the 'Paderewski' type? An answer is not forthcoming.

6. Tactile Fregeanism: Proper Names as Anaphoric Dependents

The remarks so far have simply been critical of Kripke's argument. But it should be clear that his puzzle can lead to positive conclusions, not so much about belief and its ascription as about proper names. Indeed, it is best thought of as a puzzle about proper names, rather than about beliefs—insofar as it deserves to be called a 'puzzle' at all. A good place to begin is with the observation that one cannot tell simply from the lexical type of an expression whether it is used in such a way that the disquotation principle applies to it. For instance, putting aside the category of proper names, which is currently at issue, only some occurrences of definite descriptions (the least 'Millian' of expression types) are suitable for the application of the disquotational principle. In particular, only 'attributive' uses of definite descriptions are suitable—'referential' uses are not. Thus it is possible to use the definite description 'the man in the corner drinking champagne' to refer to someone who is not in fact drinking champagne. Such 'referential' or speaker's referential uses can for many purposes be assimilated to demonstrative uses. In any case, such cases are not in general happily reported in indirect discourse by attributing "the belief *that* the man in the corner drinking champagne . . ."

Kripke offered three categories of exceptions to the disquotational princi-

ple: ambiguous, indexical, and pronominal uses. It has been suggested here that there might be reasons to assimilate the uses of proper names that are puzzling if they are assumed to fall under the disquotational principle either to the case of ambiguities or of indexical uses. What about the possibility of assimilating them to the pronominal cases? The disquotational principle does not properly describe the relation between the tokening used to avow a belief and the tokening used to report or ascribe that belief in case the former contains an anaphoric dependent, because the ascriber's repeating another token of the same type need not involve anaphorically picking up the same antecedent as the original. Thus if Jones says: "I met a man from the Biology Department yesterday" and if Smith responds with the avowal: "He must be the new statistician," it need not be correct for Brown later to report this belief with the ascription: "Smith believes that he must be the new statistician." Only if special arrangements are made to see to it that the antecedent of the pronoun tokening that appears in the content-specifying portion of the ascribing expression is either the reported tokening or its antecedent can a type-identical term be used in the indirect discourse report.

Is there any reason to think that tokenings of the type 'Paderewski' behave oddly in reports of Peter's beliefs because they are functioning in a way analogous to anaphorically dependent tokenings? Kripke describes his own approach to proper names this way: "Ordinarily the real determinant of the reference of names of a former historical figure is a chain of communication, in which the reference of the name is passed from link to link."[84] This is the "causal-historical" approach to proper-name usage. As has already been suggested, it can usefully be developed in the form of an *anaphoric* understanding of proper names. In the favored case, there is a 'baptism' of an object by using an indexical or definite description to pick out the object the name is to be attached to. This reference-fixing specification becomes the antecedent, standing at the beginning of a "chain of communication, which on the present picture determines the reference" of further tokenings of the name in the chain that is anchored by that initial antecedent.[85] These further tokenings inherit their reference from earlier elements in the chain. This is just anaphora. Anaphoric chains of tokenings create repeatables that can play the same role in substitution inference that is played by classes of cotypical tokenings, for instance, semantically definite descriptions, used attributively. Just so for chains of name tokenings. In each case, the commitments one is undertaking by using a dependent expression late in a chain are to be determined by tracing the chain back. The "legitimacy of such a chain" of name tokenings depends on how the reference is passed,[86] just as with an anaphoric chain. Kripke says that for proper names, differences in the beliefs of users do not change the reference of those name tokenings, so long as the user "determines that he will use the name with the referent current in the community."[87] This sounds like a primitive (and not ultimately satisfactory)

account of what it is to use a pronoun with a certain antecedent (as though one needed to have the *concept* of reference in order to use an expression in an anaphoric and so in fact reference-inheriting way!). In either the canonical anaphoric-dependent case or the case of proper-name tokenings, to take a tokening as continuing a chain is to take it as inheriting its substitution-inferential role from those antecedent tokenings.

Sometimes distinct anaphoric chains of proper-name tokenings are anchored in antecedents picking out different objects, as in the 'Cicero'/'Cicero' case. Sometimes distinct chains are anchored in a single object, as in the 'Paderewski'/'Paderewski' case. Both of these structures can also occur with ordinary anaphoric dependents, such as 'it'. Just as one cannot tell "by pure logic and semantic introspection" whether two chains that one is continuing are anchored in one object or in two for ordinary anaphoric dependents, so one cannot for the anaphoric chains that govern the use of proper names. Just as for canonically pronominal expression types such as 'he' or 'it', cotypicality is no guarantee of coreference—one must consult the anaphoric chain to which a tokening belongs in order to determine its reference. This is just the situation in which the disquotational principle is not applicable. The cases Kripke presents do not generate a puzzle; they just show that proper names can be used in such a way that the disquotational principle does not apply to them. Kripke's own approach to proper-name reference in terms of chains of tokenings suggests exactly why. The result is that these cases present good reason for treating proper names on an anaphoric model.

With the exception of the anaphoric initiator that is their common antecedent, all the tokenings of a single proper name are of the same lexical type. (As cases like Kripke's 'Cicero' case show, however, not all tokenings of the same lexical type need belong to one chain and so be tokenings of one proper name.) This explains why there is no need for a special category of *quasi-names* performing an expressive role in *de dicto* ascriptions of strong *de re* beliefs that corresponds to that performed by quasi-indexicals.[88] Quasi-names would be used to attribute a belief that would be properly expressed by the believer by the assertion of a sentence containing a proper name and would function as anaphoric dependents of such attributed tokenings. But if the ordinary uses of proper names are already anaphoric, with antecedents stipulated to be of the same lexical type, then proper names can function as ascriptional quasi-names without further alteration. The fact that for this reason one can often specify the content of an ascribed belief by using a tokening of the same type as would be used by the believer to express it is responsible for the impression that the disquotational principle ought to apply to proper names, as it does to *attributive* (that is, nondemonstrative, nonanaphoric) uses of definite descriptions. The underlying anaphoric structure of proper-name usage—and so the way in which the disquotational principle incorrectly generalizes from some well-behaved cases—is high-

lighted in cases where different anaphoric chains govern tokenings drawn from the same lexical type.

The invisibility of this alternative anaphoric analysis and the impression that the ^{S}puzzles primarily has to do with belief, rather than with proper names, both arise out of Kripke's commitment to a 'Millian' theory of the semantics of proper names. The opposition Kripke sets up between 'Millian' and 'Fregean' theories of proper name usage is misleading in a number of ways. It is evidently not exhaustive. It is not even clear that Kripke's own theory is best described as Millian—that is, directly referential in the sense that it is not legitimate to appeal to anything other than the referent (in particular, to an anaphoric antecedent or chain) in explaining the contribution the occurrence of a name makes to the use of a sentence in which it occurs. Nor does it seem sensible to treat other sorts of expressions as working this way. Demonstratives such as 'this' are not really "directly referential" because they require implicit sortals to pick out their referents. Pronouns such as 'he' are not, because they convey gender information, are differently used depending on what anaphoric chain they are part of (over and above what their ultimate referent is), and involve an all-but-explicit personal sortal. On the other side of the opposition, Kripke's view is that "the opposing Fregean view holds that to each proper name, a speaker of the language associates some property (or conjunction of properties) which determines its referent as the unique thing fulfilling the associated property (or properties)."[89] It is hard to see how this view can qualify as Fregean. Properties are in the realm of reference, not in the realm of sense—they are not immediately graspable, for Frege. Furthermore, how is this account supposed to extend to the senses of predicates, for Frege? Presumably they are not also specified by a set of properties that picks them out. In any case, the possibility of *de re* senses as Evans and McDowell describe them is overlooked by this way of setting up the issue.

The conceptual contents that are expressed by proper names, which determine their substitutional and so inferential roles, correspond to anaphorically structured constellations of tokenings. These constellations are like Fregean senses in that they determine the referents of the name tokenings whose significance they govern. They are also like Fregean senses in that they provide cognitive access to the particular objects to which they refer—via intersubstitutability equivalence classes of token-repeatables, some of which include demonstrative and indexical tokenings. Such chains anchor our thought and talk in particular objects that it is about—the objects that must be consulted in order to assess the truth of our claims and the beliefs they express. They determine (according to the interpretation of a scorekeeper) not only *that* we are talking and thinking about particular objects by correctly using proper names but also *what* we are thereby talking and thinking about. The connection they establish between our thought and talk and its objects

is so tight that the propositional contents we express by their means are object-involving: the singular or strong *de re* thoughts we express by using proper names when everything goes well (according to a scorekeeper) are not ones we can so much as entertain in cases that are just like those in every way possible except that the anaphoric chains do not (according to the score-keeper) provide access to a unique object.

Conceptual contents expressed by proper names—understood in terms of constellations of tokenings that are articulated by anaphoric commitments regarding the inheritance of substitution-inferential commitments—are un-like Fregean senses as traditionally conceived in that they are not epistemi-cally transparent in their individuation. We can be confused about which anaphoric chain a certain tokening is beholden to, and hence about whether two tokenings belong to the same or to different chains. In that sense we do not always know what we are saying or thinking; even where the proposi-tional contents are object-involving, confusion is possible regarding which objects our strong *de re* beliefs are about. Conceptual (because ultimately inferentially articulated) contents of this sort are best thought of on a *tactile*, rather than a *visual*, model. The Cartesian visual model is the one that gives rise to demands for transparency, to the idea that error and ignorance should be impossible regarding what is in the mind. But it is just this model of inside and outside that makes it unintelligible in the end how what is in the mind should essentially involve representational purport regarding what is outside the mind—how thought can seem to be, and when all goes well in fact be, *about* environing things.

Frege's own favorite metaphor for our cognitive relation to senses is that of *grasping* rather than *seeing.* One can grasp an anaphoric chain as one grasps a stick; direct contact is achieved only with one end of it, and there may be much about what is beyond that direct contact of which one is unaware. But direct contact with one end gives genuine if indirect contact with what is attached to the other end. Indeed, the more rigid the stick (or chain), the better are the antics of what is attached to its far end communi-cated to the end one grips, and the more control one can exert, albeit at a distance, over it. The Cartesian model of conceptual contents restricts them to the part of the stick touching one's hand, at the cost of mystery about how our cognitive reach can exceed that immediate grasp. A *tactile* Fregean se-mantic theory, of the sort epitomized by understanding proper names as constellations of singular-term tokenings articulated by anaphoric commit-ments, effaces this impermeable boundary between the transparency of the mind and the opacity of its objects. The model of thought it presents incor-porates, as two sides of one coin, both the possibility of ignorance of and error about our own concepts and the possibility of genuine aboutness of those concepts and genuine knowledge of the objects with which those concepts put us in touch.

VI. THE SOCIAL-PERSPECTIVAL CHARACTER OF CONCEPTUAL CONTENTS AND THE OBJECTIVITY OF CONCEPTUAL NORMS

1. *Weak and Strong Aboutness*

The strategy pursued in this chapter is to explain the representational dimension of discursive practice by offering an account in deontic scorekeeping terms of what is expressed by the locutions of ordinary language that are used to make it explicit. These are the expressions used to say what claims and beliefs are *of*, *about*, or *represent*, in the sense of what objects must be consulted in order to assess their *truth*. Locutions such as 'of', 'about', and 'represents' play the expressive role of representational locutions in virtue of their use in *de re* specifications of the contents of ascriptions of propositional attitude. The basic way of specifying the propositional content of an ascribed commitment is the *de dicto* style. This style of content specification corresponds to a particular way of indicating a possible sentence tokening *(dictum)* on the part of the one to whom the commitment is ascribed that that individual would, according to the ascriber, acknowledge as expressing the commitment ascribed. The ascriber does this by embedding in the ascription—in regimented ascriptions, as a clause inside the scope of a 'that'—a sentence tokening that stands in a relation to the indicated tokening that is potentially complicated (where it must be anaphorically articulated) but that in the simplest cases amounts to mere repetition from the perspective provided by the ascriber's different language or indexical circumstances, preserving the claim expressed.

The *de re* style of specifying the contents of ascriptions derives from this basic *de dicto* style by allowing substitutions that are licensed by the commitments *undertaken*, rather than those *attributed*, by the ascriber. Thus the very same commitment that S can ascribe *de dicto* by asserting "T believes that Benjamin Franklin could speak French" can be ascribed *de re* by asserting "T believes *of* the inventor of bifocals that he could speak French," provided only that S endorses the substitution-licensing identity "Benjamin Franklin is (=) the inventor of bifocals"—regardless of whether S also takes it that T endorses that identity. The *de re* form accordingly specifies what individual, according to the ascriber, it is whose properties must be investigated in order to determine whether the ascribed belief is *true*. S will endorse T's claim—take it to be true—if and only if S endorses the claim that the inventor of bifocals could speak French. *De dicto* ascriptions specify the content of the attributed commitment from the point of view provided by what the one to whom that commitment is attributed would, according to the attributor, acknowledge. *De re* ascriptions specify the content of that same commitment from the point of view provided by what the one attributing the commitment would acknowledge. These are two different sorts of ascription, two ways of specifying the content of a single commitment, not ascriptions of two different sorts of belief or commitment.

The previous section of this chapter acknowledges that besides two styles of content specification that can be employed in ascribing any belief whatsoever, one of which can sensibly be called *de re*, there are also two different sorts of belief, one of which can sensibly be called *de re*. Besides weak *de re* ascriptions of belief there are strong *de re* beliefs—which like the others can be ascribed either *de dicto* or *de re*. Distinguishing strong *de re* beliefs is the result of adopting a "frankly inegalitarian attitude" toward different ways of specifying an object—in particular, by acknowledging the special sort of rapport secured by the capacity to pick out an object by certain sorts of uses of demonstratives, indexicals, and proper names. The task of the last section was to account in scorekeeping terms for the use of the sentences that express and ascribe such strong *de re* beliefs. This story had three parts. The first examined the sort of strong *de re* rapport established by using certain kinds of singular terms to pick out objects. The prime example considered was that of first-person indexicals. Some attention was also paid to the anaphoric extension of demonstrative and indexical acquaintance, widening the category of singular or strong *de re* thought to include what is expressed in favored cases by the use of proper names. The second part of the story about strong *de re* belief concerned *de dicto* ascriptions of such beliefs. The key to understanding these was found in an anaphoric construal of quasi-indexicals—expressions used to ascribe the use of indexicals and demonstratives. (Understanding proper names as anaphoric chains of cotypical tokenings makes it clear how this one category can play both the expressive role of undertaking a commitment and that of ascribing it, so that no separate category of quasi-names is required.) The final move is then to derive the structure of *de re* ascriptions of strong *de re* beliefs from that of these *de dicto* ones.

Although the special, strong *de re* acquaintance with the individual is secured by the availability to the believer of a special sort of singular term to pick it out, what is secured is acquaintance with that individual. In *de re* ascriptions of strong *de re* beliefs, the individual in question can be specified in any way available to the ascriber; these are not restricted to the use of singular terms that themselves express strong *de re* rapport with the object of the belief. So if according to S, but unknown to T, Rosa Kleb is the only actual KGB colonel mentioned in Ian Fleming's *From Russia with Love*, then according to S, T believes of$_{strong}$ the only actual KGB colonel mentioned in Ian Fleming's *From Russia with Love* that she is a spy. As with *de re* ascriptions of ordinary beliefs, information that would be conveyed by a *de dicto* ascription about *how* the believer represents an object is suppressed, in favor of information about *what* object is represented; in the case of strong *de re* beliefs, *de re* ascriptions specify *that* the believer has a belief that is strongly of or about an object, and *which* object it is, but not (as a *de dicto* specification of the content of the same belief would) *how*—by the use of what sort of expression—that rapport is established.

Thus the second, stronger sense of 'of' is built on, and must be understood in terms of, the first and more fundamental sense. The representational dimension of discourse, its aboutness or intentional directedness at objects and states of affairs, is accordingly to be explained in the first instance as what is expressed by (weak) *de re* specifications of the contents of ascribed commitments—whether or not those commitments themselves qualify as strong *de re*. That *de re* style of content specification in turn reflects the *social* difference in doxastic perspective between the one *attributing* an ascribed commitment and the one *undertaking* it, the one keeping score and the one whose score is kept, the interpret*ing* interlocutor and the interpret*ed* interlocutor. *De re* ascriptions are the explicit expression of perspectivally hybrid deontic attitudes, the attribution by the ascriber of one discursive commitment (doxastic or practical, and so propositionally contentful) and the undertaking by that ascriber of another (substitutional) commitment. Commitments of these kinds are conceptually contentful in virtue of their inferential significance; talk about the *representational* dimension of the conceptual content of intentional states should be understood in terms of the *social* dimension of their inferential articulation.[90]

2. The Essentially Perspectival Character of Conceptual Contents

Conceptual content is understood in this work as what can be made explicit in discursive practice. Discursive practice has as its defining core *claiming*. Claims are a kind of commitment that can be understood in terms of the functional role things of this kind play in social scorekeeping practices—practices and practical attitudes that accordingly can be thought of as instituting this sort of deontic status. Commitments of this kind are in turn appealed to both to pick out specifically *propositional* contents and to pick out the specifically *sentential* locutions used to express them. Because it must incorporate performances with the pragmatic significance of claim-ings, discursive practice is *linguistic* practice. Linguistic practice is what makes it possible for something to be made *explicit* (and logical practice is what makes it possible for linguistic practice itself to be made explicit). Making anything explicit, *saying* it, requires using one linguistic expression rather than another. Specifying the content of a discursive commitment by using a particular sentence involves undertaking an *expressive* commitment concerning the inferential role of the locution employed.

Expressive commitments were mentioned already in Chapter 2, where it was pointed out that the use of any particular linguistic expression involves an *inferential* commitment—a commitment to the propriety of the inference from its circumstances of application to its consequences of application.[91] As with any sort of deontic status, expressive commitments should be understood in terms of the attitudes of *undertaking* and *attributing* them. These

differ not only in *who* endorses the propriety of the inference involved but also (typically, though not invariably) in *what* inference is taken to be involved. For what one takes to follow from what (committively or permissively) depends on what collateral premises one is committed or entitled to. This principle applies no less to the inferences incorporated in conceptual contents than to those that relate conceptual contents among themselves. To see the inferential relations internal to concepts as pragmatically and methodologically on a par with those that are external to concepts is the essence of the sort of holism about meaning that Quine introduces in "Two Dogmas of Empiricism," just as a result of considering the role of auxiliary hypotheses in determining the significance of particular commitments. Neither category can be traded in theoretically for (eliminated in favor of) the other in toto—any more than the implicit and the explicit can. But there is a deep practical equivalence between inferential commitments that are implicitly endorsed by employing particular bits of vocabulary and those that are overtly endorsed in the scorekeeping practice of taking or treating certain discursive commitments as consequentially related to, and so as providing reasons for, others.

We discursive creatures—rational, logical, concept-using ones—are construed here in *expressive* terms; we are the ones who can make it explicit. One manifestation of that overarching feature of the approach is the methodological principle that what is implicit is to be made theoretically, as opposed to practically, intelligible precisely by exercising our defining attribute—by making it explicit. When the inferences implicit in the use of a word are made explicit in the form of conditionals, the fact that the proprieties governing them are relative to a background of collateral commitments is manifest. Which conditionals one endorses depends in part on what other claims one endorses. When the circumstances and consequences of application potentially associated with a particular word are made propositionally explicit, which conditionals connecting them an interlocutor endorses varies with the other commitments, doxastic and inferential, that the interlocutor undertakes. This relativity of explicit inferential endorsements to the deontic repertoires of various scorekeepers reflects the underlying relativity of the inferential endorsements implicit in the concepts expressed by particular words, according to various scorekeepers. A word—'dog', 'stupid', 'Republican'—has a different significance in my mouth than it does in yours, because and insofar as what follows from its being applicable, its consequences of application, differ for me, in virtue of my different collateral beliefs (and similarly for circumstances of application—consider 'murder', 'pornographic', 'lyrical').

One way of characterizing this situation, which at least since Quine has been sufficiently popular among those endorsing holism that it is often thought to be one of its necessary consequences, is to say that different interlocutors—most importantly those undertaking commitments and those

keeping score by attributing them—use different concepts, attach different meanings to their words. There are not enough words to go around, however; we all use basically the same words to express the contents of our commitments. As was pointed out in Chapter 7, this way of putting things accordingly seems to threaten the intelligibility of mutual understanding and so of successful communication. The discussion of indexicals and demonstratives in the previous section provides a different model for construing communication and the conceptual contents that are communicated, in the face of variations in the significance of utterances in different mouths. The fact that the word 'I' can never have the significance in my mouth that it does in yours (not just practically or for the most part, but in principle) in no way precludes my understanding what you express by using it.

Communication is still possible, but it essentially involves intralinguistic *interpretation*—the capacity to accommodate differences in discursive perspective, to navigate across them. Anaphora can secure not only coreference but token repeatability across the different repertoires of commitments that correspond to different interlocutors. This capacity becomes explicit in the *ascription* by one individual of propositionally contentful commitments on the part of another individual. There quasi-indexicals can secure contact with and express the understanding of otherwise unrepeatable indexical and demonstrative tokenings by the one to whom a commitment is attributed. Anaphorically indirect definite descriptions make these connections explicit. Not only can one attribute an expressive commitment without undertaking it by making explicit what is implicit in the use of scare quotes—"the one John referred to as 'that murderous lowlife'"—but (at the usual risk of perhaps puzzling prolixity) by their means one can make explicit what is implicit in the use of quasi-indexicals—"the one John referred to as 'I'," "the book he referred to as 'this'."

Even where the anaphoric mechanisms that make them possible are not on the surface, the difference between the expressive commitment acknowledged by one who uses a word to make a claim explicit as an overt assertion and the expressive commitment that would be undertaken by the use of the same word on the part of one who attributes that claim is explicitly taken account of and managed in ascriptions. For the difference between the expressive commitments undertaken by the ascriber and those attributed by the ascriber is overtly marked by the difference between expressions occurring *de re* and those occurring *de dicto* in the content-specifying portion of the ascription. Thus *T* may have beliefs that could be expressed in *T*'s mouth by the assertion of the sentences

> That murderous lowlife should be locked up

or

> The seventh god has risen.

S may not be in a position to undertake the expressive commitments involved in using the expressions 'that murderous lowlife' or 'the seventh god' in the way needed to express these claims and so may not be able to ascribe them *de dicto* by asserting

> *T* claims that that murderous lowlife should be locked up

or

> *T* claims that the seventh god has risen.

Nonetheless, *S* can pick up the expressions anaphorically and express the claims *de dicto* using scare quotes or indirect definite descriptions, which explicitly disavow the expressive commitments those expressions incorporate (attributing, rather than undertaking, responsibility for their use in picking out who or what is talked about):

> *T* claims that ^{s}that murderous lowlife^s (the one *T* refers to as 'that murderous lowlife') should be locked up

or

> *T* claims that ^{s}the seventh god^s (the one *T* refers to as 'the seventh god') has risen.

It is the same anaphoric capacity to express such identities as

> He (^{s}that murderous lowlife^s, the one *T* refers to as 'that murderous lowlife') is (=) the mayor

or

> It (^{s}the seventh god^s, the one *T* refers to as 'the seventh god') is (=) the sun

that make it possible for *S* to specify the content of the very same belief *de re*, as

> *T* claims of the mayor (^{s}that murderous lowlife^s, the one *T* refers to as 'that murderous lowlife', as 'that murderous lowlife') that he should be locked up

or

> *T* claims of the sun (^{s}the seventh god^s, the one *T* refers to as 'the seventh god', as 'the seventh god') that it has risen.

Where the parenthetical expressions are omitted, the underlying anaphoric links that make such *de re* content-specifications possible are not apparent on the surface of the ascription.

In ascribing a discursive commitment, the ascriber must specify its content and to do so must undertake various expressive commitments. Since the

same words in the ascriber's mouth often do not express the same claim that they would in the mouth of the one to whom a claim is ascribed, the content specifications must take account of the difference in discursive perspective between the ascriber and the target of the ascription. Claims must often be differently expressed from the two points of view. Nonetheless, what is attributed explicitly in a *de re* ascription can be the very same claim that would be acknowledged, using different words, in an assertion by the one to whom it is ascribed. The difference in the inferential significance of words in one mouth and words in another, due to differences in collateral commitments, should not be understood as meaning that one interlocutor cannot strictly be said to understand what another says; it should only be taken to mean that the content they both grasp (if all goes well in the communication of it) must be differently specified from different points of view. Conceptual contents are *essentially expressively perspectival*; they can be specified explicitly only from some point of view, against the background of some repertoire of discursive commitments, and how it is correct to specify them varies from one discursive point of view to another. Mutual understanding and communication depend on interlocutors' being able to keep two sets of books, to move back and forth between the point of view of the speaker and the audience, while keeping straight on which doxastic, substitutional, and expressive commitments are undertaken and which are attributed by the various parties. Conceptual contents, paradigmatically propositional ones, can genuinely be shared, but their perspectival nature means that doing so is mastering the coordinated system of scorekeeping perspectives, not passing something nonperspectival from hand to hand (or mouth to mouth).

The expressive resources that have been considered here—the distinction between vocabulary used *de dicto* and vocabulary used *de re* in ascriptions, anaphorically and ascriptionally indirect definite descriptions, quasi-indexicals, scare quotes, and so on—provide mechanisms for making explicit the systematic differences in scorekeeping significance that constitute essentially socially perspectival conceptual contents. The representational dimension of such contents—what is expressed in *de re* ascriptions using terms such as 'of', 'about', and 'represents'—is one manifestation of their socially perspectival character. Thus wherever propositional or other conceptual contents are attributed, there is some implicit appeal to an *I-thou* social practice in which one scorekeeper interprets the performances of another. The representational dimension of propositional content is conferred on thought and talk by the social dimension of the practice of giving and asking for reasons, in virtue of which inferentially articulated contents are essentially perspectival. As the point was put at the opening of this chapter, the conceptual contents employed in *monological* reasoning, in which all the premises and conclusions are potential commitments of one individual, are parasitic on and intelligible only in terms of the sort of content conferred by *dialogical* reasoning, in which the issue of what follows from what essentially involves

assessments from the different social perspectives of scorekeeping interlocutors with different background commitments.

We sapients are discursive scorekeepers. We keep track of our own and each other's propositionally contentful deontic statuses. Doing that requires being able to move back and forth across the different perspectives occupied by those who undertake commitments and those who attribute them. Reidentifying conceptual contents through shifts in doxastic and practical point of view requires *interpretation* in Wittgenstein's sense—substituting one expression of a claim (he says 'rule') for another. The ordinary discursive understanding of propositional contents involved in all explicit knowing-*that* something is the case must be understood in terms of this implicit, practical-interpretive scorekeeping as knowing-*how*. Wittgenstein's pragmatism about norms—his insistence that norms made explicit in principles are intelligible only against a background of norms implicit in practices—is one of the master ideas orienting this entire enterprise. As was argued in Section II, however, one of the results of following that pragmatist idea out in detail by working out discursive scorekeeping practices sufficient to confer propositional contents that can codify norms in the form of claims and principles is that Wittgenstein is wrong to take that pragmatist methodological principle to be incompatible with understanding discursive practice as involving interpretation (in his sense) at every level, including the most basic. The capacity to interpret remarks, to substitute different expressions of a claim, rule, or principle—the propositional form in which things are made explicit—is a basic component of the fundamental practical capacity to grasp and communicate essentially perspectival conceptual contents.

Conceptual content is essentially something expressible, but it is expressible only from the point of view of some background repertoire of deontic statuses. What sentence should be used to express a particular semantic content varies from individual to individual. Assessments of the pragmatic significance of using a linguistic expression on a particular occasion—for instance what would entitle its utterer to that speech act, and what it commits its utterer to—differ depending on what claims are available to serve as auxiliary hypotheses. Conceptual content consists in the systematic relations among the various pragmatic significances. Thus grasping the semantic content expressed by the assertional utterance of a sentence requires being able to determine both what follows from the claim, given the further commitments the scorekeeper attributes to the assertor, and what follows from the claim, given the further commitments the scorekeeper undertakes. Where the expressive power of ascriptional locutions is available, this implicit capacity to keep and correlate multiple sets of books is made explicit in the availability of both *de dicto* and *de re* specifications of the contents of attributed discursive commitments. These features—the relativity to social perspective of the pragmatic significance of using a particular sentence or other linguistic expression, hence of which semantic content that expression

makes explicit, hence of which expression can be used to specify a given content—compose the socially perspectival character of semantic contents.

That perspectival character of the contents conferred on linguistic expressions by their role in deontic scorekeeping practices is the detailed consequence of the general methodological policy of subordinating semantic theory to pragmatic theory. This is the methodological (as opposed to the normative) sense of 'pragmatism'—that the point of treating sentences and other bits of vocabulary as semantically contentful is to explain their use, the practices in the context of which they play the role of expressing explicitly what is implicit in discursive deontic statuses. (Which is not to say that those statuses are intelligible apart from consideration of such expression— as always, talk of what is implicit involves tacit reference to the process by which it can be made explicit.) The perspectival character of the expression of conceptual contents is a structural manifestation of the peculiarly intimate connection between semantics and pragmatics mandated by methodological pragmatism. In such a context, particular linguistic phenomena can no longer reliably be distinguished as 'pragmatic' or 'semantic'.

A prime example is that of anaphora. This chapter and the previous one have described the crucial expressive role played by anaphoric connections in linking inherently unrepeatable performances (paradigmatically those involving demonstratives, indexicals, and even proper names) into repeatable structures that can express conceptual contents by being governed by (having their significance determined by) indirectly inferential substitutional commitments. It is anaphoric chains that tie different perspectives together and make it possible for scorekeepers to correlate them, structuring different pragmatic significances according to unified, though perspectival, semantic contents. Is anaphora a pragmatic or a semantic phenomenon? What about the distinction between *de re* and *de dicto* uses of linguistic expressions to specify the contents of ascribed commitments? Neither the latter, which expresses differences between scorekeeping perspectives, nor the former, which secures connections among them, is happily assimilated to either category. From the vantage point provided by this approach, philosophers have on the one hand maintained a tolerably clean line between semantics and pragmatics only by largely ignoring anaphoric phenomena (while lavishing attention on demonstrative and indexical ones that presuppose them) and have on the other hand been precluded by commitment to such a discrimination from a proper appreciation of the expressive function and significance of weak *de re* ascriptional locutions.

3. *The Objectivity of the Norms Governing the Application of Concepts Is Part of Their Social-Perspectival* Form

Methodological pragmatism expresses the recognition that semantics is not ultimately intelligible apart from pragmatics; semantic con-

tents can be studied in abstraction from the pragmatic scorekeeping significance of altering deontic statuses and performing speech acts that have those contents, only provisionally and temporarily, subject to a promise of subsequent reunification. For commitments and the performances that express them are conceptually contentful in virtue of the role they play in the game of giving and asking for reasons. Their contents are perspectival because the inferential articulation of deontic scorekeeping practice has an irreducible social dimension. The social aspect of inferential articulation is apparent in the variations of substitutional commitments from one interlocutor to another, which correspond to different identity beliefs, and in the anaphoric commitments that link speech acts by different interlocutors. In fact, the perspectival relation between semantic contents and pragmatic significances reflects a prior perspectival relation between deontic *statuses* and deontic *attitudes*. For deontic statuses are *instituted* by the same essentially social scorekeeping practice of adopting discursive deontic attitudes that *confers* conceptual content on them. The basic phenomenon is the social-perspectival character of discursive *norms*—in deontic terms, the way in which *being* committed is to be understood in relation to being *taken* to be committed, by scorekeepers occupying different deontic points of view.

The discussion of the social-perspectival character of the representational dimension of semantic contents, expressed explicitly in the distinction between *de re* and *de dicto* ascriptions, can shed light on important aspects of discursive scorekeeping social practices—aspects that have necessarily been left somewhat obscure in the discussion of pragmatics up to this point. Thus it is appropriate to look once more at the relations envisaged between deontic statuses, which are the counters in terms of which discursive score is kept, and the deontic attitudes, whose adoption and alteration constitute the activity of scorekeeping that has been described as instituting those statuses. In particular, the expressive and explanatory raw materials are finally in place to address the fundamental issue of the *objectivity* of conceptual norms. It was pointed out already in Chapter 1 that it is a critical criterion of adequacy on any account of concepts that it make sense of a distinction between how they *are* applied in fact, by anyone or everyone, and how they *ought* to be applied—how it would be *correct* to apply them.

Failure to satisfy this criterion of adequacy is one of the primary objections that was offered against what were characterized as *I-we* construals of the social practices in which conceptual norms are implicit. Such construals fund a distinction between what particular individuals *treat* as or *take* to be a correct application of a concept, on the one hand, and what *is* a correct application, on the other, by contrasting individual takings with communal ones. This is the standard way of understanding objectivity as intersubjectivity. The cost of adopting this way of understanding the significance of the social dimension of discursive practice is, unacceptably, to lose the capacity to make sense of the distinction between correct and incorrect claims or

applications of concepts on the part of the whole community. This sort of view was criticized on the grounds that it depends on an illegitimate assimilation of linguistic communities to the individuals who participate in them—treating communities as producing performances and assessing them, undertaking commitments and attributing them. Besides its commitment to this mythological conception of communities, this view was found to be objectionable in that it avoids *regulism* (finding norms only in the explicit form of rules) only by falling into the complementary error of *regularism*, by rendering norms implicit in practice as mere regularities—albeit regularities of communal rather than individual practice.[92] Even if these further complaints are bracketed, however, the fact remains that an essential part of the representational dimension of our concepts—the way they purport to apply to an objective world—is that they answer for the ultimate correctness of their application not to what you or I or all of us *take* to be the case but to what actually *is* the case. Part of what it is for our concepts to be *about* an objective world is that there is an *objective* sense of correctness that governs their application—a sense of appropriateness that answers to the objects to which they are applied and to the world of facts comprising those objects. Even communally sanctioned takings or regularities of takings of what is correct concept application that are universal within a community can still be *mis*takings: even if all of us agree and always will agree that the mass of the universe is small enough that it will go on expanding forever, the possibility remains that we are all wrong, that there is sufficient matter undetected by us to make it collapse gravitationally.

One of the central challenges of an account of conceptual norms as implicit in social practice is accordingly to make sense of the emergence of such an *objective* notion of correctness or appropriateness of claims and applications of concepts. To begin with, the possibility of making such a notion intelligible in the present context appears to be threatened by the ontological disparity between concepts and the objects they apply to. For concepts are essentially perspectival, yet if they are to be *objectively* true or false *of* objects, those objects must be understood as nonperspectival in a strong sense. It makes no sense to specify or express a propositional or other conceptual content except from some point of view—which is subjective, not in a Cartesian sense, but in the practical sense that it is the point of view of some scorekeeping subject. How then is one to understand such propositional contents as answering for their truth, such conceptual contents as answering for their correct application, to some *one* way things are? Traditional philosophy says that beliefs are many, but the truth is one; the same point arises here in the contrast between scorekeeping perspectives, which are many, and the world, which is one. It is part of the representational purport of conceptual contents that thought and talk give us a perspectival grip on a nonperspectival world. To say that objects and the world of facts that comprises them are what they are regardless of what anyone takes them

to be requires a notion of *objective correctness* of claiming and concept application that is *not* perspective-relative in the way that what is *taken* to follow from what is.

In fact, however, understanding what is expressed by assessments of the objective correctness of applications of concepts requires appeal neither to nonperspectival facts (= true propositional contents) nor to community-wide commitments to propositional contents. Rather, the distinction between claims or applications of concepts that are objectively correct and those that are merely taken to be correct is a structural feature of each scorekeeping perspective. Indeed, the required notion of objective correctness is just what is expressed by *de re* specifications of the conceptual contents of ascribed commitments. For ascriptions in the *de re* style specify the objects that determine the truth or falsity—that is, the objective correctness—of the ascribed claim. Suppose the Constable has said to the Inspector that he himself believes that the desperate fugitive, a stranger who is rumored to be passing through the village, is the man he saw briefly the evening before, scurrying through a darkened courtyard. Suppose further that according to the Inspector, the man the Constable saw scurrying though the darkened courtyard is the Croaker, a harmless village character whom no one, least of all the Constable (who knows him well), would think could be the desperate stranger. Then the Inspector can identify the objective representational content of the Constable's claim by an ascription *de re:* "The Constable claims *of* the Croaker (a man who could not possibly be the fugitive) that he is the fugitive." Of course he does not take it that the Constable claims *that* the Croaker (a man who could not possibly be the fugitive) is the fugitive. The Constable claims only *that* the man he himself saw scurrying through a darkened courtyard is the fugitive. For the Inspector, the contrast between the *de re* and the *de dicto* content specifications is the contrast between saying what the Constable has in fact, willy-nilly, undertaken commitment to—what object his claim is *about,* in the sense that matters for assessments of truth—on the one hand, and what the Constable *takes* himself to be committed to, acknowledges, on the other hand.

Given the way the world actually is (according to the Inspector)—in this case, given who actually was scurrying through the courtyard—the Constable has without realizing it committed himself to a claim that is true if and only if the Croaker is the fugitive. That is the objective content of his commitment. The *de dicto* ascription specifies the content of that same commitment *subjectively,* that is, from the point of view of the one who acknowledges that commitment. In keeping the two correlated sets of books on the Constable that are made explicit as ascriptions in *de re* and *de dicto* form, the Inspector implicitly distinguishes between the deontic *status* undertaken by the Constable and the deontic *attitude* adopted by the Constable. This is the distinction between what the Constable is *committed* to by what he acknowledges and what he *acknowledges.* The difference arises

because one is committed to the inferential consequences of what one acknowledges, but one may nevertheless not acknowledge those consequences. What follows from the claim expressed by a certain sentence, an integral part of its content, can be assessed (by a scorekeeper) either objectively—using facts, true claims, as auxiliary hypotheses—or subjectively, using as auxiliary hypotheses only claims that the one who endorses the original claim acknowledges commitment to. It is this difference that is expressed by the difference between ascriptions *de re* and ascriptions *de dicto*.

As the terms are used here, to have *undertaken* a commitment is just to *be* committed; it is for a commitment to be in force. One can count as having adopted this attitude, and so as occupying this status, even if there is no overt performance that is the undertaking in virtue of which one is committed. Some commitments (and entitlements) we just come with, by default. In order to understand how, in spite of this coincidence, status and attitude can nonetheless be categories distinguished from the point of view of every scorekeeper, one must appreciate the perspectival character of deontic statuses. What is causally efficacious (for instance in action and perception) is the attitude of *acknowledging* a commitment. Commitments may be undertaken by default, by acknowledgment, or consequentially, as a result of what one does acknowledge. Indeed the content of the commitment one undertakes by a speech act of acknowledgment (assertion) depends on these consequences. Acknowledging a commitment can be identified with attributing it to oneself.[93] Thus undertaking a commitment is to be distinguished from attributing it to oneself, which is only one species of that attitude. Attribution can be seen to be the fundamental deontic attitude. Undertaking commitments can be understood in terms of attributing them if the social articulation of scorekeeping attributions is kept in mind: an interlocutor can count as having undertaken a commitment (as being committed) whenever others are entitled—perhaps in virtue of that interlocutor's performances—to attribute that commitment.

The attributing of commitments (and the kind of self-attribution that is acknowledging them) is an *immediate* deontic attitude, while undertaking commitments includes not only the immediate deontic attitude of acknowledging them but also *consequential* deontic attitudes. Deontic statuses are just such consequentially expanded deontic attitudes. But the consequences of a particular acknowledgment are assessed differently from different perspectives—that is, by different attributors. From the vantage point of any particular scorekeeper, what one is *really* committed to by an acknowledgment (paradigmatically the assertive utterance of a sentence), what *really* follows from the claim (and hence its objective content), is to be assessed by conjoining it with truths—that is, statements of fact. But what plays this role for a scorekeeper is the set of sentences by the assertion of which the scorekeeper is prepared to *acknowledge*, and so undertake, doxastic commitment. Thus immediate deontic attitudes determine consequential ones, and so

deontic statuses, from each scorekeeping perspective. From the point of view of each scorekeeper, there is for every other interlocutor a distinction between what commitments that individual *acknowledges* and what that individual is really committed to—between (immediate) deontic attitudes and deontic status (or consequentially expanded attitudes). But how this line is drawn in particular varies from scorekeeper to scorekeeper. This perspectival structure is what has been implicitly appealed to throughout by talk of the *institution* of deontic statuses by deontic attitudes.

In this way, every scorekeeping perspective maintains a distinction in practice between normative status and (immediate) normative attitude—between what is objectively correct and what is merely *taken* to be correct, between what an interlocutor is actually committed to and what that interlocutor is merely *taken* to be committed to. Yet what from the point of view of a scorekeeper is objectively correct—what from that perspective another interlocutor is actually committed to by a certain acknowledgment—can be understood by us who are interpreting the scorekeeping activity entirely in terms of the immediate *attitudes*, the acknowledgments and attributions, of the scorekeeper. What appears to the scorekeeper as the distinction between what is objectively correct and what is merely taken to be or treated as correct appears to us as the distinction between what is acknowledged by the scorekeeper attributing a commitment and what is acknowledged by the one to whom it is attributed. The difference between objective normative status and subjective normative attitude is construed as a social-perspectival distinction between normative attitudes. In this way the maintenance, from every perspective, of a distinction between status and attitude is reconciled with the methodological phenomenalism that insists that all that really needs to be considered is attitudes—that the normative statuses in terms of which deontic score is kept are creatures instituted by the (immediate) normative attitudes whose adoption and alteration is the activity of scorekeeping.

On this account, objectivity is a structural aspect of the social-perspectival *form* of conceptual contents. The permanent possibility of a distinction between how things *are* and how they are *taken* to be by some interlocutor is built into the social-inferential articulation of concepts. The distinction is in the first instance available to each scorekeeper regarding the commitments of *others*, of those to whom the scorekeeper *attributes* commitments, for it just reflects the difference in perspective provided by the different inferential significance of claims in the context of auxiliary premises provided by commitments acknowledged (and so undertaken) by the scorekeeper and commitments that are attributed to others by that scorekeeper. Although grounded in essentially social, other-regarding scorekeeping, however, the possibility of a distinction between how things actually are and how they are merely taken to be by some interlocutor remains a structural feature, even, as will be seen below, in the case of attributions to oneself.

4. I-Thou Symmetry of Subjective Discursive Attitudes and Objective Discursive Statuses

It was pointed out above that the account of (weak) *de re* ascriptions as expressing the difference in social perspective between *attributing* a commitment (in the paradigmatic case, a doxastic one) and *undertaking*, by acknowledging, a commitment (in the paradigmatic case, a substitutional one) just generalizes the previous account of assessments of truth to include nonpropositional conceptual contents. The objective representational content of someone's belief or remark can be specified in a *de re* ascription:

> Max believes of quinine that it prevents jungle fever (Max represents quinine as preventing jungle fever).

Such a characterization specifies the objects that must be consulted in order to assess the objective truth or falsity of the belief—what it is true *of*, if it is true at all. This connection emerges particularly clearly if what is exported to *de re* position in the ascription is not just a singular term, predicate, or sortal but (a nominalized version of) an entire sentence:

> The senator believes of the first sentence of the Communist Manifesto that it is true.

In these cases, that the expressions occurring in the scope of the *de re* 'of' express what matters for determining the *truth* of the ascribed claim is made explicit by the fact that the ascription-structural anaphoric dependent left as a trace of the exportation within the scope of the *de dicto* 'that' is a prosentence, which accordingly is formed using 'true'.

Our practical grasp of the objective dimension of conceptual norms—normative assessments of the objective truth of claims and objective correctness of applications of concepts—consists in the capacity to coordinate in our scorekeeping the significance a remark has from the perspective of the one to whom the commitment it expresses is attributed and its significance from the perspective of the one attributing it. This requires recognizing the different specifications of the same claim that correspond to extracting its inferential consequences and antecedents in the context of other commitments that are acknowledged as true by the scorekeeper, on the one hand, and extracting them in the context of other commitments acknowledged by the target of that scorekeeping, on the other. This is just the difference between employing as auxiliary hypotheses claims that are true (according to the scorekeeper) and employing as auxiliary hypotheses claims that are merely *held* true (according to the scorekeeper) by the interlocutor whose commitments are being assessed. Thus every scorekeeping perspective incorporates a distinction between what is (objectively) true and what is merely (subjectively) held true.

This account accordingly provides one detailed way of understanding what

lies behind some of Davidson's most important claims, quoted with approval back in Chapter 3:

> If this account of radical interpretation is right, at least in broad outline, then we should acknowledge that the concepts of objective truth and of error necessarily emerge in the context of interpretation. The distinction between a sentence being held true and being in fact true is essential to the existence of an interpersonal system of communication . . . The concept of belief thus stands ready to take up the slack between objective truth and the held true, and we come to understand it just in this connection . . . Someone cannot have a belief unless he understands the possibility of being mistaken, and this requires grasping the contrast between truth and error—true belief and false belief. But this contrast, I have argued, can emerge only in the context of interpretation, which alone forces us to the idea of an objective, public truth.[94]

The notion of objective truth conditions makes explicit what is implicit in our grasp of the possibility of mistaken belief and so of the distinction between what is *merely* held (true) and what is *correctly* held (true). It emerges only in the context of interpretation—that is, discursive scorekeeping—because that is the practical activity in which the commitments acknowledged (held true) by one interlocutor are compared and contrasted with those acknowledged (held true) by another, the scorekeeper who attributes the first set.[95]

The objectivity of conceptual norms—the dimension of correctness in the application of conceptual contents assessed according to how the things they are applied to actually are rather than to how they are *taken* to be—is here given an essentially *social* construal. Objectivity appears as a feature of the structure of discursive intersubjectivity. But traditionally intersubjectivity has been understood in the *I-we* way, which focuses on the contrast between the commitments of *one* individual and the commitments of the *community* (collectively), or those shared by *all* individuals (distributively). In the broadly Davidsonian account offered here, by contrast, intersubjectivity is understood in the perspectival *I-thou* fashion, which focuses on the relation between the commitments *undertaken* by a scorekeeper interpreting others and the commitments *attributed* by that scorekeeper to those others. From the point of view of this latter sort of understanding of intersubjectivity, *I-we* accounts mistakenly postulate the existence of a *privileged* perspective—that of the 'we', or community. The objective correctness of claims (their truth) and of the application of concepts is identified with what is endorsed by that privileged point of view. The identification of objectivity with intersubjectivity so understood is defective in that it cannot find room for the possibility of error regarding that privileged perspective; what the community *takes* to be correct *is* correct. The community, it may be said, is *globally* privileged.

Understanding the sort of privilege that is at issue in these claims can, like

any other normative status, be approached by distinguishing the *circumstances* under which it is appropriately acquired, on the one hand, and the *consequences* of possessing it, on the other. The consequences of a perspective being privileged in the relevant sense—what that privilege consists in—is that one *cannot* hypothetically adopt a third-person point of view with respect to it and evaluate it from the outside by contrasting what it merely *takes* to be true with what actually *is* true. Thus a scheme that identifies truth with what is taken true by all the members of a community, or by the experts in a community, or what will always be taken true by them, or what would be taken true by them under some ideal conditions for inquiry, thereby identifies circumstances of application for a perspective to be privileged in just this sense. Truth is *defined* as how things are according to a discursive perspective that meets certain conditions; it is settled in advance that any perspective from which a distinction appears between how things seem from such a privileged point of view and how things in fact are is itself without any authority at all. According to an account that incorporates a global privilege of this sort, the distinction between what is true and what is taken-true does not vary from one perspective to another.

By contrast, according to the *I-thou* construal of intersubjectivity, each perspective is at most *locally* privileged in that it incorporates a structural distinction between objectively correct applications of concepts and applications that are merely subjectively taken to be correct. But none of these perspectives is privileged in advance over any other. At first glance this egalitarian attitude may seem just to put off the question of what is really correct—so that one must choose between relinquishing the notion of objectivity by acquiescing in a regress of attributors assessing attributions and securing it by acknowledging an infallibly authoritative taker (for instance, the community). The alternative is to reconstrue objectivity as consisting in a kind of perspectival *form*, rather than in a nonperspectival or cross-perspectival *content*. What is shared by all discursive perspectives is *that* there is a difference between what is objectively correct in the way of concept application and what is merely taken to be so, not *what* it is—the structure, not the content.

The crucial feature of the perspectival structure of objectivity is the *symmetry* of state and attitude between ascriber and the one to whom a commitment is ascribed. The author of a *de re* ascription treats it as specifying the objective representational content of the attributed commitment (as specifying the status in question, and what it is really about), and its *de dicto* correlates as specifying the subjective attitude that the target of the attribution has toward that state—what the one whose status it is *takes* it to be about. From the point of view of the one to whom the status is ascribed, however, things are the other way around. The attributee (a woman, let us say) treats her own specifications of the content of the commitment she has undertaken as authoritative—she takes it that she knows what she means,

what she has committed herself to, and what she is talking about. The *de re* ascription by the other (a man, let us say), insofar as its terms diverge from those appearing in the corresponding *de dicto* ascriptions (which answer directly to what the attributee acknowledges) expresses from her point of view simply the attitude of the ascriber—his version of her actual status. His divergent *de re* characterization of the content of the commitment she undertook just indicates a difference between the commitment he *attributes* to her, and what she actually undertook. The difference between status and attitude, between commitment and its attribution, between what is correct and what is merely taken to be correct, between objective content and subjective view of it, is as essential a feature of the attributee's point of view as it is of the attributor's—though the author and the target of ascriptions line up these distinctions with that between the author's *de re* and *de dicto* attributions in a complementary way. In each case, from each point of view, there are distinct attitudes of taking or treating as correct and taking or treating as taken or treated as correct. This symmetric pair of perspective types, that of attributor and attributee, each maintaining this fundamental normative distinction, is the fundamental social structure in terms of which communities and communal practice are to be understood. This symmetric *I-thou* social distinction is presupposed by the *I-we* social distinction appealed to by the other kind of construal of intersubjectivity.

Its symmetry ensures that no one perspective is privileged in advance over any other. Sorting out who should be counted as correct, whose claims and applications of concepts should be treated as authoritative, is a messy retail business of assessing the comparative authority of competing evidential and inferential claims. Such authority as precipitates out of this process derives from what various interlocutors say rather than from who says it; no perspective is authoritative as such. There is only the actual practice of sorting out who has the better reason in particular cases. The social metaphysics of claim-making settles what it means for a claim to be true by settling what one is doing in *taking* it to be true. It does not settle which claims *are* true—that is, are *correctly* taken to be true. That issue is adjudicated differently from different points of view, and although these are not all of equal worth, there is no bird's-eye view above the fray of competing claims from which those that deserve to prevail can be identified, nor from which even necessary and sufficient conditions for such deserts can be formulated. The status of any such principles as probative is always itself at issue in the same way as the status of any particular factual claim.

5. Objectivity Proofs

The objectivity of conceptual norms, it has been claimed, is a reflection of the perspectival distinction between undertaking and attributing inferentially articulated commitments. Attributions are made explicit by

ascriptions; the perspectival structure that objectivity consists in can be expressed explicitly in propositionally contentful commitments. Using pieces of logical vocabulary whose expressive roles have been specified systematically in these pages, the challenge to *I-we* construals of conceptual norms that was put in Chapter 1 is to show that it is the case neither that

(i) $(p)[p \rightarrow (S) (S \text{ claims that } p)]$[96]

nor, perhaps more threateningly, that

(ii) $(p)[(S) (S \text{ claims that } p) \rightarrow p]$.

These entailments together would codify the equivalence of p, the obtaining of the state of affairs represented by p, its being *true* that p (given the prosentential rendering of the expressive role of 'true'), to universal or communal *commitment* to p. One of the motivating criteria of adequacy of the present account is that it underwrite no such collapse of an objective state of affairs into a communal attitude. It should be obvious that neither of these conditionals holds on the *I-thou* deontic scorekeeping account, given the story that has been told about the expressive role of propositional-attitude-ascribing locutions as explicitating commitments and conditionals as explicitating inferences. But to say that is not to say that there is no point in offering *proofs*, which can help engender illumination as much as conviction.

The quantified conditional (i) might be called the *No Communal Ignorance Condition*. It is surely not plausible, and so not a desirable consequence of any theory. It is worth working through an argument in a bit of detail here, however, since the principles appealed to will be of use in more interesting demonstrations below. The conditionals that are of interest here are those that express incompatibility entailments (introduced in Chapter 2). For these express a *modal* force, $p \rightarrow q$ saying in effect that it is not *possible* for p to obtain without q obtaining. And this is the sense desired for the objectivity-denying principles (i) and (ii) above. The official definition of this sort of conditional is that $p \rightarrow q$ just in case everything incompatible with q is incompatible with p—so that 'Wulf is a dog' incompatibility-entails 'Wulf is a mammal', because everything incompatible with the latter is incompatible with the former.[97] To show that the conditional in (i) does not hold, then, it is enough to find some claim q that is incompatible with its consequent, and not with its antecedent.

To see that (i) is not a consequence of the present account, it suffices to notice that nothing in that account requires that one who *undertakes* a commitment to p, and is furthermore *entitled* to that commitment, therefore must *attribute* that commitment to *anyone* else (never mind to *everyone* else). Since propositional contents are incompatible just in case commitment to one precludes entitlement to the other,[98] this means that p will not in general be incompatible with ~$(S$ claims that $p)$—except for a few very special cases.[99] So it is possible to pick some particular p^{α}, S^{α} such that it is not the

case that $p^\alpha/\sim(S^\alpha$ claims that $p^\alpha)$ (using 'p/q' to express p's incompatibility with q). This is to say that $\sim(S^\alpha$ claims that $p^\alpha)$ is not incompatible with the antecedent of the conditional in (i). But it *is* incompatible with the consequent: (S) $(S$ claims that $p^\alpha)/\sim(S^\alpha$ claims that $p^\alpha)$. For the right-hand side just denies a substitution instance of the left-hand side. Thus for any (nonascriptional) p^α, this recipe shows how to produce a q^α that is incompatible with the consequent of (i) but not with the antecedent, which shows that the entailment does not obtain for *any* (nonascriptional) p.

The quantified conditional (ii) might be called the *No Communal Error Condition*. It is also not plausible, and so not a desirable consequence of any theory. It is, however, an unavoidable consequence of *I-we* theories of conceptual norms. (As pointed out in Chapter 1, this is one basis on which McDowell rightly criticizes Wright's and Kripke's Wittgenstein.) To show that this conditional is not a consequence of the deontic scorekeeping account, it is necessary to find a claim that will defeat the incompatibility entailment it makes explicit. Such a defeating claim is one that is incompatible with the consequent of the conditional, but not with its antecedent. Thus it suffices to find a p^α and a q^α such that q^α/p^α and *not* $q^\alpha/(S)$ $(S$ claims that $p^\alpha)$.

Intuitively, what is wanted is just something that is incompatible with a claim but not with everyone's believing the claim. Given that actual conceptual roles have been officially defined only for logical locutions, a bit of subtlety is required to produce a pair of claims that *provably* stand in the desired incompatibility relations. Here is a recipe for producing such a pair. Let p^α be:

(p) $(\sim!xDx$ claims that $p)$,

which claims that the unique D does not have any beliefs; and let q^α be:

$!xDx$ claims that p^α,

that is:

$!xDx$ claims that (p) $(\sim!xDx$ claims that $p)$,

which claims that the unique D believes that the unique D does not have any beliefs (so the unique D believes this *of* the unique D—but not, of course, *as* the unique D). For instance, '$!xDx$' might be 'the object weighing 100 lbs. or more that made the scrape marks on the flagstones in the garden'. Then it might be that John endorses p^α because he believes that the scrape marks were made by a rock rolling down the hillside into the garden. Since rocks do not have beliefs, he believes that what made the marks does not have beliefs. But I, who take it that John unknowingly made those marks with his hobnailed boots, may endorse q^α.

Notice, first, that q^α/p^α, since q^α attributes a belief to $!xDx$ and p^α denies that $!xDx$ has any beliefs. But second, it is *not* the case that $q^\alpha/(S)$ $(S$ claims

that p^α). For q^α is just an instance of that quantified claim and so is entailed by it: if (S) $(S$ claims that $p^\alpha)$, then $!xDx$ claims that p^α.[100] Thus q^α is incompatible with the consequent of the conditional (ii), but not with the antecedent, and so the conditional does not hold. Not all the counterexamples to (ii) have this form, but the pattern indicated shows how to construct a whole family of them (notice that it is not even essential that the singular terms involved be definite descriptions).

The explicit expression of the structural distinction of perspective between undertaking and attributing commitments is a general acknowledgment by each interlocutor of the possibility for *any* S and p that S believes that p but it is not true that p. Rehearsing rigorously the coherence of such an attitude in particular cases does not show, however, that this principle extends even to the self-attributions expressed by self-ascriptions. It does not yet show the coherence of the attitude that would be endorsed by saying:

It is possible that (*I* believe that p and it is not true that p).

It is the coherence of such an acknowledgment that one would expect to be hard to fund on an *I-thou* social account of conceptual norms.

The equation of what one is saying in claiming that p with a claim about universal or communal belief threatens *I-we* social accounts of conceptual norms; it is no surprise to find that it is not a threat to the *I-thou* variety. What would be more impressive is to see a similar sort of wedge driven between the use of $\ulcorner p \urcorner$ and $\ulcorner I$ believe (or claim) that $p \urcorner$ on the deontic scorekeeping account. That there is an apparent difficulty in avoiding a strong equivalence between these two very different claims is evident from the fact that in any language with the expressive resources both of propositional-attitude-ascribing locutions and of strong first-person pronouns, each interlocutor will be committed to what is expressed by $\ulcorner I$ claim [am committed to the claim] that $p \urcorner$ when and only when that interlocutor is committed to what is expressed by $\ulcorner p \urcorner$. An inability to distinguish these claims would show that there is a problem with the *local,* perspectival privileging of each scorekeeper's repertoire of commitments that corresponds, in the *I-thou* rendering of discursive social practice, to the failure to account for the objectivity of conceptual norms that afflicts the *I-we* rendering of social-discursive practice in virtue of its *global* (or quantificational) privileging of the communal repertoire.

Nonetheless, utterances of $\ulcorner I$ claim that $p \urcorner$ and $\ulcorner p \urcorner$, though having the same assertibility conditions, express different propositional contents; as the point is usually put, they have different truth conditions. One wants to say that their equivalence is a pragmatic rather than a semantic matter—that is, that it derives from features of claim*ings*, rather than of what is claim*ed*. The Geach-Frege test for distinguishing these was endorsed and exploited in Chapter 5. It depends on the fact that pragmatic force or significance (para-

digmatically the assertional variety) attaches to a whole speech act, not merely to the utterance of an embedded component of an asserted sentence. Thus one can strip off the pragmatic consequences of claiming by embedding a content, paradigmatically as the antecedent of a conditional. The practical cash-value of the different truth conditions associated with $\ulcorner p \urcorner$ and $\ulcorner I$ claim (or believe) that $p \urcorner$ is that they behave differently as the antecedents of conditionals.[101] In particular, $p \to p$ holds generally, while (I claim that p) $\to$ p does not. Thus the real challenge to an *I-thou* social account of conceptual norms is posed, not by the conditionals (i) and (ii), but by:

(iii) (p) $[p \to (I$ claim that $p)]$

and, more threateningly,

(iv) (p) $[(I$ claim that $p) \to p]$.

Showing that the conditional in (iii) (the *No First-Person Ignorance Condition*) does not hold in the deontic scorekeeping account requires finding a claim that is incompatible with what is expressed by $\ulcorner I$ claim that $p \urcorner$, but not with p itself. Consider $\ulcorner I$ do not claim that $p \urcorner$ (or anything that entails it, such as 'I make no claims' or 'I do not exist'). This is evidently incompatible with the ascription that appears in the consequent, since it is just the denial of that claim. Is it incompatible with the antecedent? That I cannot be entitled to both p and (I do not claim that p) is just the tenor of Moore's paradox. But this does not mean that the *contents* of these claims are incompatible. For if they were incompatible, then *anyone* committed to the one content would be precluded from entitlement to the other. But if I endorse the explicit ascriptions

S claims that I do not claim that p, and S claims that p,

I have not attributed incompatible claims to S.[102] That is, I could take S to be *entitled* to both these commitments. Thus they are not incompatible. For one involves what commitments S *attributes* to me, and the other involves what commitments S *undertakes*, and these do not collide.[103]

Ascriptional locutions make explicit the possibility of taking up hypothetically a sort of third-person scorekeeping attitude toward my own present commitments and entitlements (much as I must do for my *past* commitments and entitlements in any case). Here such ascriptions show that what precludes entitlement both to the claim that p and to my denial of a self-ascription is a pragmatic matter concerning attitudes, not a semantic matter concerning the contents to which they are addressed. My denial that I claim that p collides with what I am *doing* (claiming that p), not with what I am *saying* (that p). To distinguish these, I must look at someone *else's* attitudes toward the *same* contents. The social dimension of their inferential articulation is essential to the semantic nonequivalence of $\ulcorner p \urcorner$ and $\ulcorner I$ claim (hold) that $p \urcorner$.

The intuition underlying the distinction between what is expressed by $\ulcorner p \urcorner$ and $\ulcorner I$ believe (or claim) that $p \urcorner$ on the deontic scorekeeping account is straightforward. Although whenever one is *committed* to one of these contents, one is committed to the other, the same does *not* go for entitlements. The easiest way to see this is to think about someone with incompatible commitments p^α and q^α. One of the distinctive virtues of understanding intentional states in the first instance in terms of normative statuses rather than in causal-functional terms is precisely that sense can be made of the undertaking of commitments with incompatible contents. The subject of such commitments will not be *entitled* to any of them, but this in no way precludes nonetheless being committed. If p^α/q^α, then if I am committed to both, I cannot be entitled to either. But I can still be entitled to the ascriptional commitment expressed by $\ulcorner I$ claim that $p^\alpha \urcorner$, for this just correctly ascribes one of the commitments I in fact have. Thus though *commitment* to what is expressed by $\ulcorner p \urcorner$ and by $\ulcorner I$ claim that $p \urcorner$ go hand in hand, *entitlements* to those commitments do not.

This idea provides a general recipe for producing counterexamples to the conditional in (iv) (the *No First-Person Error Condition*). Such defeasors are claims that are incompatible with the consequent of that conditional, but not with its antecedent. But *any* claim q such that q/p can serve this role. For it will not be the case that $q/($I claim that $p)$. I *can* be entitled to the ascriptional commitment expressed by $\ulcorner I$ claim that $p \urcorner$, even if I am committed also to q, which is incompatible with p. Then I just have incompatible commitments (namely p and q) and am correctly ascribing one of them to myself (perhaps my commitment to p arises as a consequence of another commitment I have overtly avowed by asserting it, and I have just realized this about my deontic status). Thus (iv) does not hold.

The result of these discussions of conditionals (i)–(iv) is that the perspectival notion of objectivity that arises out of the deontic scorekeeping account of inferentially articulated conceptual contents can be *shown* not to collapse into *I-we* quantificational intersubjectivity, and it can be *shown* not to collapse into a mere privileging of one's own perspective. Admittedly, this is to demonstrate only that a fairly weak necessary condition on a conception of objectivity has been satisfied—that the contents of ordinary claims, such as "Snow is white" and "The mass of the universe is great enough to produce eventual gravitational collapse," are not equivalent to those of any claims about who is committed to what. But it is no small thing to be able to *prove* that a notion of objectivity that satisfies even these minimal constraints can be funded starting with a notion of *social* practice. Facts are true claims—'claims' in the sense of what is claim*ed*, which does not depend on what people do, not of claim*ings*, which do. The demonstrations just presented define a robust sense in which the facts as construed in this work are independent of what anyone or everyone is committed to. The *claim-making* practices described here are accordingly properly understood as making pos-

sible genuine *fact-stating* discourse, for they incorporate practices of assessing claims and inferences according to their *objective* correctness—a kind of correctness that answers to how things actually are, rather than to how they are *taken* to be, by anyone (including oneself) or everyone.

6. Conclusion

A central task of this work has been to show how such objective conceptual norms can be made intelligible in terms of social-deontic scorekeeping practices governing attitudes of taking or treating oneself and others as having inferentially articulated commitments and entitlements. The fine structure of the arguments presented above shows that essential appeal is made to both the *social* articulation of practical attitudes into attributing and undertaking (in addressing (iii)) and the *deontic* articulation of normative statuses into commitments and entitlements (in addressing (iv)) in the *pragmatics*, in showing the *objectivity* of the *semantic* (inferential) contents those practices confer. These are, of course, the two central structural elements in terms of which the deontic scorekeeping pragmatics has been elaborated.

The avowed aim of the model of discursive practice motivated and introduced in Part 1 and developed in detail in Part 2 has been to describe scorekeeping practices that are sufficient to confer various sorts of conceptual content: to begin with, fundamental nonlogical *propositional* contents, then the sorts of conceptual contents associated with predicates and singular terms (including pronouns, demonstratives, definite descriptions, and proper names), and finally the specifically *logical* expressive conceptual content of conditionals, negation, quantifiers, identity locutions, traditional semantic vocabulary, and ascription-forming operators, which help make explicit crucial implicit features of the inferential and social articulation of the scorekeeping practices that confer conceptual contents in the first place. A major criterion of adequacy of that account is that it explain how *objectivity* can arise as a structure within *intersubjectivity* construed in the *I-thou* fashion characteristic of the deontic scorekeeping approach to discursive practice. Satisfying this condition required showing how the *representational* dimension of conceptual contents—the sort of correctness of concept application that answers to how things are with the things represented and that contrasts in principle with the sort of correctness that answers only to how things are taken to be by some individual or whole community—can be understood in terms of the *inferential* articulation that defines those contents. The understanding that emerges from this discussion construes the objective representational dimension of conceptual contents as a reflection of the *social-deontic* articulation of the practices of giving and asking for reasons, which confer conceptual contents on the states, attitudes, performances, and expressions suitably caught up in them.

It turns out to be a consequence of the inferentialist way of conceiving conceptual content that it makes sense to specify the content of a state only from some point of view, relative to some set of collateral concomitant commitments, which can serve as auxiliary hypotheses in inferences involving it. For both the attributor and the attributee of any contentful state, there are two relevant sets of background commitments available in determining the practical significance (for what else one should go on to do or be committed to) of adopting a state with a specified content—that of the attributor and that of the attributee. Thus two socially related kinds of perspective are always in play, for each interlocutor. This fact, it has been claimed, secures and gives meaning to the possibility of a genuinely *normative* significance for the occurrence of contentful states. It is also what is expressed by *representational* idioms such as ascriptions *de re*. It is because of this fundamental social deontic structure, then, that propositional and other conceptually contentful states are always *representationally* contentful states. This analysis of the nature of the objective representational norms that govern the application of concepts makes it possible to see why *only* what plays a suitable role in essentially *social,* indeed *linguistic,* discursive deontic scorekeeping practices should count as *conceptually* contentful in the fundamental sense. The understanding of intentional or conceptual contentfulness that is finally arrived at vindicates the initial commitment to understanding discursive practice as social linguistic practice.

Appendix: The Construction and Recursive Interpretation of Iterated Ascriptions That Mix *De Dicto* and *De Re* Content Specifications

The expressions that serve to specify the content of the commitments attributed by undertaking assertional commitment to basic ascriptions can play two different sorts of expressive role. In the regimented language employed here, this distinction of roles is marked by a distinction of two sorts of position in which content-specifying expressions can occur in ascriptions. Expressions may occur either in the scope of a 'that' operator (what is called, following the tradition, "*de dicto* occurrence") or in the scope of an 'of' operator ("*de re* occurrence"). The leading idea of the explanatory strategy developed in this chapter is that the significance of an expression's occurring in *de dicto* position is that the expressive commitment to the effect that the content of the attributed assertional commitment can properly be expressed by the use of that expression is *attributed* along with the ascribed assertional commitment, while the significance of an expression's occurring in *de re* position is that the expressive commitment to the effect that the content of the attributed assertional commitment can properly be expressed by the use of that expression is *undertaken,* along with the assertional commitment to

the whole ascription. The distinction between *what* is represented, the objective, *de re*, relational content of what is ascribed, on the one hand, and *how* it is represented, its subjective, *de dicto*, notional content, on the other, is based on this fundamental social-perspectival distinction of deontic attitudes (the distinction between undertaking a commitment and attributing one). One of the distinguishing characteristics of the present approach to propositional-attitude-ascribing locutions is understanding what is expressed by the difference between ascriptions *de dicto* and ascriptions *de re* as indicating the difference between two perspectives or attitudes an ascriber can adopt when specifying the content of an ascribed state, rather than as distinguishing two *kinds* of state (as Quine does), or two *components* in the content of any intentional state (as McGinn does).[104]

One of the criteria of adequacy of any account of this difference is its capacity to deal with *iterated* ascriptions—ascriptions of assertional commitments that themselves have ascriptional contents, expressions for which accordingly contain *embedded* ascriptional expressions. For interpretation of such compound expressions requires the recognition of many more than just two ways in which expressions can function in specifying the contents of ascribed commitments. It turns out to be straightforward to extend the social-perspectival account to handle the complexities of iteration. It is much less clear how the motivation behind distinguishing *de dicto* and *de re* ascriptions as ascriptions of different kinds of belief, or of different components of beliefs, fares once iterated ascriptions are taken onboard.

How, then, can all the iterated ascriptions be constructed in the regimented language of the scorekeeping model? To keep things under control, two simplifying assumptions are adopted—working at the level of types rather than tokenings, and only considering the case of singular-term exportation. These are straightforwardly dispensable in favor of more general formulations. Consider a basic nonascriptional content expressed using an *n*-adic predicate $\Phi\langle x_1, x_2, \ldots x_n\rangle$. In ascribing commitment to a claim of this form, one might attribute the expressive commitments associated with the use of all of the terms used to specify the content, in the pure *de dicto* form

S_2: S_1 claims that $\Phi\langle t_1, t_2, \ldots t_n\rangle$.

Or one might undertake all those expressive commitments and syntactically export all of the terms to *de re* position, in the pure *de re* form

S_2: S_1 claims of $\langle t_1', t_2', \ldots t_n'\rangle$ that $\Phi\langle it_1, it_2, \ldots it_n\rangle$,

where each it_i is an ascription-structural anaphor dependent on t_i'. But such exportation need not be an all-or-none thing. Some terms may remain in *de dicto* position while others are removed to *de re* position, as in

Russell believed of Hölderlin's roommate that he was not a worthy successor to Kant,

where 'Kant' remains in *de dicto* position, but the other term has been exported to *de re* position, leaving the anaphoric trace 'he'. Thus arrayed between the pure *de dicto* and pure *de re* forms will be a variety of mixed ascriptions, in which some terms appear in *de dicto*, and some in *de re* position. So the general form of first level ascriptions is

$$S_2: S_1 \text{ claims of } \langle t_1', t_2', \ldots t_k' \rangle \text{ that } \Phi \langle t_1, t_2, \ldots t_n \rangle,$$

where k is less than or equal to n, and for all i less than or equal to k, there is a j less than or equal to n such that t_j is an ascription-structural anaphoric dependent of t_i', symbolically: Depends $(/t_j/, /t_i'/)$. Since each term can appear in two positions, either exported or not, corresponding to an n-adic predication $\Phi \langle x_1, x_2, \ldots x_n \rangle$ there will be 2^n different first-level ascriptions.

Each of these ascriptions still contains n independent terms in its content-specifying regions. The *de dicto* positions will always contain n argument places. Each of them is filled either by an independent term or by an ascription-structural anaphoric pronoun. But corresponding to each anaphoric pronoun is exactly one term that is exported to *de re* position. Such exportation accordingly does not change the total number of terms occurring in the content specification. All the forms of first-level ascriptions of commitment to a nonascriptional claim involving n argument places can then be thought of as (ascriptionally complex) $n + 1$ place predications. (The extra argument place is that which specifies the one to whom the commitment is ascribed; it occurs outside the content-specifying regions of the ascriptional expression.) These first-level ascriptions of the form $\Psi \langle t_1, \ldots t_n \rangle$ can now themselves be treated as specifying the content of commitments that can be ascribed, by second-level ascriptions. A second-level ascription is one like

Russell claims that Hegel claims of Pluto that it does not exist,

in which the commitment ascribed is itself an ascriptional commitment. Clearly all of the terms that occur independently in the first-level ascriptional content $\Psi \langle t_1, \ldots t_n \rangle$ (that is, all the terms except the ascription-structural anaphoric pronouns left as syntactic traces of terms exported to *de re* position) are available either to be left in what is *de dicto* position with respect to the outermost, second-level, ascriptional content-specification, or to be exported to the *de re* position of that outermost, second-level, ascriptional content-specification. Thus there can be second-level forms such as

$$S_2 \text{ claims of } t_1 \text{ that } S_1 \text{ claims of it that } \Phi \langle it_1, t_2 \rangle$$

and

$$S_2 \text{ claims of } t_1 \text{ that } S_1 \text{ claims that } \Phi \langle it_1, t_2 \rangle$$

that differ in that in the first what is exported to the second-level *de re* position occurred in *de re* position in the embedded first-level ascription as

well, whereas in the second what is exported to the second-level *de re* position occurred in *de dicto* position in the embedded first-level ascription.

Just as the first time around, exportation does not change the total number of independently occurring terms, and each of the $n + 1$ independent terms available for possible exportation in the first-level ascriptions can either be exported or not. So for each of the 2^n first-level ascriptions, there are 2^{n+1} second-level ascription forms, or $2^n \cdot 2^{n+1} = 2^{2n+1}$ second-level ascriptions (based on the nonascriptional n-adic expression $\Phi \langle x_1, x_2, \ldots x_n \rangle$) in all. For the general case of m^{th}-level iterated ascriptions, there will be $2^{n + (n+1) + (n+2) + \ldots (n + m - 1)}$ distinct ascription forms, which is $2^{n \cdot m + (m^2 - m)/2}$. If one treats *de dicto* and *de re* ascriptions as ascriptions of two different kinds of belief, then this is how many different kinds there are, not just two. And if one sees first-level ascriptions of the two sorts as specifying two different components of the content of the state that is attributed, then this is how many different components one is committed to discerning in the content of an m^{th}-level iterated ascription, not just two. The complication in this calculation results from the argument place for the target of the ascription—which behaves like a term in the scope of the *de re* operator 'of', except for not having ascription-structural anaphoric dependents. Putting those occurrences aside, an n-ary nonascriptional predication generates $2^{n \cdot m}$ possible mixed ascription forms, where m is the number of iterated applications of the ascription-forming locution.

A criterion of adequacy of an account of the content of ascriptions is that it determine for each of these myriad iterated ascription-forms the pragmatic significance, in context, of undertaking commitment to an assertion with that form. The theory must offer a reading of each, specifying what an ascriber becomes committed to by asserting it. In the context of the sort of pragmatics or speech-act theory in play here, this means saying how the deontic score is changed by an ascriptional undertaking, which is to say what attitudes it expresses. For instance, looking only at a single iteration of ascribing operators and at a one-place predicate, two of the forms that must receive interpretations are:

    (i) S_2: S_1 claims of t_1 that S_0 claims of it$_1$ that $\Phi(\text{it}_1)$.
    (ii) S_2: S_1 claims of t_1 that S_0 claims that $\Phi(\text{it}_1)$.

These examples are representative of the new sorts of structural anaphoric connections across ascriptional boundaries that become possible with iteration. How do the attitudes involved in these complex ascriptions unpack according to the rules for the regimentation suggested in Section IV?

Suppressing type-tokening niceties, one can begin reading (i) by stripping off the outermost *de re* occurrences:

    (i′) S_2: For some term x_2, S_1 claims that [S_0 claims of x_2 that $\Phi(\text{it}_2)$], and $t_1 = x_2$.

The second-level *de re* ascription by S_2 is interpreted in terms of a nonascriptional identity (that is, symmetric substitution-inferential) commitment undertaken by S_2 and a second-level *de dicto* ascription by S_2. The undertaking of a *de dicto* ascriptional commitment (of whatever level) is itself readily interpreted in terms of the attributing it expresses. Bracketing, in the interests of simplicity, elaborations required to deal with indexicals, and foreign languages (details that the discussion of Section IV shows how to reintroduce as needed), substitution instances of the first clause of (i') are interpreted by

(i") S_1: S_0 claims of t_2 that $\Phi(it_2)$,

that is, by attributions of first-level ascriptions, in this case, *de re* ones. And now the same procedures that were applied to turn (i) into (i') can be applied to (i"), followed in turn by the procedures that turned (i') into (i"). Those procedures suffice to interpret n^{th}-level ascriptions in terms of attitudes toward $(n - 1)^{th}$-level ascriptions. Repeatedly turning the crank on this machinery provides a recursive procedure for assigning a reading to each of the arbitrarily complex iterated ascriptions in the hierarchy.

The procedure is the one followed with the example just considered. First, strip off the terms occurring in the scope of the outermost *de re* operator, resulting in an undertaken identity commitment and an attributed *de dicto* ascriptional commitment at the same ascriptional level. Then trade that *de dicto* ascriptional commitment for an attributing at the next lower ascriptional level. The construction of the hierarchy of regimented iterated ascriptions proceeded by arbitrary repetitions of two sorts of formation rule, one corresponding to *de dicto* ascriptions, another to *de re*. So these interpretive rules match the ones used in constructing the expressions, and it follows that for any complex regimented ascription, a finite number of repetitions of the two interpretive steps will render the complex attitude in terms of simpler, ultimately nonascriptional ones. Recall that the general form of an m^{th}-level ascription is:

S_{m+1}: S_m claims of $\langle t_1', t_2', \ldots t_k' \rangle$ that $\Phi \langle t_1, t_2, \ldots t_n \rangle$,

where k is less than or equal to n and for all i less than or equal to k, there is a j less than or equal to n such that t_j is an ascription-structural anaphoric dependent of t_i' (symbolically Depends(/t_j/, /t_i'/)), and $\Phi \langle t_1, t_2, \ldots t_n \rangle$ is itself an ascriptional sentence of level $m - 1$. It is clear from this that it suffices to reduce the ascriptional complexity of the ascriptions involved, first, to trade *de re* ascriptions for *de dicto* ones plus undertaken identity commitments and, second, to trade undertaking an m^{th}-level *de dicto* ascription for attributing an $(m - 1)^{th}$-level ascription.

In this way every complex ascription is shown to correspond to a set of deontic attitudes. A converse condition holds as well. That is, starting off with any set of nonascriptional attitudes on the part of various interlocutors, it is possible to express them explicitly, from any desired point of view, by

means of iterated ascriptions, in a recursively complete fashion. The attitudes with which one starts can be any combination of undertaken and attributed identity commitments (involving singular terms) and assertional commitments (involving sentences in which those terms occur). What is being claimed is the *expressive completeness* of the regimented ascriptional idiom, over a certain domain. For consider: given a grasp of the background entailments, any interlocutor can attribute any nonascriptional assertional commitment to any other *de dicto*, with the ascriber consulting only his or her other attributions to that individual. Furthermore, given one's own undertakings of commitment (particularly to identities), one can offer *de re* versions of those ascriptions, and so indicate what, according to the ascriber, the one to whom the commitments are ascribed is talking *about.* Thus all of the first-level attitudes—that is, perspectives on the states of interlocutors—can be expressed with assertional explicitness by the use of ascriptional locutions. The undertakings of assertional commitment to these first-level ascriptional claims, however, expand the community's stock of states beyond what was present before ascriptional locutions are introduced. So these states must be explicitly ascribable in their turn, if ascriptional locutions are in fact to make possible the explicit expression of all the deontic attitudes. Applying the ascriptional expressive machinery one more time permits this, yielding second-level *de dicto* ascriptions of all of the new states generated by the first application of the machinery, and then in turn *de re* ascriptions of all of them. Repeating these two procedures inductively then permits the expression by any interlocutor of the contents of ascriptional claims of arbitrary complexity, from either the point of view of the ascriber (*de re,* inferentially expanded by commitments *undertaken* by the ascriber) or the point of view of the ascribee (*de dicto,* inferentially expanded by commitments *attributed* by the ascriber).

9

Conclusion

The meaning of words is to be determined by their use.
                    ISAAC NEWTON, Principia[1]

I. TWO CONCEPTS OF CONCEPTS

1. Three Kantian Dualisms

The semantic core of the account of discursive practice presented here is the theory of conceptual content it incorporates. The distinctive features of that theory emerge most clearly when it is contrasted with more traditional ways of thinking about concepts. The most familiar conception, one that is pervasive in contemporary philosophical thought, traces its ancestry back to Kant. Its debt to Kant is most evident in its essentially *dualistic* character: the ways in which the conceptual is contrasted with the nonconceptual. It is in this regard that traditional views of concepts differ most strikingly from the nondualist alternative endorsed here.

Kant's account begins by elaborating two of his epoch-making insights: first, that *judgments* are the fundamental form of awareness, so that concept use must be understood in terms of the contribution it makes to judging; second, that cognition and action are distinguished from their analogs in nonrational beings by their liability to certain sorts of *normative* assessment (see Chapter 1). Kant combined these insights with a *classificatory* theory of concepts, in terms of which he synthesized the teachings of his rationalist and empiricist predecessors. It is this aspect of his account that has been most influential in subsequent thought—becoming so much a matter of

course as to be almost invisible as a presupposition. Although it is based on important dimensions of ordinary concept use, the classificatory conception generalizes inappropriately as a result of running together substantially different phenomena.

For Kant, concepts provide only one of the two elements required for judgment. Concepts without intuitions are empty, and intuitions without concepts are blind.[2] Kant's theory is essentially dualistic in that the notion of the conceptual element in judgments is that of one of a *pair* of contrasting aspects. That neither is intelligible apart from its collaboration with the other is one reflection of his healthy emphasis on the primacy of judgment. It remains unclear, however, how much remains of the picture of judgment as the joint product of two distinct faculties if those faculties can be understood only by abstraction, that is, in terms of their contribution to the activity of judging. Insofar as Kant's embrace of both intellectual and sensible faculties is construed as his saying "You're both right" to Leibniz and Locke, his insistence on their mutual presupposition is bound to look like the bit where he takes it back. On the other hand, insofar as sense can be made of the notion of distinct contributions to judgment made by concepts and the unconceptualized given, the nature of their collaboration seems bound to remain mysterious. What sort of 'fit' is envisaged between concepts and intuitions, in virtue of which it is *correct* (or just possible?) to apply some, but not other, concepts to the manifold of (preconscious) representations with which intuition in some sense presents the understanding? How does intuition *constrain* the application of concepts? Kant's appeal to the schematization of the concepts just moves the bump in the rug around. The capacity for judgment, for applying rules to particular instances, subsuming intuitions under concepts, is something that in the end[3] we must just accept *that* we have, without understanding just *what* we have. A distinction becomes a dualism when its components are distinguished in terms that makes their characteristic relations to one another ultimately unintelligible. (Descartes's dualism is, as always, the paradigm.)

Essential elements of Kant's dualistic conception of concepts are still with us today. They are the basis for the suspicion evinced by some (for instance Davidson) that talk of concepts inevitably commits us to a picture in which they play the role of *epistemological intermediaries*, which stand between us and the world we conceptualize and forever bring into question the very possibility of genuine cognitive access to what lies beyond them.[4] To see why such suspicions are justified, and to bring out the contrast between dualistic and nondualistic conceptions of the conceptual, it is helpful to disentangle three different sorts of contrast between the conceptual element in thought and some nonconceptual element in thought, all of which are in play in Kant's usage. Each of these contrasts represents a genuine distinction, but as these distinctions are elaborated and run together in Kant's classificatory model of concepts, each becomes a dimension of an unworkable dualism. For

Kant, concepts contrast with intuitions first as *form* to *matter*, which they structure or organize. Second, they contrast with intuitions as *general* to *particular*. Finally, they contrast with intuitions as products of *spontaneity* or intellectual activity, as opposed to products of *receptivity*.[5]

In the first, the conceptual is distinguished from the *material*, that which provides *content*, as opposed to the form (more specifically the *normative* form or rulishness), which is the contribution of concepts. In the second, the conceptual is distinguished from the particular, as what *classifies* to what is *classified*. In the third, the conceptual is distinguished from what is imposed on us from without, as what *we do* as opposed to what is *done to* or imposed on us. It is the beginning of wisdom in reading the first Critique to distinguish the roles played in various arguments by these different distinctions. It is central to Kant's account that the three contrasts (though different) all line up together. Once this is questioned, a host of alternatives to his arguments present themselves. The lines of thought developed in this work support the conclusions that:

1. there are genuine distinctions underlying the contrasts Kant points to, but
2. far from coinciding, they are each independent of and orthogonal to the others, and
3. none of them is properly understood as distinguishing the conceptual from some nonconceptual element in judgment.

What a judgment expresses or makes explicit, its content, is conceptual all the way down.

The first idea is that of concepts as organizing something else. This can take many forms. It can be claimed that what is organized is *experience*,[6] which is carved up by concepts, or alternatively lumped together by them. The material on which concepts work can be conceived of as perceptions or observations, sense data, or patterns of sensory stimulation. The concepts are supposed to be the source of *structure*, while something else provides the *content* or matter. Davidson has this picture in mind when he objects to the "scheme/content" dualism that he takes to be implicit in talk of alternative conceptual schemes.[7] Concepts contrast with the unconceptualized matter that they conceptualize, which thereby provides content to the judgments that result. The worry inevitably raised by this picture is that unless its activity is entirely unnecessary, in conceptualizing the unconceptualized the understanding that is deploying the concepts must somehow *alter* what it works on and is therefore liable to the possibility of systematically *falsifying* that matter in rendering it digestible to the intellect.[8] It should be admitted that it always remains pretty obscure what can be meant by either the form or the matter side of this opposition. (C. I. Lewis's heroic expository effort in *Mind and the World Order* is probably as clear a setting-out of this way of conceptualizing intuitions about concepts and intuitions as can be had.)

The second idea is that concepts are something *general*, something best expressed by the use of *predicates*. Along this dimension they contrast with nonconceptual *particularity*, as expressed by the use of (at least some kinds of) singular terms. The idea here is that predicates *classify* things, *say* something about them, as opposed to simply picking them out. This thought is the heir of Kant's treatment of intuitions as representations of particularity. The association of concepts with general terms rather than proper names is pervasive in the tradition. One important example is Frege,[9] for whom concepts are functions from (sequences of) singular terms to truth-values, and so essentially things that can be *true of* the objects picked out by singular terms, by contrast to those *objects*, which concepts can be true of.

Finally, the third idea is that the *conceptual* order contrasts with something like the *causal* order, which constrains it. This distinction is the heir of Kant's distinction between judgments as the joint products of the activity ('spontaneity') of the intellect and the receptivity of the senses. According to this line of thought, whatever is conceptually articulated shows the effects of the mind working on it, whereas the nonmental world that thought is largely about is not in itself conceptually articulated. Because of special features of Kant's view, he could not put this contrast in terms of concepts versus causes (since talk of causation is for him already talk that betrays traces of the activity of the concept-mongering intellect).[10] Nonetheless the tradition he inspired contrasts conceptually articulated expressions such as definite descriptions with those that are taken merely to register causal impingements—above all the uses of *demonstratives* that are so important in expressing the noninferential reports in virtue of which our concepts have empirical content. Kaplan's work is a prime example of contemporary versions of this distinction, as he worries about how to characterize the relation between the conceptual element in propositional contents, expressed by the use of predicates and descriptions, and the *nonconceptual, contextual,* or *causal* element, expressed by the use of indexicals.

In this contemporary form, Kant's distinction survives as the contrast between the *unrepeatable* character of indexical tokenings, reflecting their token-reflexive embeddedness in a causal context, and the *repeatable* concepts that are epitomized by definite descriptions.[11] Kant ran the second and third dualistic thoughts together—that is, the distinction between predicates and singular terms, on the one hand, and between repeatable and unrepeatable elements of thought, on the other, by systematically failing to distinguish between *representations of particularity* and *particular representations* (though elsewhere he is clear-headed about the distinction between representations of relations and relations among representations). In fact, however, singular terms, which represent particulars, are typically themselves as repeatable as predicates, while unrepeatable or token-reflexive indexical expressions can be of either grammatical category.

According to this broadly Kantian, dualistic, classificatory conception of

concepts, they function as epistemological intermediaries. They stand between the understanding mind and a world that is the source of their content or matter—a world composed of particulars that are grasped by means of general concepts and that imposes itself causally on a mind obliged somehow to conform itself to those causal impingements. As long as the conceptual is conceived in this way, Davidson is quite right to object to talk of conceptual schemes by means of which we render the world intelligible to and digestible by thought. But one of the lessons that ought to be drawn from the stories told here that this is not the only way to think about concepts. In particular, this broadly Kantian approach can be laid alongside another, inspired by Sellars, which avoids the dualistic understanding of each of the three distinctions that is characteristic of the Kantian one.

2. The Inferential Conception of Concepts Is Not Dualistic in Any of the These Ways

The approach developed here thinks of concepts to begin with as inferential roles (see Chapter 2). It treats a reliably differentially elicited response as *conceptually* classifying the stimulus to which it is keyed just in case that response occupies a position in an *inferentially* articulated space of claims that can be offered as, and stand in need of, reasons. In order for it to count as a conceptually contentful performance, that response must be able to serve as a premise for inferences to the applicability of further concepts. The particular content of a given concept is accordingly thought of as the content of an inferential commitment: roughly the commitment to the propriety of the inference from any of the appropriate *circumstances of application* of that concept to any of the appropriate *consequences of application* of the concept.[12] In this way even the *empirical* content some concepts have in virtue of their connection with *non*inferential circumstances of application in *perception*, and the *practical* content some concepts have in virtue of their connection with noninferential consequences of application in *action*, can be seen to be inferentially articulated.

It is essential to this inferential approach to concepts that the inferences in question are what Sellars calls *material* inferences. This is to say that their correctness involves the particular *contents* of the concepts invoked by their premises and conclusions; it is not underwritten purely by the *form* of those premises and conclusions. A paradigm is the inference from "*A* is to the East of *B*" to "*B* is to the West of *A*," whose correctness expresses part of the *content* of the concepts *East* and *West*.

The first point to notice, then, is that thinking of conceptual contents as articulated by the material inferences that determine their role in giving and asking for reasons involves no contrast between concepts as *form* and something else as *matter* or content. The inferential role, which is the conceptual role, *is* the content. If one likes, one can say that on this conception the form

of that content is inferential. But the concept itself is identified with the particular constellation of material-inferential transitions the concept is involved in. This is not a structuring of something *else* that contrasts with the concept. The inferences materially relate one concept to other concepts, not to something of another kind. Thus the first of the Kantian dualisms, contrasting the conceptual and the material, is simply not involved in the inferential conception of nonlogical concepts.

It is possible, however, to go on to erect a superstructure of formal proprieties of inference on this base of material proprieties of inference (see 2.4.2 above). This sort of inferential articulation is an essential part of the expressive role of specifically *logical* vocabulary, by means of which we make explicit to ourselves the contents of our nonlogical concepts. So a distinction between matter and form is discerned and exploited by the inferential approach, though not in a form recognizable as distinguishing a conceptual from a nonconceptual element in judgment. Indeed, the same Fregean procedure of noting invariants under substitution that gives rise to the notion of formal proprieties of inference—an inference being valid in virtue of its form with respect to some distinguished vocabulary-kind K (paradigmatically logical vocabulary) just in case it is a materially good inference and cannot be turned into one that is not good by substituting non-K for non-K vocabulary—is what makes it possible to distinguish various formal categories of subsentential expressions, such as singular terms and predicates.

Only claims can literally function as premises and conclusions in inference; so only what is expressed by *sentences* can directly have an inferential role and so be in the most basic sense inferentially articulated. This is the version of Kant's insight concerning the primacy of judgments in cognition that survives into the inferential conception of concepts. But subsentential expressions can nonetheless be conceptually articulated according to that conception—their occurrence in a sentence can have an *indirect* inferential significance. For substitution of one subsentential expression for another in a sentence can either result in preserving the goodness of inferences in which the sentence is involved or fail to preserve it. In this way subsentential expressions can be sorted into indirect inferential equivalence classes, by noting direct inferential invariances of the sentences that result from their substitution one for another. Thus the *inferential* approach to the conceptual articulation of sentences can be extended *substitutionally* to include the conceptual articulation of *sub*sentential expressions.

When this is done, the subsentential categories of singular terms and predicates can be distinguished by the different patterns of *substitution inferences* in which they are involved. In particular, singular terms are distinguished by the de jure *symmetric* significance that their occurrence in a sentence has for substitution inferences involving it. For example, if the inference from "Benjamin Franklin spoke French" to "The popularizer of lightning rods spoke French" is a good one, then so is the converse inference.

By contrast, all predicates are involved in some *asymmetric* substitution inferences. For instance, the inference from "Benjamin Franklin could dance" to "Benjamin Franklin could move," is a good one, but the converse inference need not be. On the basis of such differences in substitution-inferential behavior, the difference between the sort of conceptual role played by singular terms and that played by predicates can be characterized (see Chapter 6).

This means that the second of the Kantian dualisms, though based on a genuine categorial distinction, also does not set off concepts as conceived by the material-inferential model. There is no restriction of the conceptual to the general, as expressed by predicates, in contrast to the particularity invoked by singular terms. Singular terms have an inferential role, represented by the set of terms intersubstitutable for them, just as predicates do. The difference between them is a formal difference of symmetric versus asymmetric substitution inference. It is not a difference that involves contrasting the conceptual as embodied in predicates, which express generalities, with something else, embodied in singular terms, which express particularity. Singular terms and predicates, the particular and the general aspects of claims, are equally (though not identically) inferentially articulated, and so equally conceptually contentful. *Particularity is as much a conceptual matter as generality*, on this conception. Thus the second dualism gets no grip on the inferential rendering of conceptual contentfulness, once that account has been extended to the subsentential level by invoking the notion of substitution inferences.

The third of the Kantian dualisms contrasts the conceptual, as the product of cognitive activity, with the nonconceptual impingement on cognitive receptivity in virtue of which that cognitive activity is constrained. Outside the strictures of Kant's own system, we can think of this as the conceptual/causal contrast, in which the application of concepts is constrained by the causal order, thought of as not itself conceptually articulated. The point of contact between the conceptual order and the causal order, according to this conception, takes place in *deixis*, where something is indicated without being characterized. In grasping this conception it is helpful to focus on the use of deictic expressions in noninferential reports, such as "This is red." For it is in such reports that the world most directly imposes itself on suitably trained concept-mongers, who find themselves passively acknowledging empirically contentful commitments.

Once again, it ought not to be denied that this sort of receptivity is essential to our empirical knowledge and that it ought to be distinguished from other, more spontaneous applications of concepts, for instance in purely inferential theorizing. Yet according to the inferentialist conception, unrepeatable deictic tokenings—for instance particular uses of 'this'—are fully conceptually articulated. Indeed, were they not, they could serve no cognitive purpose. To see how occurrences of deictic tokenings are assigned an inferential significance, and so taken to be conceptually contentful, is accordingly

to see that the third of the broadly Kantian dualisms—contrasting conceptual constraint with causal constraint on the application of concepts—fails to get a grip on the inferential conception of the conceptual. Just as the idea of *inference* needed to be supplemented by that of *substitution* in order to be brought to bear on subsentential expressions, so the idea of substitution-inferential significance needs to be supplemented by that of *anaphora* in order to be brought to bear on *unrepeatable* tokenings of subsentential expressions, rather than just on their repeatable types. To take one expression to be anaphorically dependent on another is to take it as *inheriting* its substitution-inferential role from the tokening that is its anaphoric antecedent.[13] If you say, "*That* is a porcupine" and I pick up that premise and conclude, "(so) *it* is a vertebrate," the truth of the conclusion I have drawn is to be settled (according to an interpreter) by what substitutions are appropriate (according to the interpreter) for the demonstrative tokening that serves as the antecedent for my anaphorically dependent tokening. If (according to the interpreter) what the first speaker referred to by 'that'[14] is the most cunning wooden replica of a porcupine in the room, then since this identity claim is to be understood as an intersubstitution license, I have unwittingly claimed of a cunning wooden replica of a porcupine that it is a vertebrate, and what I said is false.

Anaphora permits the formation of *chains* of tokenings, anchored by antecedents that can be deictic and therefore strictly unrepeatable. *These chains of unrepeatables are themselves repeatables and play the same role in substitution inferences that sets of cotypical tokenings play for repeatable expressions such as proper names and definite descriptions.* It is by means of anaphora, then, that substitution-inferential potential can be inherited by one expression from an unrepeatable tokening. In virtue of this mechanism, unrepeatable tokenings such as uses of demonstratives become available for service as premises in inference. In this way they acquire an *inferential* significance and so can be understood as expressing a *conceptual* content. This function of anaphora is essential to the existence of deictic expressions. For without the capacity to be picked up anaphorically, and so to have some inferential significance, deictic expressions would just be noises wrought from us by exposure to things—rather than genuinely linguistic expressions that can be used to *say* something. Thus anaphora is more basic than deixis, for there can be languages that have anaphoric mechanisms but no deictic ones, while there cannot in principle be languages with deictic mechanisms but no anaphoric ones (see Chapter 7).

In any case, with anaphora available to bring deictic expressions into substitution inferences, such expressions have indirect inferential roles, and so conceptual contents. There is no contrast between expressions like definite descriptions and those like demonstratives over the issue of whether or not they are inferentially articulated and so conceptually contentful. The structure of their contents is specifically different, for the latter are involved

in substitution inferences via anaphoric chains of unrepeatable tokenings potentially of a variety of types, while the former[15] are involved in substitution inferences via sets of repeatable, because cotypical, tokenings. But this difference plays a role analogous to that between symmetric and asymmetric substitution-inferential significances in distinguishing singular terms from predicates. In neither case is a contrast underwritten between the conceptual and something else, whether particular or causally responsive. Deictic tokenings play a role in the causal order, but they are not for that reason not also conceptually articulated.

Thus *none* of the Kantian dualisms—which contrast the conceptual as formal with the material, the conceptual as general with the particular, and the conceptual as spontaneous activity with the constraint of causes—applies to the inferential conception of concepts. That conception does not involve any commitment to a dualism of conceptual scheme and something else that it structures, classifies, or is about. So Davidson's proper fastidiousness about scheme/content dualisms and epistemological intermediaries ought not to motivate a rejection of appeals to concepts as here conceived. Concepts conceived as inferential roles of expressions do not serve as epistemological intermediaries, standing between us and what is conceptualized by them. This is not because there is no causal order consisting of particulars, interaction with which supplies the material for thought. It is rather because all of these elements are themselves conceived as thoroughly conceptual, not as contrasting with the conceptual.

The conception of concepts as inferentially articulated permits a picture of thought and of the world that thought is about as *equally*, and in the favored cases *identically*, conceptually articulated. *Facts* are just true claims.[16] Facts, like other claims, are conceptually articulated by their inferential and incompatibility relations to other claims. It is a feature of the conceptual articulation of claims, and hence of facts, that they are about particular objects.[17] (Indeed, the fact that we are accustomed to saying that facts, like claims, are *about* objects, rather than that they somehow *consist* of objects, is evidence for the correctness of identifying facts with true claims.) It is these facts and the propertied and related objects they involve that are cited as stimuli by interpreters who are specifying the reliable differential responsive dispositions in which the contents of empirical contents originate. These noninferential dispositions (the locus of our empirical receptivity) accordingly do not constitute the interface between what is conceptually articulated and what is not, but merely one of the necessary conditions for a conceptually articulated grasp of a conceptually articulated world—the world consisting of everything that is the case, all the facts, and the objects they are about.

In this way a story has been told about how the three nonconceptual poles of Kant's tripartite division of the conceptual and the nonconceptual contributions to the contents of judgments ought to be incorporated within the

conceptual realm. An approach to concepts that moves beyond exclusive focus on classification to include inferential connections among concepts as essential to their identity and individuation:

1. incorporates *content* by employing a notion of *material* proprieties of *inference*,
2. incorporates *particularity* by distinguishing between the symmetric role of singular terms in *substitution* inferences and the asymmetric role of predicates in substitution inferences, and
3. incorporates the deictic *unrepeatability* by which causal context affects conceptual content by explaining how *anaphoric* chains initiated by unrepeatable tokenings function as type-repeatables in substitution inferences.

The key theoretical concepts used to characterize the articulation of conceptual roles are material *inference, substitution,* and *anaphora,* so this can be called the ISA approach to semantics.

II. NORMS AND PRACTICES

1. The Normative and the Factual

This inferential semantics is embedded in a normative pragmatics. Material proprieties of inference are understood as norms implicit in social practices that qualify as *discursive* inasmuch as they involve treating some performances as having the significance of assertions. Such inferentially articulated practices confer propositional contents on statuses, attitudes, and performances that are suitably caught up in them (since for an expression to have a certain conceptual content just is for its use to be governed by a corresponding set of norms). In this way the semantically primitive notion of material proprieties of inference is explained in the pragmatics—in the account of linguistic practice.

Such a pragmatic theory of the relations between meaning and use raises issues about the status of implicit practical norms. Does not talk of deontic statuses as instituted by social practices involve a residual dualism? When the orienting commitment to the normative character of discursive practice was first introduced and motivated, in Chapter 1, this insight of Kant's was presented in the context of a shift from a broadly Cartesian dualism of the mental and the physical to a broadly Kantian dualism of the normative and the factual. In these crude initial terms, Descartes's opposition of two kinds of descriptive properties (corresponding to ontological kinds of substances) was contrasted with a deeper opposition between descriptive and prescriptive attitudes—between attributing properties and attributing proprieties. Thus even if Kant's *semantic* dualisms have been overcome by the ISA approach, it would appear that the pragmatics in which that semantics is embedded

incorporates an overarching dualism that distinguishes the normative and the nonnormative. How should the relations between these categories be understood?

The deontic scorekeeping idiom acknowledges a distinction between normative and nonnormative claims, explained in terms of their different roles in practical reasoning, but that *distinction* does not underwrite a *dualism* of norm and fact. Indeed, looked at more carefully, neither does Kant's. (He is large; he contains multitudes.) The initial discussion of replacing one dualism with another (in Chapter 1, Section II) was only a temporary expository device, discarded in favor of a more nuanced treatment (in Section IV) once its purpose was served. For Kant, rules are the form of the normative as such. To call something 'necessary' is to say that it happens according to a rule, and everything that happens in nature, no less than everything done by humans, is subject to necessity in this sense. Concepts are rules, and concepts express natural necessity as well as moral necessity. So according to him there is strictly no nonnormative realm—no realm where concepts do not apply. Kant's fundamental innovation is best understood to consist in his employment of a normative metalanguage in specifying *both* what merely happens and what is done.

Of course he does distinguish between the realm of regularity and the realm of responsibility. This is the distinction between that to which concepts apply and those who apply concepts—between that which can acknowledge rules only implicitly by obedience (by having concepts be applicable to it) and those who can acknowledge them explicitly by the use of concepts (by applying concepts). It is only rules as explicitly acknowledged that can be both binding and *dis*obeyed,[18] and it is the capacity for such acknowledgment—acting not just according to rules but according to *conceptions* of rules—that institutes distinctively normative statuses such as duty and responsibility. The applicable distinction is accordingly not between the normative and the nonnormative but between what can adopt explicitly normative *attitudes* and what cannot. Only we discursive (that is concept-mongering) creatures can *take* ourselves and others to be bound by the norms that are our concepts.[19]

This is the idea that is followed out in the deontic scorekeeping pragmatics presented here. The idiom in which the account of discursive commitment is expressed is normative throughout. Propositional contents are understood in terms of their explanatory role in specifying proprieties of claiming, judging, and inferring—in general, in terms of the role they play in the game of giving and asking for reasons. What it is for something to state or express a fact is explained in normative terms, and what it is for something to be stated or expressed is explained in turn by appeal to that practice. So what it is to be a fact—that is, true claim—is explained in normative terms. It is explained phenomenalistically, by appeal to the practice of fact-stating, which comprises the practical attitudes of taking a performance to be the

stating of a fact and purporting to state a fact by producing a performance. In this order of explanation, normative notions such as commitment and entitlement—which articulate implicit proprieties of practice—are more fundamental than the nonnormative properties they enable discursive practitioners to express explicitly.

However, only some of the vocabulary on which conceptual content is conferred by implicitly normative discursive practice plays the expressive role of making explicit specifically normative attitudes—for instance the attribution or acknowledgment of commitments. As explained in Chapter 4, the distinctive function of normative vocabulary is to express endorsement of patterns of practical reasoning—that is, in the first-person case, reasoning that leads from doxastic to practical commitments (presystematically: from beliefs to intentions). Social practices are implicitly normative in a way that mere behavioral regularities are not. Put phenomenalistically, that is to say that what a scorekeeper or interpreter has attributed counts as a *practice* in this sense (rather than a behavioral regularity or disposition) only if it is specified in explicitly normative terms—in terms of what, according to the practice, it is *correct,* or *proper* to do, what one *ought* to do, what one becomes *committed* or *entitled* to by a certain sort of performance, and so on. The account of practical reasoning explains in deontic scorekeeping terms how words have to be used in order to mean what such terms as 'correct', 'ought', and 'committed' do. By doing that, it makes sense of the distinction between normative statuses and attitudes, on the one hand, and nonnormative states and dispositions, on the other.

Explicitly normative vocabulary can be used to make claims (for example "Bank employees are obliged to wear neckties," "One ought not to torture helpless strangers"). Those claims can be taken-true, can be put forward as, or purport to be, true. Since facts are just true claims (in the sense of what is claim*ed,* not the claim*ing* of it), this means that norm-explicitating vocabulary is in the fact-stating line of business. That is, corresponding to the distinction between normative and nonnormative vocabulary is a distinction between normative and nonnormative facts. (Indeed, this ontologically relaxed approach to facts finds nothing mysterious about negative, conditional, or modal facts, facts about the self-identity of objects, or in general facts expressed by any sort of declarative sentence at all.) In this way the normative is picked out as a subregion of the factual.

To revert to the previous point, however, this is a distinction made *within* the encompassing normative metalanguage in which the deontic scorekeeping roles characteristic of normative and nonnormative vocabulary are specified. The distinction between normative and nonnormative vocabulary, claims, and facts is itself drawn in normative terms. In this sense, the story is one in which it is norms all the way down—a Kantian story (on the pragmatic, rather than the semantic side).[20] Far from opposing one another, the realms of fact and norm mutually include one another: fact-stating talk

is explained in normative terms, and normative facts emerge as one kind of fact among others. The common deontic scorekeeping vocabulary in which both are specified and explained ensures that the distinction between normative and nonnormative facts neither evanesces nor threatens to assume the proportions of an ultimately unintelligible dualism.

2. Where Do Norms Come From?

The story told here is Kantian not only in that it is told in normative terms but also in the pride of place it gives to normative *attitudes* in explaining how we are both distinguished from and related to the non-us that surrounds us. On the one hand, such practical attitudes—taking or treating a performance as correct, attributing or acknowledging a commitment—have been appealed to in explaining our relations in perception and action to the causal order of nonnormative facts that we inhabit cognitively and practically. On the other hand, they have been appealed to in explaining where discursive norms come from—how sapience could have arisen out of the primordial nondiscursive ooze of mere sentience. For it has been claimed not just that we discursive beings are creatures of norms but also that norms are in some sense creatures of ours—specifically, that discursive deontic statuses are *instituted* by the practices that govern scorekeeping with deontic attitudes.

Norms (in the sense of normative statuses) are not objects in the causal order. Natural science, eschewing categories of social practice, will never run across *commitments* in its cataloging of the furniture of the world; they are not by themselves causally efficacious—any more than strikes or outs are in baseball. Nonetheless, according to the account presented here, there are norms, and their existence is neither supernatural nor mysterious.[21] Normative statuses are domesticated by being understood in terms of normative attitudes, which *are* in the causal order. What is causally efficacious is our practically taking or treating ourselves and each other as having commitments (acknowledging and attributing commitments)—just as what is causally efficacious is umpires and players dealing with each other in a way that can be described as taking the score to include so many strikes and outs.[22]

It must then be asked how such an apparently reductive story about norms as instituted by social practices can be understood to be compatible with an insistence on the *irreducibly* normative character of the metalanguage in which norm-instituting social practices are specified. Here is the short answer: The work done by talk of deontic statuses cannot be done by talk of deontic attitudes *actually* adopted or relinquished, nor of *regularities* exhibited by such adopting and relinquishing, nor of *dispositions* to adopt and relinquish such attitudes. Talk of deontic statuses can in general be traded in only for talk of *proprieties* governing the adoption and alteration of deontic attitudes—proprieties implicit in social scorekeeping practices.

The crucial inferential articulation of discursive commitments consists in part in the fact that unacknowledged commitments can be (taken by other scorekeepers to be) undertaken consequentially, by acknowledging commitments to claims that (according to those scorekeepers) entail them. So according to the attributions (normative attitudes) of another, my commitments (normative statuses) outrun those I acknowledge (normative attitudes). In this way the social articulation of deontic scorekeeping attitudes is essential to the inferential (and hence discursive) articulation of the contents of the commitments they address. But this social articulation of scorekeeping practice is essentially normative in force. That I acknowledge commitment to *p* does not (according to the scorekeeper) mean that I *do* or *will* acknowledge commitment to its consequence *q*, only that I *ought to*— that I am, whether I realize it or not, *committed* to *q*.

It was shown at the end of the last chapter that the contents of ordinary empirical claims—objective proprieties governing the application of concepts—are not equivalent to the contents of any claims about who is committed to what. The implicit scorekeeping attitudes expressed by this difference in explicit contents accordingly distinguish what follows from *p* from what I or anyone *takes* to follow from *p*. What follows from *p* cannot be identified with how I or anyone *actually* keeps score; it is rather to be identified with a feature of *correct* scorekeeping (for it depends on what else is true, not on what anyone *takes* to be true). Conceptual contents on this inferential conception—and so what interlocutors are really committed to by using particular expressions (performing particular speech acts)—codify *proprieties* of scorekeeping. Any scorekeeper who attributes a conceptually contentful commitment may get these wrong, just as anyone who acknowledges or otherwise acquires such a commitment may get them wrong. Talk of inferentially articulated contents is a way of talking about implicit norms governing deontic scorekeeping practice;[23] this is the cash value of the claim that conceptual contents are *conferred* by such practice. But since commitments must be individuated at least as finely as their contents, if those contents are determined only by how it is *correct* to acquire and alter deontic attitudes, the commitments themselves must be understood as instituted also by proprieties of scorekeeping, rather than by actual scorekeeping. The scorekeeping account incorporates a phenomenalist approach to norms, but it is a *normative* phenomenalism, explaining having a certain normative status in effect as being *properly* taken to have it.

At this point it can easily look as though the account of normative statuses as instituted by social practices is marching around in an unproductive circle (at best, unilluminating; at worst viciously circular and incoherent). For clearly the prior question arises once more: What is the relation between normative specifications of practices and nonnormative specifications of behavior? Actual scorekeeping, the adoption and alteration of practical normative attitudes (acknowledgments and attributions of deontic statuses),

consists of causally efficacious events and dispositions. If normative statuses could be understood as instituted by actual attitudes of acknowledging and attributing them, then the use of normative vocabulary specifying proprieties, commitments, and entitlements would straightforwardly supervene on the use of nonnormative vocabulary specifying performances and performative dispositions and regularities. If, however, as has been claimed, the institution of discursive deontic statuses should be understood rather in terms of the implicit practical *proprieties* governing such scorekeeping—not how the score is actually kept, but how according to the implicitly normative scorekeeping practices it *ought* to be kept, how scorekeepers are *obliged* or *committed* to adopt and alter their deontic attitudes rather than how they *actually* do—then the source and status of these norm-instituting proprieties of scorekeeping practice must be inquired into.

3. Interpretation

Proprieties are normative statuses—the status a performance has as *correct* or *incorrect* according to a rule or practice. This is so even when the practice whose proprieties are in question is itself a deontic scorekeeping practice. In that case what is being evaluated as proper or improper is the acquisition and alteration of deontic attitudes—that is the acknowledgment and attribution of further deontic statuses (commitments and entitlements). The (normative) phenomenalist strategy that has been pursued throughout is to understand normative statuses in terms of normative attitudes—in terms of (proprieties of) *taking* to be correct or incorrect. This strategy dictates two questions concerning proprieties of scorekeeping practice. First (apropos of phenomenalism about norms), what must one be doing in order to count as taking a community to be engaging in implicitly normative social practices—in particular in deontic-status-instituting, conceptual-content-conferring discursive scorekeeping practices? Second (apropos of its being a normative phenomenalism), what is it about the actual performances, dispositions, and regularities exhibited by an interacting group of sentient creatures that makes it correct or appropriate to adopt that attitude—to interpret their behavior by attributing those implicitly normative discursive practices?

The first question can be addressed by considering the different sorts of intentional stance that interpreters can adopt, according to the story told here. The central task of the pragmatic part of this project (the account of discursive practice) has been to introduce the model of *deontic scorekeeping.* Keeping discursive deontic score by attributing inferentially articulated deontic statuses—propositionally contentful commitments and entitlements to those commitments—is treating the one so interpreted as being in the game of giving and asking for reasons. Social practices are *linguistic* practices when interlocutors take up the *discursive scorekeeping stance* toward one another. Adopting this stance is (implicitly, or in practice) taking or treating others as

producers and consumers of propositionally contentful speech acts. Performances count as propositionally contentful in virtue of their relation to a core class of speech acts that have the pragmatic significance of *claims* or *assertions.*

Assigning this sort of significance to performances is treating them as making *explicit* the adoption of a normative status—that is, acknowledging (undertaking) a doxastic commitment by *saying* what one is committed to. Keeping discursive score on others is adopting deontic attitudes—that is, attributing discursive commitments by implicitly or in practice taking or treating another *as* committed. Such scorekeeping (and so linguistic practice generally) does *not* require that one be able *explicitly* to *attribute* deontic statuses—to *say* (assert) *that* someone is committed to the claim that *p.* The logical locutions whose expressive role is to make the adoption of such pragmatic attitudes explicit in the form of claimable contents—propositional-attitude-ascribing vocabulary such as the regimented ". . . is committed to the claim that . . ." or its vernacular correlate ". . . believes that . . ."—form an optional superstratum whose expressive role can be understood in terms of what is implicit in ground-level linguistic practice, but which is not required for, or presupposed by, such practice.

The production and consumption of speech acts of which participants in these fundamental discursive practices are capable accordingly differ as to whether the adoption of deontic attitudes (toward normative statuses) they involve is explicit or implicit. They can *explicitly acknowledge* (and so undertake) discursive commitments, in their assertional performances, but only *implicitly attribute* them, in their scorekeeping practice. Since acknowledging a commitment (the basic sort of undertaking or acquisition of that deontic status) is producing (or being disposed to produce) performances whose pragmatic significance is to make it appropriate for scorekeepers to attribute that commitment, to take someone to be a producer of speech acts is implicitly to take that practitioner to be also a consumer of them—a scorekeeper. Givers of reasons must be able to understand what it is to give a reason. As Davidson says: "One cannot be a thinker unless one is an interpreter of the speech of others."[24]

Although performances cannot be accorded the significance of speech acts without implicitly treating the performer as a discursive scorekeeper, it is possible for those who are discursive scorekeepers to attribute a derivative sort of propositionally contentful discursive status and attitude to nonlinguistic creatures. This is adopting the *simple* intentional stance of interpreting something as a simple or practical intentional system. When this stance is adopted, the interpreter keeps a simplified sort of deontic score, by attributing propositionally contentful commitments, both doxastic and practical, which the subject is taken to acknowledge implicitly in its behavior. Its performances, dispositions, and behavioral regularities can be made intelligible by attributing sample pieces of practical reasoning, in the way Dennett

has described so well. The scorekeeping involved is simplified in that adopting the simple intentional stance does not involve attributing speech acts; it does not involve even implicitly treating the system in question as itself able to keep score (attribute, not just acknowledge deontic statuses); hence it does not involve treating it as a participant in the essentially social and linguistic game of giving and asking for reasons.

Discursive scorekeepers, participants in full-blooded linguistic practices, do two sorts of things that such simple, nonlinguistic intentional systems cannot: institute deontic statuses and confer conceptual contents. On the pragmatic side, both social flavors of deontic attitude—acknowledging and attributing—are needed to institute deontic statuses; reference to practical grasp of the possibility of *attributing* them is required to make sense of what is *acknowledged* as being inferentially articulated *commitments*. In the case of simple intentional systems, that essential pragmatic ingredient is supplied only by the interpreter, rather than attributed to the one being interpreted. On the semantic side, the social-perspectival dimension of inferential articulation is required to make sense of what states, attitudes, and performances exhibit as genuinely *propositional*, which includes having objectively *representational* conceptual content (see Chapter 8). In the case of simple intentional systems, that essential semantic ingredient is supplied only by the interpreter, rather than being attributed to the one interpreted. So the intentionality attributed by adopting this sort of stance is doubly derivative. On the side of pragmatics, the socially and inferentially articulated norms are derivative from the scorekeeping practices of the interpreter. As a result, on the side of semantics, the propositional and other conceptual contents employed to measure and systematize its behavior cannot be funded out of that behavior itself.[25]

By contrast, if one attributes genuinely linguistic practices to a community—takes its members to adopt the discursive scorekeeping stance to one another, and so to accord some performances the significance of speech acts, in particular assertional ones—one thereby takes them to exhibit *original* intentionality. The social practices one interprets them as engaging in are sufficient by themselves to institute inferentially articulated deontic statuses and so to confer genuinely conceptual contents. Describing the model of inferentially articulated deontic scorekeeping social practices is specifying in detail what one must *take* the members of a community to be doing in order for it to be *talking*—giving and asking for *reasons, making* their words and performances mean something by *their taking* them to mean something[26]— that one is thereby taking them to be doing. In short, the model specifies what structure an interpretation of the activities of a community must have in order for it to count as attributing original intentionality to that community—taking it as instituting socially and inferentially articulated deontic statuses and so conferring genuinely propositional conceptual content on them. This is adopting a further sort of stance.

So the difference between derivative simple and original discursive intentionality is presented in terms of the difference between two stances or forms of interpretation—in terms of the difference between the attitudes adopted in *attributing* them. The difference between these sorts of intentionality is not that one is construed in methodologically phenomenalist terms and the other is not. In keeping with the stance stance, this account is phenomenalist about both. The difference is that what one attributes in the case of genuinely discursive intentionality is (taken to be) autonomous in a way that what one attributes in the case of simple or practical intentionality is not.[27]

4. Semantic Externalism and the Attribution of Original Intentionality

Interpreting a community as exhibiting original intentionality is taking its members to adopt the discursive scorekeeping stance toward each other. The content-conferring norms and proprieties that an interpreter who attributes discursive scorekeeping practices takes to be implicit in them have a number of important structural features. Central among them is the fact that the conceptual norms implicit in the practices attributed to a community outrun the nonnormatively specifiable behavioral discriminations members of that community are disposed to make. For this reason, conceptual norms can be understood as *objective*, and so as binding alike on all members of a discursive community, regardless of their particular attitudes. This feature of attributions of linguistic practices secures the sense in which concepts and the commitments they involve concerning appropriate circumstances and consequences of application can be understood to be *shared*, in spite of the many differences of attitude that correspond to the different scorekeeping perspectives of the discursive practitioners who keep track of each other's statuses. This normative surplus of practice (as attributed by an interpreter) over behavior (nonnormatively specified) is also what is appealed to in responding to the issue raised by the possibility of *gerrymandering* (introduced above in 1.3.5)—the problem of what privileges one of the many ways of projecting from actual applications of concepts (and regularities and dispositions regarding such performances) commitments regarding cases that have not arisen for practical adjudication.

The reason the conceptual contents conferred by the discursive scorekeeping practices a community is interpreted as engaging in can outrun the community's capacity to apply them correctly and to appreciate the correct consequences of their application is the empirical and practical *solidity* or concreteness of those practices. The assertible contents a discursive interpretation takes to be conferred by communal deontic scorekeeping practices are inferentially articulated, but they are not merely placeholders in abstract, purely formal, relational structures—hollow shells waiting to be filled up by supplying actual facts and objects that somehow 'fit' them. For the content-

conferring practices do not relate the deontic statuses that bear those contents only to other deontic statuses. Discursive practice comprises noninferential entries and exits as well, and these (according to the interpreter attributing those practices) relate contentful doxastic and practical commitments to the worldly states of affairs that properly elicit acknowledgments of those commitments and are properly elicited by such acknowledgments, respectively. Standard discursive practices—those that encompass both empirical and practical dimensions—are solid (even lumpy), in that they involve actual objects and states of affairs, as well as the deontic statuses in terms of which score is kept.[28]

In such practices, the actual causal provenance or consequences of a deontic attitude—and not just the proprieties that connect its adoption to the adoption of other deontic attitudes—can matter (according to the external intentional interpreter attributing the content-conferring practices) for the content of the status it is an attitude toward. So an interpretation of this sort takes it that what an interlocutor who performs a certain speech act is committed to thereby, according to the practices of the relevant community, can depend on how things are in the nonlinguistic world. The interpreter takes it that the solid, corporeal communal practices determine what is being talked about (whether or not any scorekeepers in the community realize it), for those practices incorporate it. And the interpreter also takes it that what is being talked about determines what it is correct to say and infer, including practically (whether or not any scorekeepers in the community attribute the right claims and consequences). Interpretations that attribute original intentionality are accordingly *semantically externalist* in Davidson's sense.[29] This is part of what was called above (8.5.6) 'tactile Fregeanism': our practice puts us in touch with facts and the concepts that articulate them—we grasp them. But what we grasp by our practice extends beyond the part we have immediate contact with (its handles, as it were); that is why what we grasp is not transparent to us, why we can be wrong even about its individuation. How the world really is determines what we have gotten a hold of; but even though for that reason we do not know all the details about it, we still genuinely grasp *it*.

In this way the proprieties governing the application of a community's concepts are in part determined (according to the interpreter) by the actual properties of and facts concerning the things the linguistic practitioners are perceiving, acting on, and so talking about—which are just features of *their* practice (according to the interpreter). How the things and properties they are talking about actually are determines the correctness of the commitments of all community members alike. They are all bound by the *same* conceptual norms, regardless of the differences in collateral commitments that make particular claims have different inferential significances for different scorekeepers. According to the practices the interpreter takes them to be engaging in, they *share* a common set of concepts, which determines how the attitudes of those who keep score on each other are answerable to the facts.

When concrete discursive practices (including perceptual reporting and intentional agency) are ascribed to a community, the states of affairs that properly noninferentially elicit the acknowledgment of doxastic commitments and those that are properly noninferentially elicited by the acknowledgment of practical commitments are specified in the interpreter's own language. For instance, in assessing the extent to which the claims made by various community members do express facts, and so are correct uses of their concepts, the interpreter compares the commitments he or she attributes to them to those the interpreter undertakes—and similarly for assessments of their reliability as perceivers and agents. Semantic externalism is perspectival externalism.

To treat those interpreted as linguistic practitioners who use particular concepts is to treat them as bound by proprieties that project beyond their actual behavior and dispositions. The interpreter uses the norms implicit in his or her own concepts in specifying how the conceptual norms that bind the community being interpreted extend beyond the practitioners' actual capacity to apply them correctly. All the resources of the interpreter's home language are available in distinguishing one such set of proprieties from another; taking the interpreted interlocutors to have bound themselves by even a slightly different set of proprieties would be offering a different interpretation, attributing a different set of practices. The general point is that while normative interpretation of a community as engaged in one set of practices rather than another is underdetermined by nonnormatively specified actual behavior, regularities of behavior, and behavioral dispositions, *relative to such an interpretation*, concepts nevertheless are objective, shared, and unambiguously projectable.

5. Sharing Inferentially Individuated Concepts

It has been acknowledged throughout this exposition that an inferential conception of concepts raises prima facie difficulties for understanding what is involved in communication between individuals with different repertoires of commitments. The inferential significances of utterances of the same sentence produced by different performers are different—even where anaphoric and indexical phenomena are not in play. For their different collateral commitments make available different auxiliary hypotheses; hence what consequential commitments the performer undertakes by producing those performances and what would entitle their utterer to them (according to the scorekeeper who attributes the collateral commitments) are different. So something special needs to be said about the sense in which interlocutors with different collateral commitments can nonetheless be said to be able to make the same claims and express the same inferentially articulated concepts. It is worth rehearsing briefly the features of the discursive scorekeeping model that are appealed to in providing such an account.

What is from many points of view the most natural way out of this difficulty is not the path taken here. The most straightforward approach would be to adopt an inegalitarian attitude toward the different inferences a concept is involved in. A privileged class of inferences would be distinguished, which are taken to be constitutive of the concept, while the rest are accorded a secondary status as turning out to be correct ways of using the concept so constituted. There is an undeniable intuitive basis for such a distinction: The inferences from "This tractor is completely green" to "This tractor is not completely red" and from "This cloth is scarlet" to "This cloth is red," for instance, have a different status from the inferences from "This tractor is completely green" to "This tractor is made by John Deere" or from "The apple in the box is a ripe Winesap" to "The apple in the box is red." The correctness of the first inference plausibly is taken to be part of the concepts *green* and *red*, while the correctness of the second sort is equally plausibly taken to be just a matter of empirical facts about John Deere tractors and ripe Winesap apples—inferences whose correctness involves the concepts *red* and *green* without in any way constituting them.

Quine, of course, argues that one way of construing the sort of concept- (or meaning-) constitutive privilege that distinguishes the first class is defective because it does not correspond to the sort of difference in the use of the words (the practical status of the inferences) that the theory behind it entails.[30] There do not seem to be any inferential connections that are unrevisable in principle, immune to being undermined by suitable empirical evidence, and so a priori for those who grasp the concepts involved. But this is not to say that no pragmatic sense can be made of the intuitive difference in status between two sorts of inferences instanced above. Sellars,[31] for instance, does not take all the materially good inferences involving a concept to be essential to it.[32] He picks out the privileged concept-constitutive inferential connections as those that support *counterfactual* reasoning, and so count as having *nomological* force. This is a real practical difference; this way of drawing the line does not fall afoul of Quine's strictures, for it by no means follows that these conceptual matters are a priori—we need to investigate the world to find out what the laws are, as for any other facts. Since the laws involved are not a priori, unrevisable, or immune to factual evidence, this is not a version of analyticity. According to this view, not only claims but concepts can be correct or incorrect, depending upon whether the inferences they incorporate correspond to actual laws.

The difference between inferential connections among concepts that are counterfactually robust and those that are not is an important one, and this fact accounts for the felt difference between the two sorts of inferences mentioned above. Nonetheless, nothing is made of it here. This is partly because the notion of nomologicality and counterfactual reasoning, important though it is in other contexts, has not been reconstructed in discursive scorekeeping terms as part of this project (though the key notion required,

that of the *incompatibility* of claims, has been given a pragmatic interpretation). But neither this nor any other way of picking out a privileged subclass of concept-constitutive inferences has been appealed to in individuating concepts here, for two other reasons.

First, mastery of a special subset of distinguished inferences (for instance, the counterfactually robust ones) is not in general sufficient for grasp of a concept. For such grasp requires that one be hooked up to the *function* that takes as its argument repertoires of concomitant commitments available as auxiliary hypotheses and yields inferential significances as its values. Carrying on a conversation involves being able to move from perspective to perspective, appreciating the significance a remark would have for various interlocutors. (More is said about this below, in connection with the representational dimension of discourse.) The effect that various auxiliary hypotheses have on the inferential significance of a claim relative to a particular doxastic context cannot be determined just from the privileged inferences it is involved in (for instance, the counterfactually robust ones), unless it is assumed that the repertoire in question contains conditionals corresponding to all the other materially good inferences (for example from the ripeness of Winesap apples to their redness). Assuming that is contrary to the spirit of this enterprise: it depends on the formalist view of inference, which sees enthymematically suppressed conditionals behind every material propriety of inference. In particular, such a view would have the consequence that communities that do not yet have the expressive resources of logical vocabulary such as the conditional were precluded for that reason from counting as employing nonlogical concepts such as *red*.

The second reason that the inegalitarian attitude toward inferences is not taken in individuating concepts is that no matter how the privilege distinguishing some supposedly uniquely concept-constitutive inferences is construed (as counterfactual robustness or otherwise), endorsement even of these privileged inferences can still vary from perspective to perspective. There can be different views about what the laws of nature are, for instance, just as there can be differences about the colors of John Deere tractors and ripe Winesap apples. Failure to agree about such large-scale empirical matters does not preclude the interlocutors from nonetheless having a hold on the same concepts. This is the 'tactile Fregeanism' that explains why people can be counted as having radically false (nomologically precluded) views that are nonetheless genuinely *about*, say, arthritis.

Thus the response to the difficulty of reconciling the possibility of genuine communication with an account that individuates concepts by inferential roles comes in two parts. The first is the social-perspectival move. It allows inferential significances to vary with doxastic perspective, while conceptual content, which determines a *function* from perspective to significance, does not. But both the perspective-relative inferential significances of potential speech acts and the perspective-independent conceptual contents that deter-

mine them (in context) are thoroughly *normative* notions—consisting in *proprieties* of discursive scorekeeping.

The crucial second part of the response is accordingly the normative-interpretive move. It distinguishes the proprieties governing *correct* use in which the concepts grasp*ed* by individuals consist, on the one hand, from the dispositions to apply concepts, make inferences, and perform speech acts, in which an individual's grasp*ing* of a concept consists, on the other—and so distinguishes concepts from conceptions of them. Talk of grasp of concepts as consisting in mastery of inferential roles does not mean that in order to count as grasping a particular concept an individual must be disposed to make or otherwise endorse in practice all the right inferences involving it. To be in the game at all, one must make *enough* of the right moves—but how much is enough is quite flexible. One of the strategies that has guided this work is a commitment to the fruitfulness of shifting theoretical attention from the Cartesian concern with the grip we have on concepts—for Descartes, in the particular form of the centrality of the notion of *certainty*, that is infallibility about the content grasped, including its individuation (so long as we access it clearly and distinctly)—to the Kantian concern with the grip concepts have on us, that is the notion of *necessity* as the bindingness of the rules (including inferential ones) that determine how it is correct to apply those concepts.

Interpreting the members of a community as engaging in discursive practices is interpreting them as binding themselves by objective, shared concepts whose proprieties of use outrun their dispositions to apply them. There is no answer that could be given in advance as to how much one must be able to get right in order to be interpreted as hooked up to one concept or another. Massive individual differences in inferential dispositions among interlocutors are compatible with interpreting them all as nonetheless governed by (answerable to) the same set of conceptual proprieties. For it is compatible with interpreting them as talking about the same objects, answering to the same set of objective facts. In this way the perspectival account of propositional contents (and so conceptual contents generally) combines the intensional and extensional approaches to communication outlined above in 7.5.

6. *Three Levels of Norms*

The normative phenomenalist methodology applies a version of the stance stance to the problem of understanding normative statuses such as the proprieties implicit in discursive scorekeeping practices. It does so by focusing on when it is appropriate to adopt a certain sort of attitude—the stance of interpreting a community as engaged in inferentially articulated deontic scorekeeping practices that confer particular conceptual contents. It has been explained what it is for an interpreter to attribute to a community discursive practices that confer objective, shared, projectable conceptual con-

tents. The question that remains is, What is it that determines when it is *appropriate* or *correct* to adopt one rather than another of these interpretations, to attribute one rather than another of those sets of discursive practices? (Recall that the corresponding question that was asked without being answered above was rather what made it appropriate to adopt any such normative interpretation at all—to attribute practices rather than mere behavior.)

Once again, the issue of the origin of the warrant for employing a normative vocabulary seems just to have been put off. Norms have been appealed to at three different interpretive levels. First, talking and thinking, grasping and applying concepts, is described in terms of inferentially articulated norms; moves in the game of giving and asking for reasons are made intelligible in terms of alterations in what one is committed and entitled to at each stage. This is a normative reconstrual of the *discursive* in terms of *deontic statuses*. Second, what it is to *take* or *treat* interlocutors in practice *as* committed or entitled, as exhibiting deontic statuses, is explained in terms of scorekeeping practices. The norms implicit in these practices govern the alteration of *deontic attitudes*. At this stage in the account, deontic statuses are understood as instituted by proprieties of scorekeeping—of systematically altering deontic attitudes and thereby assigning pragmatic significances to performances, paradigmatically the fundamental speech act of assertion. Reference to deontic statuses is made only as the objects of deontic attitudes; the only thing one can do with a commitment is to attribute it or undertake it (perhaps, but not necessarily, by acknowledging it).

The third stage applies the methodological strategy of normative phenomenalism one more time, doing for deontic attitudes what those attitudes did for deontic statuses. The focus is now on the practices of *attributing* deontic *attitudes*—interpreting a community as engaged in implicitly normative discursive practices, as keeping deontic score by attributing and acknowledging deontic statuses. The account of deontic scorekeeping on doxastic and practical commitments explains what one must interpret a community as doing in order for it to be *talking* that one is thereby taking them to be doing. More precisely, it specifies conditions on the structure of practices a theorist attributes to a community that are sufficient for community members, so interpreted, to be treating each other as exhibiting propositionally contentful doxastic and practical commitments. Thus the relation envisaged between original intentionality and the stance of the interpreter who attributes it is analogous, at a higher level, to that obtaining between deontic statuses and deontic attitudes—for in place of a direct explanation of what commitment and entitlement are, an account of what it is to *take* someone to have such a status was offered. The phenomenalist explanatory retreat from status to attitude is applied at two levels, within the interpretation and in the relation the interpretation stands in to what is interpreted.

Norms come into the story at three different places: the commitments and

entitlements community members are taken to be attributing to each other; the implicit practical proprieties of scorekeeping with attitudes, which institute those deontic statuses; and the issue of when it is appropriate or correct to interpret a community as exhibiting original intentionality, by attributing particular discursive practices of scorekeeping and attributing deontic statuses. It is normative stances all the way down.

Regularities of communal behavior and disposition specified in nonnormative terms cannot *dictate* the attribution of scorekeeping practices that institute a particular set of normative statuses and confer a particular set of propositional contents. In adopting such a stance, the interpreter takes the interlocutors being interpreted to be *committed* to keeping score according to specific patterns, associating pragmatic significances with discursive performances that correspond to the inferentially articulated contents of the doxastic and practical commitments they express. The interpreter thereby undertakes commitments to various sorts of assessments of propriety of performance of those interpreted. Such commitments on the part of the interpreter are compatible with an indefinitely large lack of fit between the norms attributed and the actual performance of those to whom they are attributed, including their performance in assessing each other. This means that the normatively specified practices attributed by a discursive interpreter are always underdetermined by nonnormatively specified actual performances and dispositions; various sets of practices could be attributed as interpretations of the same behavior. So whenever an interpreter takes a community to be engaging in scorekeeping practices whose implicit proprieties confer one set of propositional contents on the deontic statuses they institute, there will always be alternatives, other sets of contents that could be taken to determine the pragmatic significances that scorekeepers ought to associate with discursive performances. Because of this slippage between the normative and nonnormative specifications of what community members are doing, the interpreter has considerable leeway in how to interpret them.

It remains, then, to discuss the nature of the norms that govern the choice of an interpretation of a community as engaging in one set of implicitly normative, content-conferring discursive scorekeeping practices rather than another, or rather than describing their behavior exclusively in nonnormative terms. This issue is best approached by considering the relation between the discursive scorekeeping stance adopted by the members of a linguistic community (according to an interpretation), on the one hand, and the stance adopted by the interpreter who attributes implicitly normative linguistic practices governing such scorekeeping attitudes (and so original intentionality), on the other. On the face of it, one major difference between the two stances is that discursive scorekeepers take up attitudes toward other members of their *own* communities, while an interpreter who attributes original intentionality takes the members of some *other* community to be discursive

scorekeepers. This is a misleading appearance, however. The important difference between these two sorts of norm-attributing stance is of a different sort. Indeed, under the right circumstances, the difference dissolves entirely, and the two stances coalesce. This collapse of levels provides the key both to understanding the status of the concept-articulating norms implicit in our discursive practices and also to understanding ourselves as not merely *rational,* but *logical* normative creatures, as not merely *expressive,* but *self-explicating* ones.

III. WE HAVE MET THE NORMS, AND THEY ARE OURS

1. *Original Intentionality and the Explicit Discursive Scorekeeping Stance*

The relation between the attitudes of an interpreter who attributes to a community discursive practices (and hence original intentionality), on the one hand, and the proprieties of scorekeeping implicit in those practices, on the other, is modeled on the relation between the deontic attitudes of scorekeepers and the normative statuses they attribute. What the discursive scorekeeper does *implicitly* (taking or treating others, to whom speech acts and discursive commitments are attributed, as discursive scorekeepers), the attributor of original intentionality to a community does *explicitly* (ascribing discursive scorekeeping attitudes). The underlying difference between the two stances is, not the distinction between communally *external* and *internal* attitudes or interpretations, but the distinction between *explicit* and *implicit* ones. Only a creature who can make beliefs explicit—in the sense of claiming and keeping discursive score on claims—can adopt the simple intentional stance and treat another as having beliefs implicit in its intelligent behavior. Just so, only a creature who can make its attitudes toward the beliefs of others explicit—in the sense of being able to ascribe scorekeeping attributions—can adopt the explicitly discursive stance and treat others as making their beliefs explicit, and so as having original intentionality.

Discursive scorekeeping is what the members of a community must be *doing* in order for any of their performances to have the significance (for them) of *saying* something. To take them to be a community of discursive scorekeepers whose practices confer conceptual contents, an interpreter must be capable of saying what they are doing—making *explicit* the broadly inferential proprieties that are (taken to be) *implicit* in their scorekeeping practices. For those who can adopt only the basic scorekeeping stance can attribute commitments to others (even to nonlinguistic, simple intentional systems) and can also take performances to have the significance of assertions, that is of explicit acknowledgments of discursive commitments. They

thereby implicitly recognize others as scorekeepers, and hence as attributors of commitments.

But adopting the basic discursive scorekeeping stance does not require attributing specific attributions to others; it does not require keeping score on their *attributions*, as well as their *acknowledgments* of discursive commitments. In contrast, interpreting the members of a community as engaging in discursive scorekeeping practices requires attributing to them the full range of deontic attitudes: attributing particular attributions as well as particular acknowledgments. And attributions can be attributed only by being ascribed, for it is only when made explicit in the form of propositional contents that they can be embedded in one another and so iterated. Only someone who can *say* something of the form "*S* is committed to the claim that *S'* is committed to the claim that *p*" can adopt the attitude that it makes explicit.[33]

To attribute a particular conceptual content to an expression is to say something about how it is correctly used; to attribute such content to a state or status is to say something about the circumstances under which it is appropriately acquired or relinquished and the appropriate consequences of doing so. Interpreting a community as exhibiting original intentionality is taking it that the broadly inferential proprieties that articulate the conceptual contents of their expressions, performances, and states are implicit in their deontic scorekeeping practices. So one capable of adopting that interpretive stance must be able to attribute not only scorekeeping attitudes but also those implicit inferential proprieties, which relate the adoption of one scorekeeping attitude to another. Altering a deontic scorekeeping attitude is a practical doing—the sort of thing a specification of which can play the role of the conclusion of a piece of practical reasoning. So proprieties of scorekeeping can be expressed as proprieties of practical reasoning. Again, only someone who can *say* something of the form "*S* is committed to the claim that if a scorekeeper *does* attribute to *A* commitment to *p*, then the scorekeeper *should* attribute to *A* commitment to *q*" can adopt the attitude that it makes explicit.

This is to say that interpreting a community as engaging in discursive scorekeeping practices, and so as exhibiting original intentionality, requires the full expressive resources of the logical locutions whose use has been reconstructed here in scorekeeping terms. Ascriptional locutions are needed so that both essential flavors of deontic attitude can be attributed, not just adopted: *attributions* as well as acknowledgments of commitments. Sentential logical vocabulary, paradigmatically the conditional, makes it possible to attribute acknowledgment of specifically *inferential* commitments. Normative vocabulary is required so that endorsement of a pattern of *practical* reasoning can be attributed.[34] Subsentential logical vocabulary such as quantifiers and identity locutions enable the attribution of endorsements of *substitutional* commitments, and so on. The expressive power of these logical

locutions is necessary and sufficient to make possible the adoption of the *explicit* discursive scorekeeping stance.

2. Expressive Completeness and Interpretive Equilibrium

Of course it is not just a coincidence that foregoing chapters have explained how to introduce into the basic discursive scorekeeping model just the sorts of logical vocabulary needed to make explicit the various inferentially articulated proprieties implicit in that practice—the very proprieties in virtue of which the expressions, performances, and deontic statuses governed by them count as expressing or exhibiting nonlogical conceptual contents. One of the criteria of adequacy that has guided the project from the outset is that it be possible to elaborate the model of discursive practice to the point where it is characterized by just this sort of *expressive completeness*. This means that the model reconstructs the expressive resources needed to describe the model itself. By means of these logical resources, the theory of discursive practices becomes expressively available to those to whom it applies. What is required is just that the scorekeeping practices that confer conceptual contents on the fundamental sorts of explicitating vocabulary used in stating the theory and specifying the content-conferring discursive scorekeeping practices in the first place be themselves specified within the terms of the theory. The hypothetical practitioners who play the idealized Sprachspiel of giving and asking for reasons herein described can then be understood as themselves capable of *saying* what they have been supposed to be *doing:* they can make explicit the implicit practical proprieties in virtue of which they can make anything explicit at all.

Once the expressive resources of a full range of semantically and pragmatically explicitating logical vocabulary are in play, those who have mastered them can keep discursive score *explicitly,* by making *claims* about each other's doxastic, practical, and inferential commitments. They can theorize about each other's scorekeeping attitudes. The broadly inferential scorekeeping proprieties that otherwise remain implicit, in the shadows of the practical background, are brought out into the full revealing light of explicit, public, propositional awareness. Particular ascriptions of commitment and entitlement, endorsements of consequential relations among them, and acknowledgments of and failures to acknowledge deontic statuses become topics for public challenge, justification, and debate. Though all the deontic attitudes and practical inferential know-how involved in scorekeeping cannot be made explicit in the form of claims and principles at once, there is no part of that content-constitutive practice that is in principle immune from such codification—out of reach of the searchlight of explicitation. Having been all along implicitly normative beings, at this stage of expressive development we can become explicit to ourselves *as* normative beings—aware both of the sense in which we are creatures of the norms and of the sense in which they are

creatures of ours. Having been all along implicitly discursive beings, at this stage of expressive development we can become explicit to ourselves *as* discursive beings—aware both of the sense in which we are creatures of our concepts (the reasons we produce and consume) and of the sense in which they are creatures of ours.

The members of a linguistic community who adopt the explicit discursive scorekeeping stance to one another achieve thereby a kind of *interpretive equilibrium*. Each one interprets the others as engaging in just the same sort of interpretive activity, as adopting just the same sort of interpretive stance, as one does oneself. This symmetric taking of others to adopt just the same sorts of attitudes one is oneself adopting, characteristic of the discursive scorekeeping stance, contrasts markedly with the *a*symmetric relation obtaining between an interpreter who adopts the *simple* intentional stance and the nonlinguistic creature interpreted as a simple intentional system. In that case the interpreter does not take the system being interpreted to be able to do just what the interpreter is doing, namely *attributing* (as opposed to acknowledging) beliefs, intentions, and endorsement of patterns of practical reasoning. This is one of the reasons what is attributed by such interpreters deserves to be understood as a derivative sort of intentionality.

Linguistic practitioners who have not yet deployed logical vocabulary implicitly treat other interlocutors as adopting the same interpretive stance that they do—as being discursive scorekeepers. The relations between interpreter and interpreted in such basic nonlogical discursive practices are accordingly also symmetric; an interpretive equilibrium is achieved in that case as well. Their idiom is not semantically and pragmatically explicitly complete, however; they adopt attitudes they cannot make explicit as the contents of commitments that can be acknowledged by assertion. They do not attribute the sort of attitude they are adopting just by attributing propositionally contentful commitments. They can only *implicitly* treat one another as scorekeepers, by keeping score on each other.

They treat others as in the general line of business of attributing commitments (and so being scorekeepers) by treating some of their speech acts as having the force or pragmatic significance of acknowledgments of commitments. For it is a necessary condition of being able to acknowledge (and so undertake) discursive commitments in general that one can also attribute them. So the interpretive equilibrium exhibited by basic nonlogical discursive scorekeeping practices is implicit and expressively incomplete. There is still an asymmetry between the stance such scorekeepers are interpreted as adopting by one who attributes original intentionality to the community in whose practices they participate, on the one hand, and the interpretive stance adopted by the interpreter who attributes such content-conferring practices, on the other.

That gap disappears—a complete and explicit interpretive equilibrium is achieved—for a community whose members have access to the full expres-

sive resources supplied by logical vocabulary. They can adopt the *explicit* discursive stance toward one another. Each scorekeeper can explicitly take the others to be doing just what that scorekeeper is doing: attributing discursive deontic attitudes, including that very sort of attribution. Such discursive practitioners have available as topics for explicit discussion the doings that underwrite their sayings, the practices in virtue of which anything can be explicit to or for them at all, and the interpretive stance they adopt to each other.

To the Kantian dictum that judgment is the form of consciousness has been added the claim that logic is the expressive organ of *self*-consciousness. Judging has been construed here as the practical attitude of acknowledging a certain kind of inferentially articulated commitment. Logical vocabulary then supplies the expressive resources needed to make explicit—to put in judgeable form—the semantic and pragmatic bases of judgment. By its means we come to be able to talk about proprieties of *inference,* about the structures of scorekeeping attitudes within which a performance can be accorded the significance of acknowledging or undertaking a *commitment,* and about the relations between these characteristics of specifically *discursive* practice as such. The complete and explicit interpretive equilibrium exhibited by a community whose members adopt the explicit discursive stance toward one another is social self-consciousness.[35] Such a community not only *is* a *we,* its members can in the fullest sense *say* 'we'.

3. Saying 'We'

Such a community-constitutive 'we'-saying attitude is also the one adopted by those external interpreters who attribute to a community both original intentionality and the use of logical vocabulary. In the weakest sense, we treat others as among us by attributing to them, and interpreting their performances in terms of, propositionally contentful practical and doxastic commitments—that is, by adopting the simple intentional stance. In a more basic sense, we treat others as among us by taking them in addition to perform speech acts. Keeping discursive score in this fuller sense is implicitly treating them as rational scorekeeping creatures who can appreciate the inferentially articulated pragmatic significance not only of their own nonlinguistic performances but also of their claims and of the actions and speech acts of others. At the next level, explicitly keeping discursive score on the members of a community—by ascribing not only acknowledgments but attributions of propositionally contentful commitments—is attributing original intentionality. This is explicitly treating the members of a community as among us, in the sense of being rational linguistic creatures. The richest sort of 'we'-saying is then taking those others to be in addition *logical* creatures—treating them as able to adopt, toward each other and at least potentially toward us, just the attitude we are adopting toward them.

So at the highest levels of 'we'-saying, interpretive equilibrium is achieved (whether implicitly or explicitly). The interpretive stance attributed to the members of a discursive community approaches that adopted by the interpreter who attributes original intentionality to that community. Finally, the sort of scorekeeping that is—according to the interpreter outside the community—internal to and constitutive of the community being interpreted comes to coincide with the scorekeeping of the interpreter who attributes discursive practices to the members of that community. External interpretation collapses into internal scorekeeping. Thus attributing discursive practices to others is one form or another of 'we'-saying. It is recognizing them as us.[36]

This assimilation of the external to the internal interpretive point of view means that the question of what it is to interpret the members of a community as engaged in discursive practices—what it is in this fundamental sense to say 'we' to them—has been answered by showing how the deontic scorekeeping model can be elaborated so as to make available the expressive power of logical locutions (in particular ascriptions, conditionals, and normative vocabulary). The next question dictated by the methodological strategy of normative phenomenalism about discursive norms is then, When is it *proper* or *appropriate* to adopt such an interpretive stance? When is it appropriate to say 'we' in the sense of making what others do intelligible as the acknowledging and attributing of propositionally contentful doxastic and practical commitments? When is it appropriate to *interpret* their antics, as we do for each other, rather than merely to *explain* them, as we do for nonsapients?

The collapse of the external explicit discursive interpretive stance into scorekeeping within our own expressively sophisticated practices transforms this from an abstract theoretical question into a concrete question about our own practices. Understood that way, the proper answer would seem to be latitudinarian (as suggested in the opening paragraphs of Chapter 1): one *ought* to adopt the discursive scorekeeping stance whenever one *can* adopt it. For on the one hand, the detailed requirements one must satisfy in order to count as adopting such an interpretation are stringent. Not just any group of interacting organisms can be made out to be attributing to each other commitments whose inferential and social articulation suffices to confer genuinely propositional contents on their performances. So there is little danger of such a generous policy leading to the facile or promiscuous extension of the franchise of sapience to those undeserving of it. And on the other hand, the rewards for adopting the discursive scorekeeping attitude wherever it is possible are great. Conversation is the great good for discursive creatures. Extending it increases our access to information, our knowledge, and our understanding—our semantic and pragmatic self-consciousness. Those who *can* be understood as fellow strugglers in the enterprise of making it explicit *should* be so understood.

Adopting such an inclusive demarcational attitude is saying 'we' to whoever can be understood as adopting demarcating practical attitudes—as them-

selves distinguishing by their scorekeeping a 'we' of rational agents and knowers, inhabiting a normative space of giving and asking for reasons, from an 'it' that comprises what does not live and move and have its being in such a space. Establishing entitlement to such a commitment with respect to demarcation in general would not, however, resolve the more specific issue of the status of the norms that govern the selection of one particular discursive interpretation rather than another. For the underdetermination of normative interpretation by behavior and dispositions specified in nonnormative terms means that whenever what a community does supports an interpretation of its members as engaging in discursive practices in which one set of conceptual norms is implicit, that behavior also supports rival gerrymandered interpretations of them as engaging in discursive practices in which different sets of conceptual norms are implicit. When it is possible to offer some such interpretation, how is it settled which one is most appropriate?

4. Semantic Externalism Begins at Home

The previous issue was a global one, concerning the propriety of attributing discursive scorekeeping practices at all. The present issue is a local one: assuming the global question settled in the affirmative, what is involved in choosing among various specific alternatives? Deciding to treat each of the members of some alien community as one of us (in the sense of treating them as adopting deontic attitudes, attributing and acknowledging propositionally contentful commitments) by no means settles what those contents and commitments should be taken to be. Their speech acts will typically differ in their nonnormative characteristics; they will utter different noises, make different marks (or, for all that it matters to the abstract scorekeeping model of discursive practice, turn colors, emit odors, shift voltages). What about the conceptual contents they express? What the contents of their commitments and expressions are depends on their inferential practices and on the noninferential perceptual circumstances of application and practical consequences of application implicit in their scorekeeping practices. These may differ from ours in a myriad of details and still be intelligibly interpretable. How radically different might they be?

Both the question of what makes a better discursive interpretation and the question of how different from ours the practices of the others might be taken to be before it becomes impossible to offer an intelligible interpretation of them as in the same discursive line of work as we are—as scorekeeping by changes of deontic attitude with the right social and inferential structure to confer propositional content—are questions that can be addressed only by appeal to our actual practices of interpretation in conversation. Because in a community with sufficient expressive resources the tasks of external discursive interpretation and of internal communicative interpretation are tasks of the same kind, looking at the dynamics of intralinguistic interpretation in

ordinary conversation reveals the essential features that determine also the dynamics of interlinguistic interpretation.[37] This is to say that there is no usefully general answer to the more specific interpretive question. The coalescence of external and internal discursive interpretation dictates a regress to the background language, to *our* discursive practices. The norms that determine the propriety of choices as to which discursive practices, and so which implicit conceptual norms, to attribute to those we take to be talkers are not available in advance as a set of explicit principles. They are implicit in the particular practices by which we understand one another in ordinary conversation.

The question the interpreter faces is to determine what discursive norms the members of a community have instituted, what conceptual contents they have conferred, by their linguistic practices and deontic attitudes. According to the scorekeeping model, two sorts of attribution are involved in such interpretation. The concepts according to which the truth of their claims and the success of their actions (and so their reliability as empirical reporters and practical agents) should ultimately be assessed are the ones they have committed themselves to (a matter of deontic status) by their dispositions to acknowledge some commitments in their linguistic and nonlinguistic behavior (a matter of deontic attitude). According to the interpreter, the conceptual contents practitioners have bound themselves by can outrun their discriminative dispositions to acknowledge their commitments. For this reason, objective, shared concepts can be understood as projecting beyond the dispositions to apply them of those whose concepts they are. The collapse of external interpretation into internal scorekeeping shows that this semantic externalism is just a special case of the sort of perspectival scorekeeping that has been in play all along: the commitments a scorekeeper *attributes* to someone outrun those that individual *acknowledges*. In acknowledging one discursive commitment, one is in general undertaking others, whether or not one knows what they are. This is the pragmatic (scorekeeping) significance of the *inferential* articulation of their semantic contents.

So the job of an external attributor of linguistic practices is just a special case of the job any discursive scorekeeper has: each must keep two sets of books, distinguishing and correlating the commitments interlocutors are disposed to acknowledge by overt performances, on the one hand, and those they undertake thereby, on the other. These correspond to two ways of specifying the contents of their claims—those made explicit in *de dicto* and *de re* ascriptions, respectively. For recall that *de dicto* specifications extract inferential consequences only with respect to auxiliary hypotheses (including those inferential commitments that would be made propositionally explicit in the form of conditional claims) the ascriptional target acknowledges as collateral commitments. *De re* specifications extract those consequences by appealing to auxiliary hypotheses (including inferential ones codifiable as conditionals) that are (according to the ascriber) *true*. Ordinary intralinguistic

communication—the ability to carry on a conversation across the most or-
dinary differences in doxastic perspective—requires that scorekeepers be able
to move back and forth between these two sorts of specifications of the
contents of the commitments they attribute. The contents of the commit-
ments it is appropriate to attribute to another depend both on the commit-
ment-acknowledging performances (both linguistic and nonlinguistic) the
ascriptional target is disposed to perform and on how things actually are with
the objects being talked about. Mastering our practices of attributing concep-
tually contentful commitments is learning how in particular cases to adjudi-
cate the claims of these two sources of content. Semantic externalism—the
way in which what we mean depends on how things actually are, whether
we know how they are or not—is a feature of the perspectival character of
propositional content.

So semantic (perspectival) externalism begins at home. The contents of
the commitments attributed to others, the concepts they have bound them-
selves by, cannot be specified apart from reference both to what they are
disposed to do and say and to what is true of what they are making claims
about. For what actually follows from what (according to a scorekeeper =
interpreter) depends on the facts (according to that scorekeeper = interpreter).
The point that matters here is that once the task of external interpretation
is recognized as a special case of internal interpretation (scorekeeping), the
practical norms that govern the attribution of one set of conceptually con-
tentful commitments rather than another can be recognized as just one more
instance of deciding what others of us are talking about and what they are
saying about it. Our norms for conducting ordinary conversations among
ourselves are the ones we use in assessing interpretations. There is never any
final answer as to what is correct; everything, including our assessments of
such correctness, is itself a subject for conversation and further assessment,
challenge, defense, and correction. The only answer to the question of what
makes one interpretation better than another is what makes one conversa-
tion better than another. The answer is a matter of our practical norms of
understanding one another here at home.

So the norms governing the use of the home idiom determine how to
project the concepts used to specify the content of the stranger's attitudes
(which determine how it would be proper to apply those same concepts in
novel situations) in the same way they do for the ascriber's own remarks.
This is so even in the case where the stranger is best made intelligible by
attributing concepts that differ from those used in the home community.
Thus the collapse of external into internal interpretation means that the
problem caused by the existence of gerrymandered alternatives to any par-
ticular discursive interpretation of another community from the outside is
displaced to the context of interpretation and projection within our own
community. This regress to our own interpretive practices dissolves, rather
than solves, the gerrymandering problem concerning the relation between

regularities and norms. For there is no general problem about how, from *within* a set of implicitly normative discursive practices, what we do and how the world is can be understood to determine what it would be correct to say in various counterfactual situations—what we have committed ourselves to saying, whether we are in a position to get it right or not. The account of the use of *de re* ascriptions of propositional attitude shows explicitly just what is involved in such a determination.

For our own practices come to us with the norms in; we do not just utter noises, we undertake commitments, adopt normative statuses, make pragmatically significant moves in the game of giving and asking for reasons. That there is a vocabulary, for instance any nonnormative one, that does not have sufficient expressive power to make it possible to specify our practices, make the distinctions we make, project in the way we do, has, from within our practices, no particular significance. We are always already inside the game of giving and asking for reasons. We inhabit a normative space, and it is from within those implicitly normative practices that we frame our questions, interpret each other, and assess proprieties of the application of concepts.

The account being offered is embodied in the trajectory described by attempts to answer the question, Where are the norms?

The normative first appears in the story in the guise of deontic statuses, of commitments and entitlements. Thought and talk are presented as structures of commitments and entitlements, with particular expressions having the conceptual contents they do because of the role they play in an inferentially articulated structure of such deontic statuses.

Talk of deontic statuses is then traded in, however, for talk of the attitudes of taking or treating people as committed or entitled. Deontic statuses are revealed as scorekeeping devices used for identifying and individuating deontic attitudes. In this sense the first set of norms turns out to be in the eyes of their beholders. This does not amount to a reduction of the normative to the nonnormative, however, because not only actual attitudes, acknowledgments and attributions of deontic status, but also practical *proprieties* governing the adoption and alteration of such attitudes are invoked in explaining the institution of deontic statuses by discursive scorekeeping practices.

At the next stage, these proprieties themselves are removed to the eye of the discursive interpreter, who takes a community to exhibit original intentionality by attributing to it discursive practices socially and inferentially articulated in such a way as to confer propositional contents. Once again, however, it is not only the *actual* attitudes adopted by external interpreters that must be considered but also *proprieties* governing the adoption of the discursive stance and commitment to a particular interpretation.

With the collapse of external into internal interpretation—its revelation as a special case of the sort of interpretation that goes on all the time within the practices of a discursive community—those proprieties are assimilated

to the ordinary scorekeeping proprieties in play in *our own* discursive practices. The norms turn out to be . . . here.

5. Making It Explicit

So the theoretical attempt to track down the 'source' of the normative dimension in discourse leads us right back to our own implicitly normative practices. The structure of those practices can be elucidated, but always from within normative space, from within our normative practices of giving and asking for reasons. That is the project that has been pursued in this work. Its aim is not reductive but expressive: making explicit the implicit structure characteristic of discursive practice as such.

The irreducibly normative pragmatics (theory of social practice) presented here is elaborated in terms of the basic deontic statuses of *commitment* and *entitlement* to commitments, and the essentially perspectival scorekeeping attitudes of *attributing* and *acknowledging* those deontic statuses. The semantics, or theory of the sorts of conceptual content that can be conferred by such deontic scorekeeping practices, takes the form of an account of the *inferential, substitutional,* and *anaphoric* articulation that distinguishes specifically *discursive* commitments. The result is a use theory of meaning—a specification of the social-functional roles that doxastic and practical commitments and the speech acts that express them must play in order to qualify as semantically contentful. The sorts of content addressed are those traditionally grouped together under the heading of 'intentionality'. Saying what pragmatic scorekeeping significance speech acts must have to count as assertions makes it possible to explain *propositional* contentfulness in turn as what can in that sense be made explicit—as what can in the first instance be *said* (as well as believed or meant or done). *Empirical* and *practical* contributions to such propositional (assertable, and so believable) contents are explained in terms of their conceptually articulated incorporation of the appropriate causal antecedents (in perception) and consequents (in action) of acknowledgments of discursive commitments. The *representational* dimension of propositional contents is explicated in terms of the social-perspectival character of discursive scorekeeping and the substitutional substructure of its inferential articulation. In this way it is possible to understand what is involved in assessments of judgments as *objectively* true or false—as correct or incorrect in a sense that answers to the properties and relations of the objects they are *about,* rather than to the attitudes of any or all of the members of the community of concept users.

One of the leading ideas of this enterprise is that developing an account of how semantics is rooted in pragmatics (meaning in use, content in social-functional role) is an exercise not only in the philosophy of language and the philosophy of mind but also in the philosophy of *logic*. Discursive practice is understood in terms of reasoning and representing, but above all in terms

of *expressing*—the activity of making it explicit. The expressive role distinctive of logical vocabulary is its use in making explicit the fundamental semantic and pragmatic structures of discursive practice, and hence of explicitness and expression. Pursuing the ideal of expressive completeness requires working out an account of the practices of using various particular logical locutions—paradigmatically those used to express inferential, substitutional, and anaphoric commitments and those used to ascribe discursive commitments to others.

In the end, though, this expressive account of language, mind, and logic is an account of who *we* are. For it is an account of the sort of thing that constitutes itself as an expressive being—as a creature who makes explicit, and who makes itself explicit. We are sapients: rational, expressive—that is, discursive—beings. But we are more than rational expressive beings. We are also *logical, self*-expressive beings. We not only make *it* explicit, we make *ourselves* explicit *as* making it explicit.

Abbreviations

Notes

Index

Abbreviations

The following abbreviations appear frequently throughout the Notes:

BGS   Gottlob Frege. *Begriffsschrift.* 1879.

"EPM"   Wilfrid S. Sellars. "Empiricism and the Philosophy of Mind." Reprinted in *Science, Perception, and Reality.* London: Routledge and Kegan Paul, 1963.

FPL   Michael Dummett, *Frege's Philosophy of Language.* New York: Harper and Row, 1973.

GL   Gottlob Frege. *Die Grundlagen der Arithmetik.* 1884. English translation, *Foundations of Arithmetic,* by J. L. Austin. Evanston, Ill.: Northwestern University Press, 1959.

PI   Ludwig Wittgenstein. *Philosophical Investigations.* Translated by G. E. M. Anscombe. Oxford: Basil Blackwell, 1953.

PPPW   J. Sicha, ed. *Pure Pragmatics and Possible Worlds: The Early Essays of Wilfrid Sellars.* Reseda, Calif.: Ridgeview Publishing, 1980.

PW   H. Hermes, F. Kambartel, and F. Kaulbach, eds. *Gottlob Frege: Posthumous Writings.* Chicago: University of Chicago Press, 1979.

RFM   Ludwig Wittgenstein. *Remarks on the Foundations of Mathematics.* Edited by G. H. von Wright, R. Rhees, and G. E. M. Anscombe. Translated by G. E. M. Anscombe. Cambridge: MIT Press, 1972.

"STSSD"   Robert Brandom. "Singular Terms and Sentential Sign Designs." *Philosophical Topics* 15, no. 1 (Spring 1987): 125–167.

"ÜSB"   Gottlob Frege. "Über Sinn und Bedeutung." In *Philosophical Writings of Gottlob Frege,* ed. P. Geach and M. Black, pp. 56–78. Oxford: Blackwell, 1969.

Notes

1. Toward a Normative Pragmatics

1. The particular way in which Kant understands what theoretical and practical concepts are rules for doing—namely manipulating representations either by synthesizing many under one or as determinations of the will—depends on further, independent commitments that are not here in question.

2. That is the lesson of his "Was ist Aufklärung?" Indeed, for this reason the line between Cartesian and Kantian approaches should not be drawn so sharply as to imply that Descartes had no inkling of the significance of normativity, which becomes an explicit concern for Kant. His idea of the mental as a special stuff can be seen as a response to those issues, as yet only dimly appreciated. Descartes's sense of the mental as special is precisely an inchoate awareness that its essence lies in rational, hence normative, interconnectedness. This makes it impossible to fit into what we now think of as nature, according to a conception of nature that was being formed around Descartes's time. (Thanks are due to John McDowell for emphasizing this important point.)

3. In the unpublished 1897 draft of "Logic," in *PW*, p. 147.

4. From another fragment on logic, ibid., p. 4.

5. Ibid., p. 145.

6. Ibid., p. 144.

7. Ibid., p. 145.

8. Ibid., p. 128.

9. Ibid., p. 4. Sometimes the point is put in terms of reasons, correct inference, or

justification: "Logic has a closer affinity with ethics [than psychology] . . . Here, too, we can talk of justification, and here, too, this is not simply a matter of relating what actually took place or of showing that things had to happen as they did and not in any other way" (ibid.).

10. Crispin Wright calls this Wittgenstein's "contractual" model of meaning and understanding (though for reasons that will emerge, the overtones of explicitness incorporated in this way of talking about the normative dimension are less than happy) (*Wittgenstein on the Foundations of Mathematics,* [Cambridge: Harvard University Press, 1980], p. 19). John McDowell describes it as the idea that we are "committed to certain patterns of linguistic usage by the meanings we attach to expressions" ("Wittgenstein on Following a Rule," *Synthese* 58 [1984] 325–363). As Saul Kripke puts it: "The relation of meaning and intention to future action is normative, not descriptive" (*Wittgenstein on Rules and Private Language,* [Cambridge: Harvard University Press, 1982], p. 37).

11. For example, at *PI,* § 146.

12. For example, at *RFM,* I121.

13. *PI,* § 193.

14. *RFM,* I116. One crucial difference is that the *laws* of society are explicit—they *say* what is correct and what is not. The assumption that all the laws of inference (not just those of logic) are explicit in this sense generates a regress, discussed in the next section.

15. *PI,* § 195.

16. Ibid., § 217.

17. Two sorts of norms have been pointed out as involved in attributions of intentional states. On the one hand, intentional states stand in normative relations to each other: acquiring one belief commits one to believing its inferential consequences, intending to make-true a certain claim commits one to intending the necessary means, having certain constellations of beliefs and desires can commit one to form corresponding intentions, and so on. On the other hand, intentional states stand in normative relations to states of affairs that are not intentional states: there is a certain sort of normative *accord* between a belief and the state of affairs that must obtain for it to be true, between a desire and the states of affairs that would satisfy it, between an expectation and the states of affairs that would fulfill it, between an order and the performances that would count as obeying it. The first sort of normative relation is broadly *inferential,* the second is broadly *referential.* Although Wittgenstein invokes both sorts, his primary concern is with the latter. The strategy of this work is to start with the former kind of norm and to explain the latter kind in terms of it.

18. Thus the norms incorporated in the content of a belief concern not only what other beliefs one is committed to by having that belief (and in the context of other intentional states, how one is committed to act) but also how one thereby is committed to the world's being—to be assessed by determining what objects one's belief is *about,* and what is *true of* them.

19. The more general Kantian view at stake is that concerning the normative character of concept use. The more specific view is the understanding of norms as having the form of explicit rules. The juridical idiom he employs systematically obscures the distinction between these two commitments.

20. This is a different sense from the one that Sellars, whose views are discussed below, attaches to this expression.

21. *PI*, § 201.
22. Ibid., § 84.
23. Ibid., § 198.
24. Ibid., § 201.
25. Ibid., § 289 and *RFM*, V33.
26. *PI*, § 202.
27. Ibid., § 199.
28. *Mind* 4 (1895): 278–280.
29. Gilbert Ryle, *The Concept of Mind* (New York: Barnes and Noble, Harper and Row, 1949), chap. 2.
30. *PI*, § 78.
31. What matters for the present project is the opposition between these two orders of explanation. But since they have been set out in connection with actual historical figures, as the lesson Wittgenstein has to teach Kant, it should at once be acknowledged, if only parenthetically, that when one looks at the details, Kant is somewhat better off than he appears in this sketch, for he does appreciate the point that Wittgenstein is making. Kant's acknowledgment of the possibility of a regress of rules appears in his discussion of the faculty of *judgment (Urteilskraft)*: "If understanding in general is to be viewed as the faculty of rules, judgment will be the faculty of subsuming under rules; that is, of distinguishing whether something does or does not stand under a given rule *(casus datae legis)*. General logic contains and can contain no rules for judgment . . . If it sought to give general instructions how we are to subsume under these rules, that is, to distinguish whether something does or does not come under them, that could only be by means of another rule. This in turn, for the very reason that it is a rule, again demands guidance from judgment. And thus it appears that, though understanding is capable of being instructed, and of being equipped with rules, judgment is a peculiar talent which can be practised only, and cannot be taught" (*Critique of Pure Reason*, A132/B171). The regress-of-rules argument is here explicitly acknowledged, and the conclusion drawn that there must be some more practical capacity to distinguish correct from incorrect, at least in the case of applying rules. Very little is made of this point in the first two Critiques, however. Kant's own development of this appreciation of the fundamental character of this faculty of acknowledging norms implicit in the practice of applying explicit rules, in the third Critique, has an immense significance for Hegel's pragmatism, but only his formulation of the issue seems to have influenced Wittgenstein's. The Appendix to this chapter discusses Wittgenstein's use of 'rule' in more detail.
32. P. 60 of "Realism and the New Way of Words," in *PPPW*, pp. 219–256. Another early paper that is important in this connection is "A Semantical Solution to the Mind-Body Problem" (also in *PPPW*), which argues for the paired claims (1) that mental concepts are semantic, metalinguistic concepts and (2) that semantic concepts are normative concepts.
33. "Some Reflections on Language Games," in *Science, Perception, and Reality* (London: Routledge and Kegan Paul, 1963), p. 321.
34. Ibid.
35. From "Language, Rules, Behavior," *PPPW*, p. 155. In a similar vein he says: "The mode of existence of a rule is as a generalization written in flesh and blood, or nerve and sinew, rather than in pen and ink" (from the same essay, p. 139). Talk

of rules as generalizations, even incarnate ones, is dangerous in this connection, however, for it flirts with a reductive regularism (about which more below) that identifies proprieties of practice with regularities of conduct.

36. Sellars, "Some Reflections on Language Games," p. 322.
37. *Logic*, trans. R. S. Hartman and W. Schwarz (New York: Dover Publications, 1974), p. 3.
38. Kripke, *Wittgenstein on Rules.*
39. McDowell, "Wittgenstein on Following a Rule," p. 342. It should be acknowledged that McDowell construes the structure of Wittgenstein's argument differently from the way it is presented here. He takes it that the identification of understanding with interpreting presents two unacceptable alternatives: either the regress of rules does not stop, in which case the norms evanesce, since every action is in accord with any given norm on some interpretation and fails to accord on some other, or platonistic, self-applying norms are imagined as the *last* interpretation. Here this platonistic "rails laid out to infinity" misconstrual was presented as arising independently of identifying understanding with interpreting—as a way of misunderstanding norms on a quasi-*causal* model.
40. It should be clear that to insist on this point is *not* to claim that one *cannot* explicitly say what ought to be done, say by promulgating a rule or giving an order. Nor is it to claim that where one *does* follow such an explicit rule, one must be interpreting it. Precisely not. In the typical case the understanding of what is explicit, the following of a rule, is itself practical—the exercise of implicit understanding or "know-how." One of the central tasks of this work is to say what one must be able to *do* in order to count as in this sense understanding an explicit claim, rule, or order.
41. This usage of 'discursive' is Kant's. See for instance pp. 21, 34, 82 of his *Logic.*
42. Making out this distinction is really the subject of the whole of the *Critique of Practical Reason*. The rational will is defined this way in Section 7 of Part I, Book 1, Chapter 1, p. 32 of the Akademie Textausgabe.
43. "Heidegger on Being a Person," *Nous* 16 (1982): 16.
44. This point is related to McDowell's criticism, discussed below, of social regularity theories of the sort Kripke and Wright attribute to Wittgenstein, which make the community of assessors incorrigible.
45. Although this seems the natural way to elaborate the picture, it is not evidently incoherent to imagine one organism shaping its own behavior by responding to its responses with positively and negatively reinforcing behavior. What makes such a suggestion odd is that one would think that the capacity to distinguish correct from incorrect performance that is exercised in the postulated responsive disposition to assess would also be available at the time the original performance is produced, so that no behavior-shaping ground would be gained by the two-stage procedure. But this need not be the case; the assessment might be addressed toward the performance as characterized by its consequences, discernible more readily in the event than the advent. It is no doubt more difficult to tell a story about how such self-reinforcing patterns of behavior might come about in one animal than in a group, for the behavior-shaping in question is not here, as it is in the regulist versions, *deliberate,* a matter of explicitly expressible intentions. Yet the issue of what it would be for there to be norms implicit in practice ought to be kept distinct from the issue of how such practices might in

fact plausibly arise. If the intra-organism reinforcement story is coherent, then regularity versions of the sanctions approach to implicit norms need be social only in the sense that they essentially involve the distinction of perspective between producing performances and assessing them. This contrasts with Sellars's story, in which the behavior shaping by reinforcement is deliberate and the regularity of conduct aimed at is accordingly explicitly expressible by the assessors, even though the assessors and the assessed may be time-slices of the same organisms (and full-fledged membership in the community may require playing both roles at some time). That account seems genuinely to require that there be performances where the assessing individual and the individual producing the performance being assessed are distinct. For only cases of this sort can be appealed to in making intelligible the norms implicit in grasp of a concept in such a way as to have any leverage at all against the regress-of-interpretations argument that motivates this approach. So the diachronic regulist sanctions theory sees norms as implicit in specifically *social* practice in a stronger sense.

46. For instance, McDowell "Wittgenstein on Following a Rule," p. 350: "If regularities in the verbal behaviour of an isolated individual, described in norm-free terms, do not add up to meaning, it is quite obscure how it could somehow make all the difference if there are several individuals with matching regularities."

47. Construing communal assessment regularity theories (paradigmatically those Kripke and Wright attribute to Wittgenstein) as offering an account of what it is for norms to be implicit in practice is implicitly disagreeing with one of McDowell's central criticisms of the relevance of such theories to Wittgenstein's text. McDowell objects: "The fundamental trouble is that Kripke makes nothing of Wittgenstein's concern to reject the assimilation of understanding to interpretation" (ibid., p. 343). He is right that neither Kripke nor Wright makes anything of this crucial motivating line of thought, which is rehearsed in Section II of this chapter. But he overlooks the fact that the theory they do elaborate can nonetheless be understood as an attempt to address just the considerations that are motivated by the regress-of-interpretations argument. For they can be seen as concerned to provide a notion of what it is for norms to be implicit in practice or for practice to be implicitly norm-governed rather than explicitly rule-governed. As will emerge, this repudiation is consistent with endorsement of McDowell's other criticisms of this line of thought and interpretation.

48. Kripke, *Wittgenstein on Rules*, p. 108.

49. Wright, *Wittgenstein on the Foundations of Mathematics*, p. 220. Although Wright explicitly addresses only the significance of linguistic performances, his point applies more generally to acting correctly according to one's intentional states. A communal assessment regularity theory is also put forward in the author's "Freedom and Constraint by Norms," *American Philosophical Quarterly*, April 1977, pp. 187–196, reprinted in *Hermeneutics and Praxis*, ed. R. Hollinger (Notre Dame, Ind.: University of Notre Dame Press 1985).

50. Davidson is a notable exception, taking linguistic practice and therefore intentionality to be essentially social only in the sense that it can be made intelligible only in the context of mutual interpretation—an *I-thou* relationship, in the current terminology.

51. One example of how this demand could be met by defining community membership in such a way as to preserve the distinction between those governed by

the practice and those whose practice it is they are governed by, without disjoining the groups, is provided by Sellars's account of pattern-governed practice. This is what going intergenerational does for him—the judgments of the assessors who train new community members are authoritative, and those they assess and train are the community members subject to their authority.

52. These correspond to the two sorts of objections to individual regularity or dispositional theories that Kripke *(Wittgenstein on Rules)* offers. McDowell ("Wittgenstein on Following a Rule") argues that the social regularity theory Kripke then suggests Wittgenstein endorses in response is subject to an objection of the second sort, namely that it fails to distinguish between a claim's being correct (normative status) and its being taken to be correct (normative attitude) by the community as a whole. It is argued above that this approach also falls foul by importing illicit notions of communal assessment, normative statuses such as community membership (being subject to communal authority), and expertise (exercising communal authority).

53. Some account along these lines has been a popular post-Enlightenment reading of what is being allegorically communicated by the supernatural retributive strand in Christian ethical theory.

54. Mill, in *Utilitarianism* (reprinted in *Essential Works of John Stuart Mill*, ed. Max Lerner [New York: Bantam Books, 1965], 3:215), introduces the vocabulary of internal and external sanctions, but to point to a different distinction than that intended here. His "internal" sanctions are internal to the individual (rather than to the space of norms). A paradigm would be feelings of shame or guilt.

55. *PI*, § 201.

56. Typically, though not in every case, by not letting it begin—since in the commonest cases we understand explicit claims, rules, principles, orders, and so on *without* interpreting them.

57. The theory developed in this work incorporates both of these suggestions. But at this point in the exposition no specific interpretation of either has yet been endorsed.

58. Wright *(Wittgenstein on the Foundations of Mathematics)* and Kripke *(Wittgenstein on Rules)* offer interpretations along these general lines.

59. This reading is closely related to McDowell's criticism of readings of passages such as these: "If I have exhausted the justifications I have reached bedrock, and my spade is turned. Then I am inclined to say: 'This is simply what I do'" (*PI*, § 217); "When I obey a rule, I do not choose. I obey the rule blindly" (§ 219); "How do I know [how I intend the pattern to be continued]?—If that means 'Have I reasons?' the answer is: my reasons will soon give out. And then I shall act, without reasons" (§ 211). These succumb to the temptation to conclude that, "at the level of 'bedrock' (where justifications have come to an end), there is nothing but verbal behavior" (McDowell, "Wittgenstein on Following a Rule," p. 341). That is to think of the bedrock of unreflective practice exclusively in nonnormative terms of behavioral dispositions and regularities. But as McDowell points out, one should not conclude that where justification has run out, normative assessments no longer apply. That is just what the regress-of-rules argument for the existence of norms implicit in practice shows. Wittgenstein says, in a claim important enough to appear verbatim in both *PI* and *RFM*: "To use the word without a justification does not mean to use it wrongly [*zu*

Unrecht gebrauchen]" [*RFM*, V33; *PI*, § 289]. As McDowell says: "It seems clear that the point of this is precisely to prevent the leaching out of norms from our picture of 'bedrock'—from our picture, that is, of how things are at the deepest level at which we may sensibly contemplate the place of language in the world" ("Wittgenstein on Following a Rule," p. 341).

60. P. 5, par. 2, of *On the Law of Nature and of Nations*, trans. of 1688 ed. C. H. Oldfather and W. A. Oldfather; vol. 2 in the Classics of International Law series (reprint, New York: Oceana Publications; London: Wiley & Sons, 1964).

61. Ibid., par. 3.

62. Ibid., par. 4.

63. Ibid., p. 6, par. 5.

64. One of Heidegger's central concerns in *Being and Time* (trans. J. Macquarrie and E. Robinson [New York: Harper and Row, 1963]) is to deny this characteristic Enlightenment thought, by describing how the value-free presence-at-hand *(Vorhandenheit)* studied by the physicist is abstracted from the value-laden readiness-to-hand *(Zuhandenheit)* of everyday life. Here is a characteristic statement: "In interpreting we do not so to speak, throw a 'signification' over some naked thing which is present-at-hand, we do not stick a value on it" (sec. 32, p. 190). This view is discussed in detail in the author's "Heidegger's Categories in Being and Time," *Monist* 66, no. 3 (July 1983): 387–409.

65. The evolution of physics from its "atoms in the void" conception has not appreciably altered the difficulty of fitting norms into the natural scientific world-picture. It is this difficulty that motivates both the Kantian dualism of norm and fact and the Kantian normative idealism that subordinates the latter to the former. Since the normative force of the better reason is not easily understood in terms of the sort of causal forces invoked by Newton, a normative conception of the way in which the necessity codified in laws outruns mere regularities is called in to support an understanding of causes in terms of pro-prieties governing the employment of concepts. Whatever one thinks of this heroic inversion strategy for reuniting the disparate elements of the Kantian dualism, its motivation underscores the difficulty of accommodating the norma-tive within the natural.

66. It is somewhat disingenuous to characterize his view in terms of *our* attitudes. Although his primary concern in this work is with the moral attributes insti-tuted by human beings, Pufendorf also acknowledges (as which seventeenth-cen-tury philosopher did not?) that God is also an intelligent being and can also impose or institute moral attributes. In this sense, God is treated as one of us. But even for God, in this respect primus inter pares, creation of the physical world is one thing, imposing moral attributes on it something else.

67. *Leviathan*, chap. 6, p. 24.

68. David Gauthier, *Morals by Agreement* (New York: Oxford University Press, 1986), p. 21.

69. He does not endorse a corresponding thesis for the merely prudential or instru-mental norms according to a conception of which we are also capable of acting. But he also holds that there could not be a being that had a rational will in the sense of being able to act according to a conception of a prudential rule or maxim, but did not have a rational will in the sense of being subject to moral norms. So even though not all the rules we acknowledge or act according to

conceptions of are moral rules, we can still be demarcated as the ones who act according to moral rules, for which he does endorse a version of the thesis being discussed.

70. Kant acknowledges his most immediate debt to Rousseau. (It has seemed incongruous to some that a portrait of that wild, intemperate, irregular figure should have provided the sole adornment in the study of the excruciatingly continent and excessively rule-governed Kant.) This tradition is treated as the organizing theme of the Enlightenment in Kant's "Was ist Aufklärung?"

71. Pufendorf, *Law of Nature,* chap. 2, par. 6, p. 27.

72. Ibid., chap. 5, par. 4, p. 89.

73. Ibid., par. 9, p. 95.

74. Ibid., par. 14, p. 107.

75. This is, of course, just as one would expect for an approach that takes its point of departure in construing norms from the example of explicit positive law. "So there are two parts of a law, one defining the offence, and one setting the penalty or the penal sanction; two parts, I say, and not two kinds of laws. For it is idle to say, 'Do this', if nothing follows; and it is equally absurd to say 'You will be punished', if the reason is not added, why punishment is deserved. It must, therefore, be borne in mind that the whole power of a law properly consists in its declaring what our superior wishes us to do or not to do, and what penalty awaits its violators" (ibid.). Thus the superior must have "the strength to threaten some evil against those who resist him" (par. 9, p. 95). Besides construing authority in terms of sanctions, Pufendorf also endorses two other central theses considered in the previous section of this chapter. For he takes it that by an obligation "we are bound by the necessity of doing something; for by it some moral bridle, as it were, is slipped over our liberty of action, so that we cannot rightly turn to any other quarter than that to which it directs. An obligation, however, can in no way so bind the will that it cannot, indeed, go contrary to it, although at its own peril" (chap. 5, par. 5, p. 90). Thus his conception of norms treats as essential the possibility of a distinction between what is in fact done and what ought to be done. Perhaps more remarkably, he develops his retributive picture of the practical expression of assessments by endorsing the idea of normatively *internal* sanctions: "an obligation affects the will morally, and fills its very being with such a particular sense, that it is forced of itself to weigh its own actions, and to judge itself *worthy* of some censure, unless it conforms to a prescribed rule . . . Again, an obligation differs in a special way from coercion, in that, while both ultimately point out some object of terror, the latter only shakes the will with an external force, and impels it to choose some undesired object only by the sense of an impending evil; while an obligation in addition forces a man to acknowledge of himself that the evil, which has been pointed out to the person who deviates from an announced rule, falls upon him *justly*" (ibid., emphasis added).

76. Ibid., par. 12, p. 101.

77. A contemporary version of this view—in particular of the sort of positive freedom (freedom *to* do new sorts of things, rather than freedom *from* constraint) that results from constraining oneself by specifically *linguistic* norms—is presented in the author's "Freedom and Constraint."

78. *PI,* § 258.

79. *Wittgenstein on the Foundations of Mathematics.*
80. Ibid., p. 220.
81. McDowell, "Wittgenstein on Following a Rule," pp. 333–334.
82. Most of the discussion of Dennett refers to views propounded already in his "Intentional Systems," *Journal of Philosophy* 68, no. 4 (1971): 87–106; reprinted in *Brainstorms: Philosophical Essays on Mind and Psychology* (Montgomery, Vt.: Bradford Books, 1978); page references are to the reprint edition.
83. Ibid., p. 221.
84. Davidson calls constellations of beliefs and pro-attitudes of this sort "primary reasons" for action, originally in "Actions, Reasons, and Causes," *Journal of Philosophy* 60 (1963), reprinted in *Actions and Events*, (New York: Oxford University Press, 1980), p. 4. This sort of intentional explanation is discussed in more detail in Chapter 4.
85. "Intentional Systems," p. 13.
86. Ibid., p. 17.
87. John Searle, *Intentionality* (Cambridge: Cambridge University Press, 1983).
88. Berkeley: California University Press, 1969.
89. To say that the norms implicit in practices confer conceptual content is to say that having such content just consists in being governed by those proprieties.
90. *PI*, § 54. This catalog might be taken to refer only to what he calls "definite" rules, as also in ibid., § 81.
91. As he does at ibid., §§ 224 and 225.
92. For example, at ibid., § 142.
93. At ibid., § 198, but his better wisdom may be expressed rather in the converse proposition at § 85.
94. Ibid., § 653.
95. For example, at ibid., § 237.
96. Ibid., § 199.

2. Toward an Inferential Semantics

1. This choice of terminology follows Hegel's use of *anerkennen* in his *Phenomenology.*
2. Franz Brentano, *Psychology from the Empirical Standpoint*, trans. D. B. Terrell, quoted on pp. 119–120 in *Introduction to the Philosophy of Mind: Readings from Descartes to Strawson*, ed. H. Morick (Glenview, Ill.: Scott, Foresman, 1970).
3. John Searle, *Intentionality* (Cambridge: Cambridge University Press, 1983), p. 2.
4. Robert Stalnaker, *Inquiry* (Cambridge: MIT Press, 1984), p. 2.
5. Searle, *Intentionality*, p. 17.
6. It is because of the distinction indicated by Brentano's reservation that the hyphenated phrase "object-representing" is used here, rather than more committal talk of "representing objects," to mark the categorial contrast with propositional contentfulness.
7. Cited by Roderick Chisholm, on p. 140 in *Introduction to the Philosophy of Mind*, ed. H. Morick (Glenview, Ill.: Scott Foresman, 1970).
8. "The regular connexion [*Verknüpfung*] between a sign, its sense, and its reference is of such a kind that to the sign there corresponds [*entsprechen*] a definite

sense and to that in turn a definite reference" ("ÜSB," p. 58). As an abbreviation only, Frege also allows talk of the expression, rather than the sense it expresses, designating or referring to what it represents: "To make short and exact expressions possible, let the following phraseology be established: A proper name (word, sign combination, expression) expresses its sense, stands for or designates its reference. By means of [*mit*] a sign we express its sense and designate its reference" (p. 61).

9. "Artificial Intelligence as Philosophy and as Psychology," in *Brainstorms: Philosophical Essays on Mind and Psychology* (Montgomery, Vt.: Bradford Books, 1978), p. 122.

10. It may be helpful in clearing up an incipient misunderstanding to remark here that in the official idiom to be developed and employed in this work, linguistic *expressions,* in the sense of marks and noises, do not need to be separately mentioned at this point. For it is not tokens but token*ings* that are in the first instance considered as contentful. Sign-designs, the linguistic vehicles of content, are meaningful only at one remove, in virtue of their involvement in linguistic performances that express intentional states and attitudes.

   The token/tokening distinction can often be overlooked (so that the theoretical decision as to explanatory priority alluded to here does not even arise) in the case of evanescent tokens such as utterances; the uttering/uttered ambiguity need not be resolved. The issue becomes more evident if one thinks about more permanent tokens, as when the religious enthusiast walks around the city with a sign in the shape of an arrow, inscribed "YOU are a sinner!" and points it at various passersby. In such a case the different tokenings have different contents, even though only one token is involved. The payoff of the policy mentioned here accordingly comes when the use of indexicals and other tokenings that are in principle unrepeatable becomes a topic, in Chapters 7 and 8.

11. *Critique of Pure Reason,* A97.

12. Bertrand Russell, in *The Philosophy of Logical Atomism,* ed. D. Pears (Lasalle, Ill.: Open Court, 1985), is a case in point.

13. *Critique of Pure Reason,* A69/B94.

14. Ibid.

15. Ibid., A68/B93.

16. Ibid., A69/B94.

17. Ibid., A126.

18. Ibid., A79/B104–105. The "transcendental element" introduced in this way is just reference to objects.

19. " Notes for L. Darmstädter," in *PW,* p. 253.

20. "Boole's Logical Calculus and the *Begriffsschrift,*" in *PW,* pp. 16–17.

21. The concept of substitution and its significance in such a *de*compositional semantic program are investigated in detail in Chapter 6.

22. *GL,* sec. 60. The claim that "only in the context of a proposition [*Satz*] does a name have any meaning" is enunciated also in the Introduction (p. x), and in secs. 46 and 62.

23. *PW,* p. 232.

24. "ÜSB," p. 63.

25. Ibid., pp. 57 and 58.

26. Ibid., pp. 62–63.

27. *PW*, p. 144.

28. "My Basic Logical Insights," *PW*, p. 252. See also the opening pages of the "Logic" of 1897, beginning at *PW*, p. 129.

29. This is what Kant is getting at in seeing the "transcendental element" of referring to objects as introduced into representations by their role in judgment, in the passage quoted above.

30. Desiring a mouse or desiring relief from hunger are best thought of as elliptically specified desires that one have or eat a mouse, that one's hunger be relieved. This should become apparent in the initial discussion of practical reasoning in Chapter 4. In any case, it is sufficient for the point being made here that the contents of the corresponding beliefs must be specified by sentential clauses.

31. A clue that is exploited in the account of this relation endorsed further along can be gleaned from looking at how to make explicit what a theorist becomes committed to in taking one complex object (for instance, a map) to be a representation of another (for instance, terrain). The theorist is claiming that from a certain kind of fact about the representing object (corresponding to a privileged vocabulary for describing it), it is possible to *infer* a certain kind of fact about the represented object. Thus from the fact that the blue squiggly line passes between a round dot and a square one, it is possible to infer that there is a river between a city whose population is less than 100,000 and one whose population is greater than 100,000. (This is not to say that when the representational relation is acknowledged only implicitly in the practice of someone using a complex object as a representation of another, the practitioner must be able to state explicitly the premises and conclusions of these inferences. See below at Chapter 8, Section II, Subsection 4 [such cross-references are abbreviated hereafter as 8.2.4].)

32. It is worth pointing out that this is not a difficulty that automatically confronts any theory that invokes a special ontological category of propositions in its account of claiming, judging, and believing. Obviously such accounts can accommodate the special status of propositional contents. A theory such as Stalnaker's, which understands propositions as sets of possible worlds and construes the attribution of propositionally contentful intentional states in terms of the use of the structure of possible worlds to *measure* those states for the purpose of explaining actions, is not vulnerable to the charge of semantic nominalism, of being in thrall to the model of designation. Such theories need involve no inappropriate assimilation of propositionally contentful states, attitudes, and performances to representings thought of as naming what they represent. They can respect the primacy of the propositional. They can do so precisely because they begin with the idea of an utterance expressing a proposition or a state exhibiting a propositional content. The question that then arises is what *expressing* a proposition has to do with *representing* anything. An account is required in any case of the relation between propositional contentfulness and object-representing contentfulness (purporting to represent objects). But only obfuscation results from talking in addition of sentences not only as expressing propositions and beliefs as having propositional contents but also of their *representing* propositions.

33. Although the point is put here in terms of cognition, a parallel point can be made on the side of rational action. For Kant understands the rational will as a faculty

of causally determining particular acts through the conception of a general rule (*Critique of Practical Reason*, sec. 7).

34. It is abstracted by a comparative analysis, the forerunner of Frege's substitutional or functional method of analysis of the conceptual contents of judgments, which is the concern of Chapter 6 of this work.

35. *Phenomenology*, par. 109. The erotic theory of classificatory consciousness arises in the order of exposition of the *Phenomenology* as the introduction to the theory of self-consciousness. Heidegger's successor concept of understanding in terms of taking something as something in practice is discussed in the author's "Heidegger's Categories in Being and Time," *Monist* 66, no. 3 (July 1983): 387–409.

36. To make this point is not to claim that Hegel's erotic model does not have more resources (for instance for funding a distinction between correct and incorrect taking of something to be food) than are made available in the inorganic case. Consideration of inanimate objects suffices for the contrast of interest here, however.

37. P. 262 of "Inference and Meaning," reprinted in *PPPW*.

38. This brief sketch can no more than gesture at the rich development of these ideas in the *Phenomenology*. A fuller discussion of this important chapter in the tradition inherited by the approach pursued here lies outside the scope of this work. (It will be pursued on another occasion.) The few cryptic characterizations offered here are intended to serve only as placeholders, whose significance will become somewhat clearer as the way in which material contents can be construed in terms of inference and incompatibility and expressed by means of logical vocabulary are filled in as this chapter and the rest of the work proceed.

39. In the autobiographical sketch in *Action, Knowledge, and Reality*, ed. H.-N. Castañeda (Indianapolis: Bobbs-Merrill, 1975), p. 285.

40. A detailed accounts of their efforts, understood along these lines, is offered in the author's "Leibniz and Degrees of Perception," *Journal of the History of Philosophy* 19, no. 4 (October 1981): 447–479; and "Adequacy and the Individuation of Ideas in Spinoza's Ethics," *Journal of the History of Philosophy* 14 (April 1976): 147–162.

41. *BGS*, sec. 3. Frege's word *richtig* here is usually misleadingly translated as 'valid'. The discussion below of the relation between materially and formally good inferences is intended to explain why 'correct' is a better translation here.

42. Ibid., sec. 2.

43. "Boole's Logical Calculus and the *Begriffsschrift*," *PW*, pp. 16–17.

44. *BGS*, sec. 3.

45. *Logical Syntax of Language* (London: Routledge, Kegan, Paul, 1964), p. 175. Sellars's reference is in "Inference and Meaning," *PPPW*, p. 266.

46. *FPL*, p. 432.

47. Ibid., p. 433.

48. As will become clear, the idiom of material *inference* is not to be understood in relation to the use of the so-called material *conditional*.

49. "Inference and Meaning," in *PPPW*, p. 261.

50. Reprinted in *Brainstorms* (Montgomery, Vt.: Bradford Books, 1978), pp. 10–11.

51. "Intentional Systems," p. 11.

52. Ibid.

53. "Inference and Meaning," in *PPPW*, p. 265.
54. Ibid., p. 284. This talk about the "framework" of logical transformation rules is just one expression of the attitude toward the relation between formal and material inference considered here. It would not be underwritten by the approach endorsed below, where logical vocabulary is picked out by its expressive role and then used to derive a notion of formal validity from material correctnesses of inference.
55. Ibid., pp. 270–271.
56. Ibid., p. 273.
57. Ibid., p. 274.
58. It should be noticed that the point being made here has nothing to do with the relation in mathematical logic between proof theory and model theory. In particular, it is *not* being claimed that one need be concerned only with the former, to the exclusion of the latter. The concepts of arithmetic cannot be fully specified by finitely stateable rules of inference. Nevertheless, we do grasp those concepts. But this is just to say that we do in fact understand their inferential significance. To make explicit the inferences that articulate the concepts of arithmetic, we must employ model-theoretic metalanguages. This fact in no way impugns the inferential conception of conceptual content; it merely shows that traditional proof-theoretic metalanguages are not sufficiently expressively powerful to make such inferential roles explicit. The additional (inferential) expressive power added by metalanguages that employ the traditional semantic vocabulary of truth, denotation, and satisfaction is discussed in Part 2 below.
59. "Language, Rules, and Behavior," in *PPPW*, p. 136 n. 2.
60. From "Concepts As Involving Laws, and Inconceivable without Them," in *PPPW*, p. 122. The remark of A. J. Ayer's referred to is from p. 17 of *Language, Truth, and Logic* (New York: Dover Publications, 1952).
61. "Boole's Logical Calculus and the *Begriffsschrift*," in *PW*, pp. 12–13.
62. Ibid., p. 13.
63. Ibid., p. 46.
64. *BGS*, Preface, in *From Frege to Gödel*, ed. Jean van Heijenoort (Cambridge: Harvard University Press, 1967), p. 7.
65. Ibid.
66. "Boole's Logical Calculus and the *Begriffsschrift*," in *PW*, p. 16.
67. This is a reason to reject the quasi-Tractarian view according to which nothing can count as claiming or asserting (and so nothing can count as inferring) unless the repertoire already contains logical vocabulary, so that the simplest claiming (the making explicit of anything) already presupposes the whole of logic.
68. See n. 28 above.
69. It will emerge in Chapter 3 that *entitlement*-preserving inferences are also important. They correspond roughly to inductive inferences in the same way that commitment-preserving ones correspond to deductive inferences.
70. In his fragment "Logic," Frege seems to endorse this order of explanation. He says: "To make a judgment because we are cognisant of other truths as providing a justification for it is *inferring*. There are laws governing this kind of justification, and to set up these laws of correct [*richtigen*] inference is the goal of logic . . . It would not perhaps be beside the mark to say that the laws of logic are nothing other than an unfolding of the content of the word 'true'" (*PW*, p. 3).

71. Only the sentential logical connectives are being addressed here. Identity and quantification, which raise special formal and philosophical difficulties, are discussed in later chapters.
72. 'Extensional' in this context can be made sense of in purely substitutional terms, without having to appeal to the sort of representational concepts in terms of which it is usually explicated (see 6.2 below).
73. The original investigation is in the author's "Varieties of Understanding," in *Reason and Rationality in Natural Science*, ed. N. Rescher (Lanham, Md.: University Press of America, 1985), pp. 27–51. The treatment there is cleaned up, corrected, and substantially extended in Mark Lance's "Normative Inferential Vocabulary: The Explicitation of Social Linguistic Practice" (Ph.D. diss., University of Pittsburgh, 1988), where the relevant completeness results are proven. The most interesting logical systems result from a semantics that combines pragmatically conferred incompatibility relations with pragmatically conferred entailment relations.
74. The actual procedure defines the introduction of a connective only as the principal connective in a formula and defines how to eliminate only principal occurrences. Full generality is nonetheless assured by working recursively. It should be remarked that according to the approach developed here, the standard Gentzen-style definitions for logical connectives are still possible for conjunction and disjunction, but the expressive role of conditionals, negation, and many other bits of logical vocabulary requires that they be understood as having quite another sort of introduction rule.
75. *FPL*, p. 453.
76. Noninferential reports are discussed in more detail in Chapter 4.
77. The *empirical* contribution to conceptual content made by noninferential *circumstances* of application in *perception* and the *practical* contribution to conceptual content made by noninferential *consequences* of application in *action* are discussed in Chapter 4.
78. The significance of this sort of example is explored in the author's "Truth and Assertibility," *Journal of Philosophy* #73, no. 6 (March 1976): 137–149. Ingredient contents are discussed below at 6.1.2.
79. *FPL*, p. 455; the following passage is on pp. 453–454.
80. Ibid., pp. 456–457.
81. Ibid., p. 455.
82. Ibid., p. 454. It should be noted that inferential conservativeness is a weaker condition than derivability of circumstances from consequences (or vice versa). Showing how to derive one aspect from the other, using logic or prior inferential commitments, is sufficient but not necessary for conservativeness. I am grateful to Michael Kremer for this point.
83. N. Belnap, "Tonk, Plonk, and Plink," *Analysis* 22 (1962): 130–134, commenting on A. N. Prior's "Runabout Inference Ticket," *Analysis* 21 (1960–1961): 38–39.
84. *FPL*, p. 454.
85. Jonathan Bennett suggested this illustrative anecdote.
86. *FPL*, p. 455n.
87. Ibid., p. 358.
88. In Quine's *From a Logical Point of View* (Cambridge: Harvard University Press, 1953), pp. 20–46.

89. This discussion addresses only versions of the project of offering truth conditions that envisage employing *other* concepts than those expressed by the words appearing in the sentences for which one is offering truth conditions, so as to offer *substantive* explications of those concepts. If truth conditions are conceived *modestly*, so that one is allowed to specify the truth condition for the sentence "Luther was a Schwabian" as simply as that Luther was a Schwabian, then the consequences pointed to in the text do not arise.

90. It should be acknowledged that although the discussion of this chapter has been framed throughout in terms of a stark opposition between two complementary orders of explanation—the representationalist and the inferentialist—these alternatives are not exhaustive. Other possibilities include treating neither representation nor inference as explanatorily prior to the other. One might then go on to explain both in terms of some third notion, which is treated as more fundamental. Or one might eschew reductive explanations in semantics entirely and remain contented with describing the relations among a family of mutually presupposing concepts—a family that includes representation, inference, claiming, referring, and so on.

91. Recall from the discussion in 1.4 above that the most serious objection McDowell levies against the social-practice theories of norms put into Wittgenstein's mouth by Wright and Kripke is that they have no room for the idea of proprieties of concept use that the whole community could be wrong about. As Wright puts it, these theories jettison the intuitive "ratification independence" of concept use for the special case where the ratifying attitudes of taking particular candidate applications of concepts to be correct or incorrect are those of the community as a whole.

92. The obligations involved in this order of explanation have just been indicated. It was suggested above that the corresponding explanatory demands on the contrary directions of explanation pursued by the intellectualist about norms, the formalist about logic, and the representationalist about content are difficult to meet. The intellectualist about norms has trouble explaining the norms governing the use or application of rules, principles, claims, and concepts. The formalist about logic has trouble explaining nonlogical content. The representationalist has trouble explaining specifically propositional content and its grasp. Of course there are various strategies for meeting or evading these demands. The present assembling of reminders and considerations intends only to sketch an alternative; it does not pretend to offer all-purpose refutations of the various contrary explanatory strategies that might be adopted.

93. This is an important point for Kant as well. His terminology in the *Logic*, where one-premise inferences are called "immediate" and multipremise inferences are called "mediated," greatly influences Hegel's use of those central technical terms of the *Phenomenology*.

3. Linguistic Practice and Discursive Commitment

1. The same structure is exhibited even if sentences are not taken as the primitive interpreted expressions. For example, if the basic stipulated assignment is of objects to singular terms, and sets of objects to predicates (or sets of sequences of objects, for multiplace predicates), then the results of simple syntactic predi-

cations may be assigned truth-values as derived interpretants accordingly as the objects (or sequences) corresponding to the term(s) appearing in the predication are or are not included in the sets corresponding to the predicates.

2. *FPL*, p. 413. The view Dummett expounds in this passage differs from the one to be developed in this work in that he is concerned only with the contentfulness of linguistic expressions, not with that of intentional states and attitudes more generally; he considers only truth-conditional semantic interpretation; and he does not make clear the essentially normative character of the linguistic practices that constitute the use or working of the language.

3. A paradigm of the recognition of this promissory note implicit in, for example, possible-worlds semantics is David Lewis's "Languages and Language," in *Language, Mind, and Knowledge*, Minnesota Studies in the Philosophy of Science, vol. 7, ed. Keith Gunderson (Minneapolis: University of Minnesota Press, 1975), pp. 3–35; reprinted in Lewis's *Philosophical Papers*, vol. 1 (New York: Oxford University Press, 1983), chap. 11. Lewis's account of what it is for semantic interpretants to be appropriately associated with expressions by the use of language turns on his notion of *convention*, which appeals to propositionally contentful *intentions* and beliefs and so is not suitable to be extended to an account of the pragmatics corresponding to the contentfulness of such intentional states. In *Inquiry* (Cambridge: MIT Press, 1984), Robert Stalnaker provides an account of what it is for sets of possible worlds to be associated as the propositional contents of intentional states such as belief, appealing only to the possibility of intentional interpretation of intelligent behavior.

4. H. P. Grice, "Meaning," *Philosophical Review* 66 (1957): 377–388. Also see Stephen Schiffer, *Meaning* (New York: Oxford University Press, 1972).

5. David Lewis, *Convention: A Philosophical Study* (Cambridge: Harvard University Press, 1969), and "Languages and Language"; Jonathan Bennett, *Linguistic Behavior* (Cambridge: Cambridge University Press, 1976); John Searle, *Speech Acts* (Cambridge: Cambridge University Press, 1969); and *Expression and Meaning* (Cambridge: Cambridge University Press, 1979).

6. Davidson takes this model of communication to entail that sharing a language is a merely practical, hypothetical necessity—it is convenient for members of a linguistic community to use the same noises to express the same thoughts because it minimizes the need for explicit theorizing about the intentions with which the noises are produced. The expectations or customs of other speakers, however, have no authority over how any individual is correctly understood. What matters is how the speaker intends to be understood or interpreted. Davidson's subtle position is different from the others mentioned in this connection in important ways, however, as he does not take it that the contents of these communicative intentions can be made sense of antecedently, in abstraction from interlocutors' interpretation of one another. Put another way, the intention to be interpreted one way rather than another that Davidson rightly takes to be essential to the meaningfulness of ordinary discourse can be understood to be implicit rather than propositionally explicit, an intention-in-action rather than a separately individuatable prior intention. That is, its involving such an intention can be conceived as an automatic compliment paid to a performance in virtue of the fact that were the issue to be raised and the speaker sincerely to *disavow* the intention to be understood in a certain way, that would be evidence

that it had been misinterpreted. Such a view is importantly different from the Gricean picture of meanings as imposed on utterances by the antecedently contentful intentions of speakers, although Davidson is not always careful to register the distinction.

7. Jay Rosenberg, *Linguistic Representation* (Dordrecht: D. Reidel, 1975), chap. 2.

8. This thought leads John Searle (*Intentionality* [Cambridge: Cambridge University Press, 1983]) to insist that the intentionality of intentional states, unlike that of linguistic expressions, must be *intrinsic*. That is, Searle endorses the view, characteristic of agent semantics, that "I impose Intentionality on my utterances by intentionally conferring on them certain conditions of satisfaction which are the conditions of satisfaction of certain psychological states." (p. 28). Accordingly, he is committed to a strong distinction between the contentfulness of utterances and that of the intentional states they express: "Since sentences— the sounds that come out of one's mouth or the marks one makes on paper—are, considered in one way, just objects in the world like any other objects, their capacity to represent is not intrinsic, but is derived from the Intentionality of the mind. The Intentionality of mental states, on the other hand, is not derived from some more prior forms of Intentionality but is intrinsic to the states themselves. An agent uses a sentence to make a statement or ask a question, but he does not in that way use his beliefs and desires, he simply has them" (pp. vii–viii). Thus, "That the belief has those conditions of satisfaction is not something imposed on the belief by its being used at all. A belief is intrinsically a representation in this sense: it simply consists in an Intentional content and a psychological mode . . . [It is false that] in order for there to be a representation there must be some agent who uses some entity as a representation. This is true of pictures and sentences, i.e. of derived Intentionality, but not of Intentional states" (p. 22). The doctrine that the intentionality or contentfulness of intentional states and attitudes is *intrinsic* is in some ways the correlate in the domain of philosophical semantics of the method of *stipulation* in formal semantics. The theorist takes it that intentional states are just the sort of thing that comes with an intentional (paradigmatically propositional or representational) content. This is Descartes's strategy—the mental is distinguished as being naturally *about* other things (it is the sort of thing other things have *objective* reality in). The theorist can talk about the consequences of such contentfulness but cannot be expected to have anything substantive to say about what that contentfulness consists in, apart from those consequences. Thus Searle denies that it is possible to give an analysis of intentionality, taking it rather that "Intentionality is, so to speak, a ground floor property of the mind" (pp. 14–15).

9. Robert Stalnaker, *Inquiry* (Cambridge: MIT Press, 1984), p. 4.

10. Ibid., p. 5.

11. Ibid., p. 6.

12. In spite of the acknowledgment of these alternatives implicit in his disjunctive formulation, however, Stalnaker restricts his consideration (and his arguments against the linguistic approach) to resemblance theories. Thus he concludes his description of the linguistic approach: "It is not essential to the linguistic picture that every thinking creature be capable of outward speech or that every one of our thoughts be expressible in our public language. All that is essential

is that thought be explained by analogy with speech" (ibid., p. 5). This is true of resemblance theories—but not relation theories—of the significance of linguistic practice to intentional states. His further subdivision of the linguistic picture is phrased so as to accord with the restriction of consideration to analogical theories, in which thought is understood on the model of speech. "The development of the linguistic picture leads in two quite different directions which emphasize different analogies between speech and thought. One hypothesizes a language of thought, which may be different from any language used for communication; the other argues for the dependence of thought on the social activities of speech" (pp. 5–6). But in fact thought or the possession of contentful intentional states might be taken to depend on the social activities of speech *either* because the contents of intentional states must be modeled on and so understood in terms of the contents expressed by public speech acts in general *or* because it is essential to a state's having such content that the state can issue in a speech act by which it is publicly expressed—that is, either according to a linguistic theory of intentionality structured by resemblance or according to one structured by relation.

13. "Thought and Talk," in *Inquiries into Truth and Interpretation* (New York: Oxford University Press, 1984), p. 156.

14. This and the subsequent passages are from ibid., p. 170.

15. *FPL*, p. 362.

16. Hartry Field, "Mental Representation," *Erkenntnis* 13 (1978): 9–61. This strategy was put forward by Sellars already in 1953, in his "Semantical Solution to the Mind-Body Problem" (in *PPPW*, pp. 219–256; see especially secs. IV–VI), and remained at the center of his approach thereafter.

17. *Beurtheilbarer Inhalt*, introduced in the *BGS*, sec. 2.

18. This is not to say that in the full-fledged language game any particular move that is taken to be correct by the practitioners—even *all* the practitioners—thereby counts as correct. (That this promissory note is eventually redeemed is demonstrated by the objectivity proofs presented below in 8.6.5.) One of the primary tasks of this work is to begin to explain the way in which the linguistic community can institute incompatibilities relating subpractices, which then constrain its own assessments where those subpractices interact—that is, the way in which noninferential reporting practices and deliberate actions performed as a result of practical reasoning conspire to confer *objective* empirical content on the concepts they are inferentially linked to—the way in which what is said can come to answer for its correctness not to the ones using the language but to what they use it to talk *about*.

19. (Formal) inconsistency is to (material) incompatibility as formal logical validity of inference is to material correctness of inference. In each case the former should be defined in terms of the latter, for the reasons discussed above in 2.4.4. (The official interpretation of the ^{s}scare quotes^s employed here is presented below in 8.4.5.)

20. In order to get a more realistic model, a shell might be added around this practice. A man, for example, who routinely fails to fulfill promises might be held responsible not only for failing to recognize himself as having undertaken a commitment (not recognizing his entitlement to entitle others to rely on him) but also for abusing the community practice of promising by his repeated attempts to claim or pretend to an authority that he has then withheld. Holding

such a man responsible in this way might consist in shunning him or in beating him with sticks, as in the original case. It is possible, in other words, to combine a general, externally defined sanction, with a specific, internally defined one. The former becomes a shell around the latter. The general sanction might be an instance of holding responsible for any abuse of linguistic authority, which might eventually lead to expulsion from the linguistic community (being treated like a parrot or a pariah). (This addition was suggested by Michael Kremer.)

21. The language of scorekeeping is suggested by David Lewis's "Scorekeeping in a Language Game," reprinted as Chapter 13 of his *Philosophical Papers*, pp. 233–249. More is made of this notion below, in Section IV.

22. The responsibility characteristic of action and the authority characteristic of perception are discussed in Chapter 4.

23. Mark Lance noticed that this definition permits one to consider *asymmetric* incompatibility relations as well as symmetric ones. He exploits this possibility formally in his development (in "Normative Inferential Vocabulary: The Explicitation of Social Linguistic Practice" [Ph.D. diss., University of Pittsburgh, 1988]) of the sort of incompatibility semantics originally suggested in the author's "Varieties of Understanding," in *Reason and Rationality in Natural Science*, ed. N. Rescher (Lanham, Md.: University Press of America, 1985), pp. 27–51. *Inconsistency* is the formal correlate of incompatibility. It is a logical notion, to be understood in terms of *negation*. But what makes a bit of vocabulary express negation is itself to be understood in terms of its relation to material incompatibility. The negation of a claim is defined as its minimal incompatible, the inferentially weakest claim that is entailed (in the commitment-preserving sense) by everything that is incompatible with the original claim.

24. As the examples discussed in the previous section indicate, it is not impossible to have authority without responsibility, or responsibility without authority. As students of organizational behavior will attest, however, it is a basic principle of social engineering that the stability and effectiveness of a practice are undercut if the authority accorded to some practitioners outruns their corresponding responsibilities, or vice versa. Linguistic practice as here construed is well designed in this respect.

25. Kurt Baier, "Responsibility and Freedom," in *Ethics and Society*, ed. R. T. DeGeorge (New York: Anchor Books, 1966), pp. 49–84.

26. John Searle, *Speech Acts*, p. 96.

27. Justificatory practices depend on entitlement-preserving inferences. But commitment-preserving inferences are also entitlement-preserving (though not conversely). If anyone who is committed to p is thereby committed to q, the only case in which entitlement to p plausibly would not carry with it entitlement to q is one in which the interlocutor is precluded from entitlement to q by concomitant commitment to something incompatible with it. But if p commitment-entails q, anything incompatible with q is incompatible with p, so under the circumstances described, the interlocutor could not be entitled to p.

28. The justificatory mode of entitlement inheritance requires that one invoke claims with *different* contents, for otherwise the 'stuttering' inference, from p to p, would count as a justification of p.

29. This is, as will become clear in terms of the model, compatible with the possibility in the fully developed practice of an interpreter correctly taking the entire community to be wrong about what commitments they are entitled to—but

such a judgment will always be that they are in some sense wrong by their own lights, that is, wrong given how they have committed themselves to its being proper to settle such questions and assess the answers.

30. Lewis, "Scorekeeping in a Language Game," p. 236.
31. Ibid.
32. Ibid., p. 237.
33. Ibid., p. 238.
34. Chapter 9 (Sections II and III) discusses the essential role played in such objective constraints by the fact that the interpreter who attributes discursive deontic scorekeeping practices to a community can use nonscorekeeping vocabulary with an antecedent use in specifying those practices.
35. Lewis, "Scorekeeping in a Language Game," p. 239.
36. *FPL*, p. 361.
37. Ibid.
38. "To begin with" because of the contribution made to semantic contents by the role sentences play in noninferentially elicited (but inferentially articulated) observation reports and in their role in giving rise to actions.
39. It is assumed throughout (though this requirement could be relaxed) that incompatibility is a symmetric relation. Also, if it is assumed that commitment-preserving inferences are entitlement preserving, in the absence of incompatible defeasors it follows that if everything incompatible with q is incompatible with p, then the inference from p to q is good both committively and permissively.
40. The correctness of such an inference according to A depends not only on the commitments and entitlements to commitments that A attributes to B but also on the commitments that A undertakes. When the expressive resources are available for explicit challenges to reliability inferences, these background premises of A will be cited in justification of those inferences and their conclusions. Under those circumstances they will be cited as constituting *standard conditions* for A's observational authority with respect to this sort of content. This point is adverted to in the text further along.
41. Such a mapping directly characterizes what corresponds to the appropriate consequences of application of the expression. The circumstances of appropriate assertion (according to the scorekeeper in question) can be recovered from the full mapping, however, as the subset of initial scores in which the consequences of assertion include attribution, not only of commitment, but of entitlement to the assertion (according to the scorekeeper in question). The way in which entitlement is attributed to noninferential reports (thereby treating the reporter as reliable) shows that this class of appropriate circumstances of assertion is wider than what would result from taking the assertion to be in order only when the prior score already included an attribution of entitlement.
42. Philip Kremer and Mark Lance present some fascinating results concerning the explicit codification of commitment consequences in the form of logical conditionals in "The Logical Structure of Linguistic Commitment, I: Four Systems of Non-Relevant Commitment Entailment" and "The Logical Structure of Linguistic Commitment, II: Relevant Commitment Entailment," both forthcoming in the *Journal of Philosophical Logic.*
43. If one is entitled to p and p commitment-entails q, one is entitled to q—any entitlement-defeating incompatibilities to q equally defeat entitlement to p.

44. One can be (taken to be) entitled to claims one is not (taken to be) committed to—these are conclusions one is entitled to draw but has not yet committed oneself to. In this way one may be entitled to each of two mutually incompatible claims, so long as neither has been endorsed and commitment to it undertaken. So one might have good inductive reasons for believing that the barn is on fire (smoke, the particular noises that would usually accompany it, and so on) and a different set of good inductive reasons for believing it is not on fire (the alarm has not rung, it is pouring rain, the barn was just inspected, and so on). Either conclusion by itself could be defended, though one would cease to be entitled to it if already committed to the conclusion of the other argument. Attribution of entitlement to the consequences of conjoining incompatible contents to which one is severally entitled (because not committed to either) is avoided by closing entitlements without commitment only under committive inferences (if not defeated by incompatibilities) and closing only entitlements to commitments actually undertaken (according to the one keeping score) under permissive inferences.

45. Under suitable conditions. Considerations having to do with the possibility of coercion, insincerity, shyness, and so on are systematically suppressed in presenting the model of assertional practice, on the grounds that they are intelligible only against a background of propositional contents conferred by the sorts of interactions considered here. Real-world phenomena such as these (which presumably will be present at every stage in the development of actual practice) create play between the proprieties of practice (especially scorekeeping practice) an interpreter takes a community to be bound by and their actual behavior. The status of this discrepancy—a normative generalization of the competence-performance distinction—is discussed in Chapter 9.

46. Daniel Dennett, "Intentional Systems," in *Brainstorms* (Montgomery, Vt.: Bradford Books, 1978), p. 19.

47. Ibid., p. 20.

48. Ibid.

49. There are subtleties that require qualifying this formula, some of which are discussed below in 8.5.2.

4. Perception and Action

1. The further conditions that have been suggested in response to examples of the sort Gettier first presented are not discussed here. Those impressed by the significance of these counterexamples to the sufficiency of JTB analyses of "S knows that p" may want to treat the term 'knowledge' as it appears in the rational reconstruction presented here as really invoking knowledge* (which is defined simply as justified true belief). The considerations motivating many of the proposals for a fourth condition on knowing can be straightforwardly transposed into the idiom of this work, so those who believe that the status reconstructed here must be further specified in order to deserve the central role it is given might try the experiment of seeing how their favorite candidate looks in deontic scorekeeping guise. The basic elements of the social practice model's construal of the statuses corresponding to justification, truth, and belief do not turn on the kind of niceties concerning their interaction that attempts to for-

mulate a fourth condition must address. It is explained below why knowledge*
deserves to be accorded a fundamental explanatory role, regardless of its relation
to what is expressed by the English word 'know'.

2. Of course, there is no necessity to adopt either direction of explanation; it may
be that neither term is intelligible apart from its relation to the other.

3. This claim about the conceptually basic level of practices is compatible with the
institution of a distinction in more sophisticated practices between *claiming*
that *p* and claiming to *know* that *p*. This point is discussed further along.

4. The distinction being made here between the present approach and standard
ones is different from, though intimately related to, the fundamental difference
between understanding belief, justification, truth, and so knowledge as kinds of
normative status rather than as kinds of *natural state*, so that one looks for
proprieties, rather than *properties*, corresponding to them. The connection be-
tween the issues consists in the difficulty of appreciating the significance of the
social distinction of attitude between acknowledging and attributing unless one
is already thinking of what these are attitudes toward in normative terms, as
commitments and entitlements—that is, deontic statuses, not descriptive states.

5. Being a *logical* being—having the expressive resources to make propositionally
explicit crucial semantic and pragmatic features of the social practices in virtue
of which one is a linguistic, rational, cognitive being—is a further, optional stage
of development, which presupposes this fundamental one.

6. Insofar as it addresses sapience rather than sentience—that is, insofar as it is
concerned with intentionality in the sense defined by possession of *proposi-
tional* attitudes.

7. One source of these views is the "thermometer view of knowledge" put forward
by D. Armstrong in *Belief, Truth, and Knowledge* (Cambridge: Cambridge Uni-
versity Press, 1973). F. Dretske defends such a view in *Knowledge and the Flow
of Information* (Cambridge: MIT Press, 1981), and more recently in "The Need
to Know," in *Knowledge and Skepticism*, ed. M. Clay and K. Lehrer (Boulder,
Colo.: Westview Press, 1989). A. Goldman's development of a reliabilist theory
is expounded in *Epistemology and Cognition* (Cambridge: Harvard University
Press, 1986), and again in "Précis and Update of *Epistemology and Cognition*,"
in *Knowledge and Skepticism*, ed. Clay and Lehrer. R. Nozick puts forward a
version of reliabilism in *Philosophical Explanations* (Cambridge: Harvard Uni-
versity Press, Belknap Press, 1981). M. Swain has a version that takes reasons
more seriously than the others, in *Reasons and Knowledge* (Ithaca: Cornell
University Press, 1981), and more recently in "Justification, Reasons, and Reli-
ability," *Synthese* 64, no. 1 (1985): 69–92. This list is intended only to provide a
reasonably representative sample.

8. This sort of talk involves a promissory note—only in Chapter 6 is the official
account developed to the point that it explains (in terms of substitution infer-
ences) the relation between applying a *predicate* and making a *claim*.

9. Alvin Goldman, "Discrimination and Perceptual Knowledge," *Journal of Phi-
losophy* 73, no. 20 (1976).

10. J. L. Austin "Other Minds," in his *Philosophical Papers* (New York: Oxford
University Press, 1961).

11. Not objective "to begin with" because, as was acknowledged already in Chapter
1, it is a critical criterion of adequacy of any account of the use of empirical

concepts that it be able to explain how *in the end* objective proprieties governing that use can come into play—how claims can be understood as true or false regardless of whether anyone or everyone takes them to be so, depending rather on how things are with what the claims are about. Not until Chapter 8 will assembly be complete of the raw materials necessary to explain what is involved in this sense of objectivity. One important clue has been put on the table already, however. For it has been pointed out how a distinction can arise between what someone to whom a commitment is attributed is *justified* or *entitled* to believe, and what is in fact *true.* This is just the social-perspectival distinction of attitude between *attributing* deontic statuses and *undertaking* them. It is in terms of this fundamental *social* articulation of deontic attitudes that the possession by claims and concepts of *objective* representational content is eventually to be understood.

12. "EPM," sec. 32.
13. Ibid., sec. 34.
14. Ibid., sec. 35.
15. Ibid., sec. 32.
16. Ibid., sec. 35.
17. Ibid.
18. Ibid., sec. 36.
19. This possibility was floated in the author's original discussion in "Asserting."
20. It is a commonplace among teachers of mathematics that students often profess to be completely unable to deal with problems of a certain sort long after they are in fact able to solve them reliably.
21. Assuming that the attributor of knowledge both attributes to Monique commitment to the claim and undertakes such commitment (that is, endorses the claim and so takes it to be true).
22. Sellars, "Some Reflections on Language Games," in Sellars, *Science, Perception, and Reality* (New York: Routledge and Kegan Paul, 1963).
23. "EPM," sec. 32.
24. "There are two ways in which a sentence token can have credibility: (1) The authority may accrue to it, so to speak, from above, that is, as being a token of a sentence type all the tokens of which, in a certain use, have credibility, e.g. '2 + 2 = 4'. In that case, let us say that token credibility is inherited from type authority. (2) The credibility may accrue to it from the fact that it came to exist in a certain way in a certain set of circumstances, e.g. 'This is red'. Here token credibility is not derived from type credibility" (ibid.).
25. "Epistemology Naturalized," in *Ontological Relativity and Other Essays* (New York: Columbia University Press, 1969).
26. As they are called in the author's "Asserting," p. 643. Cf. Hegel's remark in a related context that "one barren assurance is of just as much worth as another," in the Introduction to the *Phenomenology* (par. 76).
27. Recall from the earlier discussion that Stalnaker introduces the linguistic approach in a way that leaves room for relational linguistic accounts such as Davidson's and the one pursued here, but he then discusses and argues against only the reductive versions.
28. Though just how this latter possibility should be understood is not officially addressed until Chapter 8.

29. As will become clear, it would be another sort of reductive mistake to *identify* rational agency with preference-maximizing.

30. David Lewis, *Convention: A Philosophical Study* (Cambridge: Harvard University Press, 1969), also "Languages and Language," in *Language, Mind, and Knowledge*, Minnesota Studies in the Philosophy of Science, vol. 7, ed. Keith Gunderson (Minneapolis: University of Minnesota Press, 1975).

31. "Communication and Convention," in *Inquiries into Truth and Interpretation* (New York: Oxford University Press, 1984), p. 280.

32. *FPL*, pp. 298ff., 354–356.

33. "Some Reflections on Language Games."

34. Though as will emerge below, in some cases the nonlinguistic intentional performance just *is* (has the scorekeeping significance of) the acknowledgment of a practical commitment, rather than being a *response* to such an acknowledgment. The difference corresponds to that between intentions-in-action and prior intentions.

35. Section V below discusses some different patterns of availability of practical reasons across interlocutors.

36. Here 'theoretical' is opposed to 'practical', as pertaining to relations exclusively between doxastic discursive commitments. This use ought not to be confused with the sense of 'theoretical' that is opposed to 'observational', within the doxastic sphere. In the latter usage, following Sellars's practice (in "Empiricism and the Philosophy of Mind" and elsewhere), theoretical claims are distinguished as those one cannot become entitled to noninferentially, by the exercise of reliable differential responsive dispositions to acknowledge doxastic commitments. Theoretical vocabulary is then distinguished as that which appears only in claims that are theoretical in this sense. It is this usage that stands behind Sellars's claim that the distinction between the observable and the theoretical is not ontological but only methodological. Neptune was a theoretical entity so long as claims about it could be arrived at only inferentially, as based on its perturbation of the observable orbits of other planets. It became observable, however, once we built telescopes powerful enough to make it subject to noninferential reporting. Something is theoretical or observable in this sense only relative to our practices; nothing is "intrinsically" theoretical.

37. Davidson wants to analyze intention in terms of reasons. A representative formulation is "Someone who acts with a certain intention acts for a reason" ("Intending," reprinted as Chapter 5 of *Actions and Events* [New York: Oxford University Press, 1980], p. 84). Irrational actions accordingly pose a problem for him (which he addresses in "How Is Weakness of the Will Possible?" reprinted as Chapter 2 of *Actions and Events*). According to the present view, he is mistaking a global condition on intention for a local condition, as a result of failing to distinguish commitment from entitlement.

38. This formulation is intended to encompass both committive and permissive inferences. The latter can be thought of as conveying *entitlement* to the commitments that are their premises to *entitlement* to the commitments that are their conclusions. In scorekeeping terms, for an attributor to endorse such a permissive practical inference is to take it that anyone who is committed and entitled to the premises is entitled, though not committed, to the conclusion. Nonetheless, the conclusion can be thought of as a commitment, because what one is entitled to is in the first instance a commitment.

39. This use of 'shall' is based on Sellars's technical usage of 'shall' as an expression of intention. The treatment of action, intention, and practical reasoning presented here owes a great deal to Sellars's seminal work on the topic, "Thought and Action," in *Freedom and Determinism,* ed. Keith Lehrer (New York: Random House, 1966).

40. The point just made about the inheritance by practical commitments of inferential relations from the corresponding doxastic commitments can be illustrated in these terms by noting that if "I *will* wear a necktie" commitment-entails "I *will* wear something around my neck," then "I *shall* wear a necktie" commitment-entails "I *shall* wear something around my neck."

41. Beginning with "Actions, Reasons, and Causes," reprinted as Chapter 1 in *Actions and Events.* Quotations are from p. 4.

42. The expressive role of such ascriptions should be understood by analogy to that of the explicit ascriptions of doxastic commitment discussed below in 8.1–4.

43. "Intending" p. 86.

44. Hempel discusses this feature of inductive arguments in detail in *Aspects of Scientific Explanation* (New York: Free Press, 1965), pp. 394–403. An example he offers (recast in the idiom of this work) is that one can have good evidence both for the inference codified by the conditional "If the barometer falls, it almost certainly will rain" and for the inference codified by "If the sky is red at night, it almost certainly will not rain." Since there are occasions on which one can be entitled to commitment to both of the antecedents, the incompatibility of the conclusions shows that these inferences cannot be commitment preserving. But they can each be entitlement preserving, even though in the situation where one is entitled to both antecedents, assertion of either can serve as a challenge to the conclusion of the other inference.

45. Literally at the end of this story, in Chapter 9, building on the social-perspectival account of objectivity developed in Chapter 8.

46. "Intending," p. 84.

47. Ibid., p. 85.

48. G. E. M. Anscombe, *Intention* (Oxford: Blackwell, 1959), and Davidson, originally in "Actions, Reasons, and Causes."

49. Davidson, "Actions, Reasons, and Causes," p. 8.

50. Ibid., p. xiii.

51. This and the next passage quoted are from John Searle, *Intentionality* (Cambridge: Cambridge University Press, 1983), pp. 84–85.

52. Sellars, "Thought and Action," p. 110.

53. Ibid., p. 109.

54. This is the conclusion of Davidson's "Intending," summarized at pp. 100–101.

55. Sellars's central text on this topic is "Thought and Action." The view is introduced in his earlier "Imperatives, Intention, and the Logic of 'Ought'," *Methodos* 8 (1956): 228–268, and is developed at greater length in the Tsanoff lectures, delivered at Rice University in 1978, entitled "On Reasoning about Values." Castañeda's treatment dates to his 1952 University of Minnesota master's thesis, "An Essay on the Logic of Commands and Norms." The fullest statement of his view is in *Thinking and Doing: The Philosophical Foundations of Institutions* (Dordrecht: D. Reidel, 1975). A concise summary can be found in "The Two-Fold Structure and the Unity of Practical Thinking," in *Action Theory,* ed. M. Brand and D. Walton (Dordrecht: D. Reidel, 1976), pp. 105–130.

56. Indeed, in his extremely useful summary of Sellars's views on this topic (pp. 149–188 in *The Synoptic Vision: Essays on the Philosophy of Wilfrid Sellars*, by C. F. Delaney, Michael J. Loux, Gary Gutting, and W. David Solomon [Notre Dame, Ind.: University of Notre Dame Press, 1977]), W. David Solomon paraphrases Sellars's view about the intentions expressed by his regimented 'shall' locutions in just this way: "Whereas expressions of intention manifest my commitment to act at some future time (perhaps precisely datable, perhaps not), volitions are commitments on my part to act here and now" (p. 163). Sellars's treatment of pro-attitudes is quite different from that presented here, and in any case, he has no general account of beliefs as commitments, to which intentions might be assimilated.

57. Davidson, "Intending," p. 90.

58. Ibid.

59. Ibid.

60. The analogy is not exact, for 'shall' indicates a kind of pragmatic *force*, the acknowledging of a practical commitment. A 'shall' statement is not an ordinary assertion in which that force becomes explicit as part of the content, for instance in the way in which 'believe' can be used to make the attribution of a doxastic commitment explicit as part of the assertible content of an ascription (as discussed in Chapter 8). As Sellars says: "'Shall', in spite of its logical role, can be said to be a *manner* rather than a *content* of thought" ("Thought and Action," p. 109). For 'shall' statements do not embed in more complex statements—paradigmatically as the antecedents of conditionals—in the way ordinary assertible contents must. Such embedding strips off the pragmatic force associated with the utterance of the embedded sentence. Thus when I say "If I believe that Kant liked turnips, then I believe that he liked some tuber," I am not saying that I believe either claim. The force of the (self-)ascription "I believe that Kant liked turnips" has been lost, but 'believe' still means exactly what it does in unembedded contexts, even though the pragmatic significance of uttering it is different. It is different with 'shall'. "I *shall* marry, so I won't be a bachelor" is a good inference, but when codified as a conditional, it takes the form "If I *should* marry, I won't be a bachelor," not the ungrammatical "If I *shall* marry, I won't be a bachelor," which contains a defective use of 'shall'. The rest of this section indicates why this shift from 'shall' to 'should' is required in embedded contexts. The result is that 'shall' statements do not make acknowledgments of practical commitments explicit as assertible contents in quite the same sense that 'believe' statements make attributions of doxastic commitment explicit as assertible contents.

61. Self-attribution of a deontic state is what is made explicit by certain uses of 'I' and is not equivalent to merely attributing it to oneself. For the latter may be attributing it to someone who is, as a matter fact, though one is not aware of this fact, oneself—as I might attribute a certain commitment to whoever wrote the words appearing on a scrap of paper I find in the street, not realizing that I wrote them a year ago. It is not at all surprising that one can attribute a practical commitment to oneself without acknowledging it; it is more surprising that one can even self-attribute such a commitment without acknowledging it (though not without undertaking it). This issue is discussed in more detail below in 8.5.2.

62. *Foundations of the Metaphysics of Morals*, trans. Lewis White Beck (Indianapolis: Bobbs-Merrill, Library of Liberal Arts, 1959), p. 29; p. 412 of vol. 4 of the Akademie Textausgabe.
63. Ibid., p. 30; p. 413.

5. The Expressive Role of Traditional Semantic Vocabulary

1. In an extended sense: intentions have propositionally expressible conditions of satisfaction corresponding to the claims the agent is practically committed to making-true.
2. Permissive or entitlement-preserving inferences, which are of the first importance in justificatory practices—particularly inductive ones—are notoriously difficult to parse in terms of truth conditions. It is because of the hope that the notion of *reliability* can supply what is wanted that that concept assumes the significance it has for epistemologists who understand the contents of knowledge claims in terms of truth conditions. That notion, and the gerrymandering difficulties to which it is subject, are discussed in Chapter 4.
3. "The Concept of Truth in Formalized Languages," in *Logic, Semantics, Metamathematics: Papers from 1923–1938* (New York: Oxford University Press, 1956), pp. 152–278.
4. Thomas Gray, "Elegy Written in a Country Churchyard."
5. An analogy is the way in which, once the expressive role of 'looks' or 'seems' is properly understood, it becomes apparent that the incorrigibility of statements formed by the use of these operators is trivial in a way that makes them unsuitable precisely for the role of foundation of knowledge that Descartes assigned to them. The analysis required for this argument is presented in Sellars's "EPM" and is sketched below in Section II.
6. Part of what is distinctive about the present approach, however, is that what are here treated as *semantic* primitives are themselves explained in terms of a prior *pragmatics*, which in turn appeals to *normative* primitives, themselves made available by mapping the theoretical idiom onto our ordinary talk.
7. Indeed, that is why semantic theorists, as opposed to linguists, have been in general so little interested in this notion—which will be taken here to be of absolutely the first importance to semantic theory. Chastain (whose work is cited below in Section IV) is a notable exception, as is Hintikka.
8. The material presented in this section and the next originally appeared as "Pragmatism, Phenomenalism, and Truth Talk," in *Realism (Midwest Studies in Philosophy* 12 [1988]), pp. 75–93.
9. The point here does not concern merely the *senses* of the contrasted expressions, but the *extensions* they determine. The appropriateness of this question would have to be defended by adducing cases in which a belief apparently "worked" and was not true, or vice versa. Such cases are not far to seek. This sort of argument is considered more carefully below.
10. *Pragmatism* (Cambridge: Harvard University Press, 1978), bound with its sequel, *The Meaning of Truth*, which as here interpreted ought to be titled *The Meaning of Taking-True*. For an important assessment on a larger scale, see R. Rorty's "Pragmatism, Davidson, and Truth," in *Philosophy of Donald Davidson: A Perspective on Inquiries into Truth and Interpretation,* ed. E. LePore (Oxford:

Basil Blackwell, 1986). The essay that originally presented the material from which the first part of this chapter is drawn was written as a tangential response to this piece of Rorty's.

11. For instance, in "Truth," reprinted in *Truth*, ed. G. Pitcher, (Englewood Cliffs, N.J.: Prentice-Hall, 1964), pp. 32–53. Strawson's view is often also referred to as a 'redundancy' account, in the sense of an account focusing on redundancy of force.

12. *Logic: The Theory of Inquiry* (New York: Holt, Rinehart, Winston, 1938).

13. Anthony Appiah, *Assertion and Conditionals*, Cambridge Studies in Philosophy (Cambridge: Cambridge University Press, 1985), p. 103. J. T. Whyte develops this line of thought in "Success Semantics," *Analysis* 50, no. 3 (June 1990): 149–157, and "The Normal Rewards of Success," *Analysis* 51, no. 2 (March 1991): 65–73. The criticisms mentioned in the text are applied to this version in the author's "Unsuccessful Semantics," forthcoming in *Analysis*.

14. It is in many ways more natural to think of the approach considered here as "expressivism," in Allan Gibbard's sense (see *Wise Choices, Apt Feelings* [Cambridge: Harvard University Press, 1990], pp. 7ff. and throughout), rather than as "phenomenalism." There are differences, however, and in any case the account offered here of expression as making explicit forbids the more familiar rubric.

15. In Quine's *Ontological Relativity and Other Essays* (New York: Columbia University Press, 1969).

16. As Sellars argues in "Phenomenalism," in *Science, Perception, and Reality* (London: Routledge and Kegan Paul, 1963).

17. "Ascriptivism" and "Assertion," reprinted in Geach's *Logic Matters* (Berkeley: University of California Press, 1972), pp. 250–253 and 254–269.

18. "Reference and Understanding," in *Meaning and the Moral Sciences* (London: Routledge and Kegan Paul, 1978), p. 108.

19. The objectivity proofs below at 8.6.5 show that this difference is preserved by the deontic scorekeeping model of discursive practice.

20. D. Grover, J. Camp, and N. Belnap, "A Prosentential Theory of Truth," *Philosophical Studies* 27 (1975): 73–125.

21. See P. Geach, *Reference and Generality* (Ithaca: Cornell University Press, 1962), pp. 124ff. For some complications, see B. Partee, "Opacity, Coreference, and Pronouns," in *Semantics of Natural Language*, ed. G. Harman and D. Davidson (Dordrecht: D. Reidel, 1972), pp. 415–441.

22. This discussion must be merely preliminary, since singular terms, pronouns, and substitution have not yet been officially discussed—the model of asserting and inferring stays resolutely at the level of sentences. These deficiencies will be remedied further along. The point of this discussion is in part to motivate those later discussions and also to set up some strategic criteria of adequacy for them. Anaphora in the grammatical category of singular terms is considered in greater detail later in this chapter, though it is only in Chapter 7 that it is shown how such a notion can be defined in the official model of linguistic practices. Substitution is introduced, and singular terms are defined, in Chapter 6.

23. This claim is qualified below in 7.3–4, where an indispensable expressive role is discerned even for 'lazy' anaphora—not only in spite of, but because of this redundancy.

24. Grover, Camp, and Belnap, "Prosentential Theory of Truth," p. 87.

25. The authors of the original theory may believe that for any syntactic prosentence type (paradigmatically "That is true") and any declarative sentence tokening there is potentially an anaphorically dependent prosentence tokening of that type that has the declarative tokening as its antecedent. On the disquotational account of lazy prosentences offered in emendation below, that formulation would not hold.

26. Grover, Camp, and Belnap, "Prosentential Theory of Truth," p. 91.

27. Ibid., p. 109.

28. The material presented in this section originally appeared as "Reference Explained Away," *Journal of Philosophy* 81, no. 9 (September 1984): 469–492, reprinted in *Philosopher's Annual* "Best Ten Philosophical Articles of 1984."

29. "Truth and Correspondence," *Journal of Philosophy* 59, no. 2 (January 1962): 29–56, reprinted in *Science, Perception, and Reality;* and Chapter 4 of *Science and Metaphysics* (London: Routledge and Kegan Paul, 1968).

30. In *Language, Mind, and Knowledge,* Minnesota Studies in the Philosophy of Science, vol. 7, ed. Keith Gunderson (Minneapolis: University of Minnesota Press, 1975), pp. 194–269.

31. More exactly, it is tokenings that stand in anaphoric relations, though this nicety will continue to be suppressed in the text. Written expressions are particularly vulnerable to reuse of tokens, which is what enforces the distinction. Tokens of pronouns are as vulnerable to this as are tokens of demonstratives (where the same "you are here" sign might be moved from one map to another on campus).

32. Not all elements of anaphoric chains need be understood as singular referring terms. Chastain says that quantificational, modal, and hypothetical contexts are "referentially segregating" and that the syntactically singular expressions that occur inside them should not in general be understood as singular referring terms. The claims made here do not require special treatment of such segregated occurrences.

33. The significance of this intersubstitution test is pursued at length in Chapter 6. It is suggested below at 7.4.3 (and argued in 8.5.4–6 in connection with Kripke's puzzle about belief) that proper names should themselves be understood as functioning anaphorically, an idea that Chastain himself suggests. For the purpose of introducing anaphoric ideas, however, it does no harm to pretend that proper names function in the idealized way they have traditionally been conceived of as functioning.

34. It would be less clumsy to call the expressions not all cotypical tokens of which are coreferential, 'token reflexive'. That policy is not adopted here, because the phrase has an established usage (due to Reichenbach) and is not generally thought of as applying to expressions like 'the Senator', even when such expressions are used as anaphoric dependents, which are among the paradigmatic cotypically nonintersubstitutable term occurrences. It is argued in Chapter 7 that the special substitution conditions applying to anaphorically dependent expressions are more explanatorily fundamental than indexicality of canonical token-reflexives.

35. One difference that might be remarked between ordinary pronouns and those formed by indirect description concerns *backward anaphora*, in which the anaphorically dependent occurrence precedes its antecedent in the discourse.

Such cases are unusual, though by no means always deviant or strained, for ordinary dependents such as 'he' and 'the man'. Consider the intertwining anaphoric chains in the opening sentences of Henry James's "The Beast in the Jungle":

What determined the speech that startled *him* in the course of *their* encounter scarcely matters, being probably but some words spoken by *himself* quite without intention—spoken as *they* lingered and slowly moved together after *their* renewal of acquaintance. *He* had been conveyed by friends an hour or two before to the house at which *she* was staying; the party of visitors at the other house, of which *he* was one, and thanks to whom it was *his* theory, as always, that *he* was lost in the crowd, had been invited over to luncheon. There had been after luncheon much dispersal, all in the interest of the original motive, a view of Weatherend itself and the fine things, intrinsic features, pictures, heirlooms, treasures of all the arts, that made the place almost famous; and the great rooms were so numerous that guests could wander at their will, hang back from the principal group and in cases where they took such matters with the last seriousness give themselves up to mysterious appreciations and measurements. There were persons to be observed, singly and in couples, bending toward objects in out-of-the-way corners with their hands on their knees and their heads nodding quite as with the emphasis of an excited sense of smell. When they were two they either mingled their sounds of ecstasy or melted into silences of even deeper import, so that there were aspects of the occasion that gave it for *Marcher* much the air of the 'look round' previous to a sale highly advertised. (In *Henry James: Selected Fiction,* ed. Leon Edel [New York: G. P. Dutton, 1964], p. 482. There are four more sentences before the initial 'she' is anchored to a token of 'May Bartram', thus allowing 'their' and 'they' to become attached as well.)

Indirect descriptions, in virtue of the explicit way they pick out the tokens they depend upon anaphorically, exhibit no such prejudice for the discursive past and, accordingly, often possess 'antecedents' only in the broader sense of anaphorically inheriting content from another tokening.

36. This last sort of example shows that the term token can be picked out in a variety of ways, in particular, by citing a predicate type to pick out a sentence tokening that is the characterizing or calling and that contains the term token (not otherwise specified) that on the present account is the anaphoric antecedent of the indirect description in the example. It may also be noted that *in*definite descriptions can similarly be constructed from indirect anaphoric sortals, as in "A woman the lawyer referred to as 'she who must be obeyed' explained the matter to us," which both initiates a new chain *and* characterizes the referent of that chain by relation to some antecedent tokening of 'she who must be obeyed' by the lawyer in question. See the discussion of referential predication below.

37. Speakers' reference is discussed briefly below as an anaphoric phenomenon, and at greater length at 7.5.6 and 8.2.2.

38. It may seem that the presence of a sortal restriction on indirect descriptions causes difficulties. The issue can be avoided, as in (*t*), however, by using an *anaphorically dependent* sortal. 'One' anaphora has long been recognized by

linguists as permitting anaphoric proforms as stand-ins for common nouns, as in "There were red pens as well as green ones on the table." In fact philosophers have made up the expression 'referent of t' to mean 'the one referred to by t'. So the fact that indirect descriptions are sortally restricted, as are ordinary descriptions (or for that matter quantifications) in natural language, adds no new difficulties to an anaphoric analysis of 'refers'.

39. Chapters 6 and 7 make explicit how these intersubstitutability claims are to be understood.

40. The identification of indirect descriptions by the iteration test provides a sense in which one term *type* can be considered to be anaphorically dependent on another *type*, as the iterated indirect description is on the type of its antecedent token. But this notion of type-anaphora is entirely derivative from the basic notion of token-anaphora, a derivation made possible by the existence of operators that form lexically complex pronouns that are invariant under cotypical intersubstitution.

41. It may seem that talk here and elsewhere of 'picking out', 'coreferring', and so on begs important questions. But it will be shown below that such talk can be understood in a way that does not commit one to reference relations.

42. Joe Camp pointed out the need to deal with these relations, which 'piggyback' psychological relations on semantic ones.

43. The significance of this principle for Frege is discussed at greater length in the author's "Frege's Technical Concepts," in *Frege Synthesized: Essays on the Philosophical and Foundational Work of G. Frege*, ed. L. Haaparanta and J. Hintikka, Synthese Library (Dordrecht: D. Reidel, 1986), pp. 253–295. The motivation for it is pursued at some length in Chapters 6 and 7 below.

44. Compare Frege's similarly motivated paraphrase of "Jupiter has four moons" into "The number of Jupiter's moons is (=) four," in *GL*, sec. 57.

45. "Tarski's Theory of Truth," *Journal of Philosophy* 69 (1972): 347–375.

46. *Journal of Philosophy* 74, no. 10 (October 1977). Also in this paper Grover independently states a weaker version of the observation exploited in the present account: "Descriptive phrases such as 'just mentioned', 'have been talking about', 'are referring to', which ostensibly describe discourse may often be used merely to locate an antecedent piece of discourse from which a referent is inherited" (p. 594). See also Grover's "'This is False' on the Prosentential Theory," *Analysis* 36 (1976): 80–83.

47. The special anaphoric nature of the disquotation involved is most apparent in the case where the indirect description picks up the speaker's reference of its mentioned antecedent.

48. Though it may be recalled that a sense is given above to the claim that "truth is one" (in 4.4.3, appealed to again in 4.5.4).

49. The specification that it is to be a language-*in-use* is intended to indicate that it should be thought of as comprising not only words and sentences but also the practices in virtue of which they mean what they mean.

50. Paul Boghossian, "The Status of Content," *Philosophical Review* 99, no. 2 (April 1990): 157–184.

51. Ibid., p. 181.

52. As Frege says in "The Thought," in *Essays on Frege*, ed. E. O. Klemke (Urbana: University of Illinois Press, 1968), pp. 507–536.

53. In fact, Boghossian does not so much argue for this claim as assume it, for lack

of an alternative: "For the remainder of the paper, therefore, I will assume that contents just are truth conditions" ("The Status of Content," p. 173).

54. Compare Sellars's claim in "EPM" that once the expressive role of 'looks' or 'seems' is properly understood, its unsuitability for service in characterizing an epistemological foundation for empirical knowledge becomes apparent, in spite of the incorrigibility of 'seems' claims that so impressed Descartes.

55. A long story would be required to distinguish properly between cases where it would be correct to take the proprieties governing the use of the concept to be determined in this way from those where the alien substance should be taken to be a new kind of water. Our use of 'water' involves commitment to all samples of water having the same fundamental physical constitution, and we have sufficiently developed conceptions of what that means to rule out XYZ and H_2O correctly being put in the same box, even in a case where we could not tell the difference. It is also possible, of course, to describe cases in which 'water' is so used as properly to be applied to any clear, tasteless, odorless, thirst-quenching liquid, regardless of its composition. It is not easy to say what it is about the prior *behavior* (specified in nonnormative terms) of a tribe that has experience only of, say, warm brackish water that determines whether they are (whether they know it or not) committed to apply the same term they use to the sort of cold, pure water that gushes from mountain springs they have never seen. But this sort of difficulty stems from the fact that many interpretations couched in normative vocabulary are compatible with any given specification of behavioral dispositions or regularities that is couched exclusively in nonnormative terms. The point being made here is rather that discursive practices (which must be specified in normative terms) and the concepts they make available are individuated in part by the causal commerce with objects of various kinds that they incorporate, and that for this reason the proprieties that are implicit in those practices may (according to an interpreter) outrun the explicit discriminative capacities of those being interpreted.

56. So to speak—for according to the idiom being recommended, there is nothing that is in principle ⁵outside⁵ discursive practice.

57. *PI*, sec. 95.

58. John McDowell has discussed this issue extensively and insightfully, beginning in "Singular Thought and the Extent of Inner Space," Chapter 5 in *Subject, Thought and Context*, ed. P. Pettit and J. McDowell (New York: Oxford University Press, 1986); in the manuscripts "Knowledge and the Internal" and "Scheme-Content Dualism, Experience, and Subjectivity"; and in his 1991 Locke lectures.

59. Tennyson, *The Idylls of the King: The Last Tournament*.

6. Substitution

1. If this sort of approach is possible, why is it not the one pursued here? Not because of inferentialism per se, but because of the dual collateral commitments: to understanding semantics in terms of pragmatics—that is, to conceiving the use of expressions as conferring their contents—and to treating asserting as the fundamental linguistic performance (the pragmatic priority of the propositional).

2. This issue is independent of the desirability of an account in nonsemantic terms

(though it is important not to confuse this with the demand for naturalistic—that is, nonnormative—terms) of what capacities the semantic capacities in some sense consist in the deployment of. It is one thing to claim (how could it be denied?) that causal interactions of various sorts with particular objects is a necessary condition of being able to represent empirical states of affairs; it is another to claim that some of these interactions ought to be understood as semantically primitive, in that what it is to represent such a state of affairs ought to be understood in terms of them.

3. *GL*, sec. 65 is one place the distinction is explicitly acknowledged.

4. *FPL*, p. 417.

5. In the standard situation, it has been stipulated which multivalues are designated, so the second assignment determines the first. It is important, however, to keep the two sorts of assignment conceptually distinct.

6. *FPL*, p. 361.

7. Ibid., Appendix to Chapter 12 (on multivalued logic).

8. Thus the designatedness value *undesignated* need not be preserved by good inferences, and substitutions that turn undesignated into designated sentences are strictly irrelevant for assessments of validity of compound sentences.

9. *FPL*, chap. 12.

10. The expressive role of such propositional-attitude-ascribing locutions is explained in terms of the social articulation of substitution-inferential commitments in Chapter 8.

11. It is important that this formulation is in terms of deontic status and does not indicate how to translate the remarks in terms of social deontic attitudes. It does not indicate who is attributing the state and to whom it is attributed. It matters how one goes about eliminating this vagueness. That issue is addressed in Chapter 8, on ascribing propositional deontic attitudes. Notice also that this issue concerns local interdefinability of assertional and inferential commitments. To say that (commitment to) the goodness of inferences is not settled by (commitment to) the truth of the premises and conclusions is not yet to deny that inferential commitments as a whole might be fixed by assertional commitments as a whole, or vice versa. According to the picture presented in Chapter 3 of the sort of indissoluble conceptual package that assertional and inferential commitments and entitlements make, however, the intimate mutual presupposition of discursive and inferential commitments precludes such global reducibilities.

12. See the author's "Semantic Paradox of Material Implication," *Notre Dame Journal of Formal Logic* 22, no. 2 (April 1981): 129–132.

13. *FPL*, p. 401.

14. Strictly, what is referred to by a singular term is a particular. Not all particulars are *objects*; there are also events, processes, and so on. The present argument does not turn on the differences among these sorts of particulars, and it will often be more convenient simply to talk of objects, where in fact any sort of particular can be involved.

15. *Word and Object* (Cambridge: MIT Press, 1960), p. 96; see also p. 90.

16. For present purposes, the fact that the same adverb can modify multiplace predicates—and so should more generally be construed as a $((\langle T_1 \ldots T_n \rangle \to S) \to (\langle T_1 \ldots T_n \rangle \to S))$—can be ignored.

17. David Lewis's "General Semantics" (in *Semantics of Natural Language*, ed.

G. Harman and D. Davidson (Dordrecht: D. Reidel, 1972) is a masterful account of the ideas behind functional-categorial semantics.

18. *GL*, p. x, secs. 46, 60, 62.

19. Two works that endorse this order of explanation are D. Davidson, "Reality without Reference," and J. McDowell, "Physicalism and Primitive Denotation: Field on Tarski," both in *Reference, Truth, and Reality*, ed. Mark Platts (London: Routledge and Kegan Paul, 1980).

20. "Compound Thoughts," *Mind* 72 (1963): p. 1.

21. In a sense, of course, we do not know how many such sentences there are, even restricting ourselves to a basic vocabulary, since we do not have a syntactically adequate grammar for any natural language. But there are grammars that will generate *only* sentences of English. The difficult thing is getting one that will generate *all* of them, without generating all sorts of garbage as well.

22. For instance, Davidson emphasizes this point in his influential "Theories of Meaning and Learnable Languages," in *Proceedings of the 1964 International Congress of Logic, Methodology, and Philosophy of Science* (Amsterdam: North Holland Publishing, 1965).

23. Notice that the problem of projection such a strategy addresses concerns moving from *proprieties* governing the use of one set of sentences to proprieties governing the use of a superset. A quite different issue concerns the relation between the correct use even of the sentences in the initial subset and the actual occasions of use or dispositions of the community to use them. These puzzles must be sharply separated, for the first remains within the normative dimension, asking about the relation between two different sets of practically embodied norms, while the second asks about the relation between such norms and the nonnormative happenings that express them.

24. Type/token issues are suppressed for the purposes of this chapter. The complications they introduce are the topic of Chapter 7.

25. Corner quotes should be discerned as required here, so that '*p*' refers to the quote-name of the sentence the variable *p* stands for, not for the quote-name of the variable letter.

26. This requirement is not absolute. The author's "STSSD" shows how to make do just with substitutional relations among substituted-in expressions and how to do without antecedently distinguishable substituted-for expressions.

27. Strictly speaking, this is true only of what Dummett calls "complex" predicates, by contrast to "simple" ones, about which more below. But as Dummett points out in making the distinction, Frege "tacitly assimilated simple predicates to complex ones" (*FPL*, p. 30).

28. *Subject and Predicate in Logic and Grammar* (London: Methuen, 1974).

29. From Frege's mature point of view, this qualification does not need to be made: sentences *are* singular terms, and the frames *are* predicates. This is what motivates Frege's classification of sentences as singular terms. As will be pointed out below, this need not be the whole story about sentences, a fact that immunizes Frege somewhat from Dummett's scandalized response to this point. Qua subsentential expressions, sentences are singular terms; the thesis is innocent of the objectionable implications Dummett complains about (missing the special role of sentences as usable to make moves in the language game—as though Frege had no idea of force, and as though being a name of the True or the False did not

play a very special role for him) because sentences are not *essentially* subsentential expressions, and it is not as subsentential expressions that they have their special pragmatic position. I am grateful to John McDowell for pointing this out.

30. This point is distinct from, although related to, the distinction Dummett makes, in Chapter 2 of *FPL*, between simple and complex predicates. Dummett there points out (following Geach's discussion in "Quine on Classes and Properties," *Philosophical Review* 62 (1953): 409–412) that there is no simple *part* or subexpression common to "Rousseau admired Rousseau" and "Kant admired Kant" that is not also a part of "Kant admired Rousseau." Yet the first two share with each other a complex predicate that they do not share with the third. One of Frege's great discoveries was that one must be able to discern predicates in this sense (complex, or substitutionally derived ones) in order to appreciate the inferential role of sentences like "Anyone who admires someone admires himself." For one must appreciate the different patterns they instantiate in order to see that in the context of that quantificational claim, "Kant admired Rousseau" entails "Kant admired Kant." Thus the status of predicates as playing derived substitution-structural roles is what lies behind the second of Strawson's stigmata distinguishing predicates from singular terms: that they are subject to quantification. Concern with quantification, in particular with codifying the inferential role of quantificational claims, enforces the distinction between simple and complex predicates, between expressions that can be substituted for and those that are substitutional frames. But the need for this distinction is not, as Dummett claims (pp. 28, 30), simply a consequence of the presence of quantificational locutions in a language. Complex predicates must be discerned by anyone who has mastered the sort of pattern of inference that is typically made explicit by a quantificational expression, such as $(x)(y)[Rxy \rightarrow Rxx]$. Such inferential connections can be important already in a language, even though quantifiers have not yet been introduced to codify them explicitly as the contents of claims. Nontrivial work must be done (and "STSSD" shows that it can be done, and how) to turn the notion of predicate as equivalence class of substitutionally variant sentences, defined here, into the full-blooded notion of a cross-referenced predicate, as will be required for the introduction of quantifiers. Appendix I at the end of this chapter discusses some related points.

31. Truth-preserving and assertional commitment-preserving inferences (whether they are substitution inferences or not) include deductive inferences but define a wider category, for there is no implication in their case that the goodness of such inferences must be underwritten by their *form* (never mind their specifically *logical* form). Justification-preserving and, more generally, entitlement-preserving inferences include inductive inferences, but once again define a wider category.

32. It should not be thought that all goodnesses of inference must conform to the preservation model, in that there is a kind of status such that the inference is good iff the conclusion has the same status as the premises (any more than it should be thought that all good inferences have some sort of substitutional goodness). The notion of "transmission" of status is intended to indicate that the possession of a certain status by the premise (for instance, that S is assertionally committed to it) guarantees or provides the reason for the possession of

that status by the conclusion. The remarks in the text apply to commitment-preserving inferences (the genus of which deductive inferences are a species), but it should be noted that they need not apply to *entitlement*-preserving inferences (the genus of which inductive inferences are a species). I am grateful to Ernie LePore for pointing this out.

33. The restriction to substitution inferences is required because one may, for instance, infer asymmetrically from the applicability of a singular term to the applicability of a predicate—from "The inventor of bifocals is Benjamin Franklin" to "The inventor of bifocals is an American." These do not count as substitution inferences, even in the extended sense allowing replacement of frames, because they cross syntactic categorial boundaries.

34. Sortals, such as 'dog' and 'mammal', might seem to contradict this claim. For they are distinguished from predicates precisely in having associated with them not only criteria of application but also criteria of identity, and yet they can be materially involved in weakening inferences: "Wulf is a dog, so Wulf is a mammal." But their criteria of identity apply not to substitutions materially involving the sortals themselves but to those materially involving the singular terms to which the sortals are applied.

35. A fuller account would address incompatibility relations, as well as strictly inferential ones. These are made explicit in quantified negations: $(x)[Px \rightarrow \sim Qx]$. A story along these lines is worked out in Mark Lance's "Normative Inferential Vocabulary: The Explicitation of Social Linguistic Practice" (Ph.D. diss., University of Pittsburgh, 1988).

36. This following according to a general rule, albeit one that depends on the particular content nonlogical expressions are taken to have, is the source of the *modal* force or flavor that even material proprieties of inference exhibit. For it provides a sense in which a particular inference can be seen to be an instance of a valid type, even though what is quantified over is substitutional situations. This is a point about material inferences (here, the substitutional species of that genus) that can otherwise appear puzzling, since it is clear that the sort of inferences instanced in the text need not be good ones with respect to other possible worlds.

37. Of course what is at issue here is an inferentialist version of the distinction between extensional and nonextensional (or transparent and opaque) occurrences of, typically, singular terms, as discussed in Section II.

38. It need not be denied that occurrences whose significance is not governed in this way are semantically significant in a secondary sense, which can be explained only once the primary sense is understood. This is discussed further along.

39. These examples can represent the asymmetries only at the level of sentences. Singular terms *do not* behave asymmetrically, so real examples of asymmetrically behaving substituted-fors are not forthcoming. Probably the closest one can get in real grammar is sortals. Since they have associated with them criteria of identity for the singular terms they qualify, they are more termlike than predicates. Yet they do have proper inclusions, and a straightforward notion of inferential weakening applies to them, as to predicates. (The objection may now occur that these examples show that expressions like predicates, whose occurrences *do* have asymmetric significances, *can* occur embedded in inferentially inverting contexts, showing that something must be wrong in the analogous

argument to the conclusion that substituted-for expressions must have symmetric substitution-inferential significances. This legitimate worry is addressed further along, where the distinction between basic subsentential expressions, which can be substituted for, and derived subsentential expression-patterns [frames, of derived substitutional category], which can only be replaced [as outermost, hence never embedded], will be invoked.]

40. Recall that to take it that q is incompatible with p is to take it that commitment to q precludes entitlement to a commitment to p. In this way acknowledgments of material incompatibilities are implicit in the practices governing adopting attitudes (for instance, undertaking or attributing) toward the same pragmatic statuses of commitment and entitlement that inferences can be distinguished as preserving.

41. In conversation, Ken Manders has suggested the language of projective geometry as an example that is interesting in this connection. Sometimes 'general points' are appealed to, whose projective properties form a proper subset of the projective properties of some other point or points and so are asymmetrically inferentially related to each other the way sortals can be. The language in which projective properties are specified does *not* have negation or the conditional in it. The present argument explains the unobvious connection between these facts.

42. Indeed it could be argued that possession of this reflexive expressive capacity and all that goes with it makes so much difference that it provides a plausible place to draw the line between the linguistic and the nonlinguistic. The line between logical and prelogical languages is in any case important enough that researchers investigating what sorts of languages chimps and dolphins can be taught would be well advised to postpone trying to teach them an extra two hundred terms and predicates, and instead try to teach them to use conditionals and quantifiers.

43. Notice that this characterization of the conclusion could be accepted even by someone who was not persuaded by the expressive approach to understanding the demarcation of specifically logical vocabulary and so the function of logic.

44. Sentences displaying multiplace predicates in which the same term occupies several distinct argument places may seem to be an exception. A general technique for recognizing adicities in purely substitutional terms in spite of such cases is offered in "STSSD." There also are predicates, such as '. . . carried the piano up the stairs', that can be understood as taking variable numbers of arguments ('Bruno', 'Bruno and Betty', 'Bruno, Betty, and Bill' . . .). The treatment of these is complex, but their existence does not affect the general point being made here about the discriminability of substituted-fors and substitution frames.

45. Frege, in the *Grundgesetze*, forbids substituted-in expressions that cannot be substituted for. That is, he insists that functions must have "complete" values. There is no technical reason for this restriction, and in the present context, no philosophical reason either.

46. I am grateful to Igal Kvart for raising this possibility.

47. This example is due to Bill Lycan.

48. Walter Edelberg has argued ("Intentional Identity" [Ph.D. diss., University of Pittsburgh, 1984]) that there is a further class of cases of "asymmetric identities"

involving singular terms. These are very interesting cases, but they are not relevant to the present point, since they essentially involve embeddings, indeed multiple embeddings, in propositional attitude constructions. Thus the occurrences of singular terms involved in the Edelberg examples are not primary semantic occurrences, as here defined. In any case, these cases are also sensitive to the particular sentence frames involved, and so also do not generalize in the way required for government by asymmetric SMSICs.

49. This requirement suffices to rule out cases corresponding to the 'cat' in 'cattle'—though it would have to be more carefully stated, to acknowledge the fact that well-formedness *is* preserved in this context by *some* substitutions for 'cat', for instance 'bat' and 'rat'.

50. "To begin with" because these equivalence classes will not without further conditions on the substitution relations respect cross-identifications of arguments. "STSSD" shows one way in which one might go on to get the fully individuated complex predicates needed to codify material inferences essentially involving the contents of predicates. See also Appendix I at the end of this chapter.

51. Defining a replacement relation on frames, derivative from substitution for basic expressions, involves defining an isomorphism from one equivalence class to another, preserving these sorts of substitution. Thus to replace '*P*αβ' by '*Q*αβ' requires defining a mapping *h* from {*s*/*s* is in '*P*αβ'} into {*s*/*s* is in '*Q*αβ'} that is an isomorphism with respect to the substitution relations among the sentences within those two sets. To say that *s*, *s'* are in the same set '*P*αβ' is to say that there is some substitution of terms for terms that turns the one into the other. Those substitution relations are indexed by sets of pairs of terms. For instance, the relation indexed by {⟨'Carlyle', 'Ruskin'⟩, ⟨'Hegel', 'Schopenhauer'⟩} turns "If Carlyle wrote *Sartor Resartus*, then Carlyle had an ambiguous relationship with Hegel" into "If Ruskin wrote *Sartor Resartus*, then Ruskin had an ambiguous relationship with Schopenhauer." The isomorphism must be such that if *s* can be turned into *s'* in '*P*αβ' by a mapping Sub with an index of *t*, then *h*(*s*) can be turned into *h*(*s'*) in '*Q*αβ' by the mapping Sub with the same index *t*. Only equipped with such a mapping can one replace "If α wrote *Sartor Resartus*, then α had an ambiguous relationship with β" by "If α wrote *Stones of Venice*, then α was influenced by β." No such isomorphism across frames (only the indexing of Sub relations within frames) is required to define substitution of one singular term for another. Appendix I expands on this topic.

52. "By and large" because terms can contain other terms. (Indeed, this condition can be used to diagnose the occurrence of terms in which other terms occur, as it is in "STSSD.") Thus although both 'Carlyle's friend' and 'Carlyle' are terms occurring in "John Stuart Mill was Carlyle's friend," substitution for either in general eliminates the occurrence of the other. (Only "in general" because of special cases such as substituting 'Carlyle's editor' for 'Carlyle's friend'.)

53. This is the sense, explained above in Section II, subsection 2, according to which the fundamental sentential logical vocabulary must be 'extensional' in this sense.

54. Negated sentence frames must be discerned if the negation locution that is used to make explicit material incompatibility relations among sentences is to be able also to make explicit material substitution-incompatibility relations among sentence frames. An example of the latter would be what is made explicit by the

claim $(x)[Px \rightarrow \sim Qx]$, that is, that nothing that has property P has property Q. The sets of SMSICs that govern frames must include commitments regarding incompatibility as well as those regarding inferential relations among frames. Where the pragmatic significance of sentential inferential relations consists in preservation of some deontic status, whether commitment or entitlement, the pragmatic significance of sentential incompatibility relations consists in commitment to one content precluding entitlement to the other.

55. Though Appendix I shows how to define the second-order role of frames anyway.

56. Notice that this procedure is just a way of keeping track of the interactions between the two fundamental sorts of inverting contexts—negations and antecedents of conditionals. It is a technical device whose use for these purposes does not involve commitment to its being proper for *other* purposes to analyze $p \rightarrow q$ as $\sim p \lor q$.

57. This objection is developed from a suggestion originally offered by Larry Sklar, to whom much thanks.

58. Or the singular terms can be individuated by the transformations. This is the route taken in "STSSD." Appendix I offers some help with this point.

59. *Tractatus*, 5.62.

60. See the general discussion above at 3.1.3 of the significance of the distinction between associating semantic contents by stipulation and by conferral.

7. Anaphora

1. *GL*, sec. 60.

2. This can be only a preliminary discussion of the topic, for Frege does not explicitly take account of the *social* dimension of language use, which, it will be argued in Chapter 8, is essential for a successful account of the representational dimension of discursive practice. (As will emerge, however, he does *implicitly* acknowledge the social dimension in his concern with the epistemic *fruitfulness* or information content of identity claims.)

3. *GL*, sec. 84.

4. Ibid. On the use of 'idea', see for instance the footnote to sec. 58.

5. Ibid., sec. 51.

6. Ibid., sec. 57.

7. Ibid., sec. 74n.

8. All phrases from ibid., sec. 62.

9. See, for example, ibid., sec. 56, 104.

10. Ibid., sec. 74n. Frege uses 'concept' as the semantic correlate of predicates—once the distinction is firmly in place, in the 1890 essays, as what their referents are rather than as the senses they express. This usage derives from Kant and, partly due to Frege's endorsement, still has considerable currency. It is not a harmless or philosophically neutral choice of terminology, however. John McDowell, in "De Re Senses," *Philosophical Quarterly* 34, no. 136 (July 1984): 283–294, indicates some of the reasons why not. In the context of the present project, whatever is inferentially articulated is conceptually articulated; since the use of singular terms, like that of predicates, is determined by substitution inferences, singular terms are conceptually articulated every bit as much as predicates are. The discussion of anaphoric inheritance of the determinants of substitution-inferential commitments, in the latter part of this chapter, shows how even

deictic tokenings get to be inferentially, and so conceptually, articulated. The way in which such an approach overcomes the dualism of cause and concept about which McDowell complains is addressed below in 9.1.

11. *GL*, sec. 74n.
12. Ibid., sec. 62.
13. Ibid., sec. 84.
14. Ibid., sec. 56.
15. Ibid., sec. 46.
16. Ibid., sec. 106.
17. If in addition the singular term occurs essentially in some *true* claim, then it not only *can* be a way an object is given to those who use the term, it also *is* a way an object is given to them. See the discussion of existential commitments in Section II of this chapter.
18. *GL*, sec. 62.
19. Ibid., sec. 109.
20. Which at this stage in the development of Frege's terminology is equivalent to requiring that they have a determinate inferential role.
21. *GL*, sec. 76.
22. Ibid., sec. 62 (section heading).
23. Ibid., sec. 63.
24. Ibid., sec. 104.
25. Ibid., sec. 65.
26. Ibid., sec. 107.
27. As required, for instance, in the passage quoted above from ibid., sec. 84.
28. As required by the passage quoted above, from ibid., sec. 62: "If we are to use a symbol *a* to signify [*bezeichnen*] an object, we must have a criterion for deciding in all cases whether *b* is the same as *a*."
29. Introducing such terms is accordingly introducing the objects they refer to as well. This is the reason for Frege's otherwise peculiar talk in this context (for example in the footnote to *GL*, sec. 74) about defining *objects* rather than just terms.
30. Ibid., secs. 64, 65.
31. Ibid., sec. 68. Frege uses extensions of concepts rather than collections of objects. Though this difference is important to him for various reasons, it can be ignored for present purposes. Note 34 below indicates some of the subtleties that are being ignored here.
32. Ibid., secs. 63, 104.
33. Ibid., sec. 65.
34. This is of course a tendentious description of what Frege says, one that specifies a lesson that is only implicit in his discussion. For the formulation offered conflates the second definition Frege considers (in the crucial substantive section of the *Grundlagen* that comprises secs. 62–69), which he explicitly rejects, with the third, which he endorses. The reason for the conflation is to avoid having to discuss idiosyncratic features of Frege's notion of concepts that are otherwise irrelevant to the points being extracted here from his treatment of objects. (Admittedly, disentangling these two strands of thought is a delicate matter, given how tightly bound up the two categories are for Frege.) The cardinal point of divergence from Frege's avowed position concerns the question of whether the sense of a numerical identity can be regarded as having been

settled, and so numerical expressions introduced as ways in which objects are given to us, when the truth-values of sortally heterogeneous identities (in the general case, those in which only one of the flanking terms is of the form *fa*) have *not* been fixed, while those of the sortally homogeneous identities (in the general case, those of the form *fa* = *fb*) have been. Frege indeed explicitly denies this; it is the reason he gives for rejecting his second definition and replacing it by his third: "This means does not provide for all cases. It will not, for instance, decide for us whether England is the same as the direction of the Earth's axis" (sec. 66). The justification for nonetheless taking the position pursued in the text as in a real sense implicit in Frege's discussion is that Frege never finds a satisfactory solution to the problem of fixing the sense of sortally heterogeneous identity claims—neither in the third and final definition endorsed in the *Grundlagen* nor in his subsequent writings, culminating in the crucial introduction of courses of values in sec. 10 of the *Grundgesetze*. The history of his attempts to come to grips with this issue, and the argument (requiring too many details to be rehearsed here) that those attempts ultimately fail, are presented in the author's "Frege's Technical Concepts," in *Frege Synthesized: Essays on the Philosophical and Foundational Work of G. Frege*, ed. L. Haaparanta and J. Hintikka, Synthese Library (Dordrecht: D. Reidel, 1986), pp. 253–295. The lesson extracted below follows from Frege's discussion using only the auxiliary hypothesis that, for the reasons offered in the essay just cited, the truth-values of the sortally heterogeneous identities *cannot* be settled by stipulation when terms referring to novel abstract objects are introduced.

35. See "Frege's Technical Concepts."
36. All involving the same predicate, that is, sentence-frame. This is a somewhat loose formulation of the situation he is ruling out, but it is a straightforward matter to use the vocabulary developed in Chapter 6 to make it precise.
37. *GL*, sec. 67.
38. Ibid., sec. 76.
39. Like the 'cat' in 'cattle', to use Quine's example.
40. It was observed above that if all the antecedent terms have a complete set of identities associated with them, introducing a new term by stipulating one nontrivial identity relating it to a term in the antecedent vocabulary suffices to associate with it a complete set of identities. However, if their senses are in this way incomplete (*not*, it is being insisted, indefinite)—if each has associated only *some* nontrivial identities and nonidentities—then the new term acquires in this way also only an incomplete sense.
41. Sections III–V of this chapter will show that anaphoric structures essentially involve a corresponding holism.
42. "The Conditions of Thought," in *Le Cahier du College International de Philosophie* (Paris: Editions Osiris, 1989), pp. 165–171, reprinted in *The Mind of Donald Davidson*, ed. J. Brandl and W. Gombocz (*Grazer Philosophische Studien* 36 [1989]), pp. 193–200.
43. In *Knowledge and the Flow of Information* (Cambridge: MIT Press, 1981).
44. Since the two routes intersect in two places, it is necessary to think of arcs rather than lines as intersecting, and the term 'triangulation' may be less than happy.
45. See Appendix II to Chapter 6.
46. *Intending* to use an expression as a singular term is adopting a practical com-

mitment to doing what is necessary to make it appropriate for oneself and others to adopt such a stance toward one's performances. How to understand the scorekeeping attitudes such a stance consists in and the proprieties that govern adopting it are accordingly issues that must be addressed before such intentions can be made intelligible. That the order of explanation in this way dictates that semantic intentions not be appealed to as fundamental at this point is just one particular instance of the difference of explanatory strategy, insisted upon elsewhere, that divides the present approach from that of agent semantics.

47. *GL*, sec. 74n.
48. Ibid.
49. Ibid., sec. 95.
50. Substitutional construals of quantification are defended against this and other objections, and shown to be not technically inferior to objectual construals, in J. Camp, "Truth and Substitution Quantifiers," *Nous* 9 (1975), and S. Kripke, "Is There a Problem about Substitutional Quantification?" in *Truth and Meaning*, ed. G. Evans and J. H. McDowell (Oxford: Clarendon, 1975).
51. Geach discusses some of the reasons for treating sortally restricted quantification as more fundamental than unrestricted quantification in *Reference and Generality* (Ithaca: Cornell University Press, 1962), for instance at pp. 150–151.
52. This claim, and other similar ones used as examples at various points in the text, are taken from David Wells's *Penguin Dictionary of Curious and Interesting Numbers* (1986).
53. *GL*, sec. 62.
54. Ibid., sec. 22.
55. Ibid., sec. 54. See also his criticism of appeals to 'unit' in place of substantive sortals, at secs. 25–26, 33, and 44.
56. Ibid., sec. 54.
57. A. Gibbard, "Contingent Identity," *Journal of Philosophical Logic* 4 (May 1975): 187–221.
58. For this example and much other wisdom about sortals and sortally restricted singular terms, see A. Gupta, *The Logic of Common Nouns: An Essay in Quantified Modal Logic* (New Haven: Yale University Press, 1980).
59. Geach initiated discussion of this topic. See for instance his "Ontological Relativity and Relative Identity," in *Logic and Ontology* ed. M. K. Munitz (New York: New York University Press, 1973).
60. For the origin of this phrase, see below at 8.5.1.
61. Strictly, it is an abbreviation for one. But so is $100 \cdot 365 \cdot 24 \cdot 60 \cdot 60$. Just as for some purposes the difference between true successor numerals and their decimal abbreviations can be ignored, so for others can the difference between their decimal and their more complex arithmetic abbreviations. This is a matter of what assumptions one is willing to take for granted about various operations in treating expressions as canonical, and accordingly varies with the context of inquiry. The commitments involved in treating the power-set successor sequence as forming canonical designators of transfinite cardinals, and so the utility of such an attitude, appeared differently before the demonstration of the independence of the continuum hypothesis than they did afterward.
62. This account of physical existence in terms of coordinates ultimately mappable onto egocentric space is meant to piggyback on the sort of story Gareth Evans

tells about the relations between public and egocentric specifications in his discussion of demonstratives in Chapter 6 of his *Varieties of Reference*, ed. John McDowell (New York: Oxford University Press, 1982).

63. The use of the expression 'I', which anchors these trajectories, is discussed below in 8.5.2.

64. Strictly, specifications of *regions* are required, but this complexity can safely be suppressed here. Thus the point-specifications offered here should be understood as standing in for the much more complex coordinate specifications of space-time regions.

65. Hard cases, such as beams of light and noises, will not be dealt with here—only paradigms of physical existents. These others often need to be individuated by further physical-dynamic variables. See the discussions in Gibbard, "Contingent Identity," and Gupta, *The Logic of Common Nouns*.

66. Lakatos's fascinating case studies in *Proofs and Refutations*, ed. J. Worral and E. Zahar (Cambridge: Cambridge University Press, 1976), indicate how various candidates for canonical designators of polyhedra suffered this fate.

67. This account leaves entirely open the question of whether and in what sense various *properties* might be said to exist. One could take a Fregean line and restrict the notion of existence to objects, while denying that properties can be picked out by singular terms. If the introduction of such terms is admitted, then the problem becomes one of determining an appropriate class of canonical designators.

68. This is a consequence of what Gareth Evans calls the "Generality Constraint" (*Varieties of Reference*, sec. 4.3).

69. Schröder, in an early misguided attempt at greater rigor, includes in his algebraic axioms, along with commutativity and associativity of his operators, the stipulation that wherever in his text he uses a symbol such as '×', it is to be understood to have the same meaning. It should be clear that this principle ought to be accorded quite a different status from the others; for instance, commutativity makes sense apart from associativity, but what sense do the statements of those conditions make independently of Schröder's terminological stipulation?

70. Although holding the inheritance relation constant, if the SMSICs governing the second were different, those governing the first would have had to be different; but this is a different counterfactual. The latter resembles the claim that if Oswald *didn't* shoot Kennedy, then someone else *did* (since someone shot him), while the former resembles the claim that if Oswald *hadn't* shot Kennedy, then someone else *would* have.

71. This is not to deny that in some cases (for instance where syntactically definite descriptions are used as anaphoric dependents) what one takes the recurrence class to be may depend on what substitutional commitments one acknowledges.

72. Although the examples considered here have been in the grammatical category of singular terms, it should be clear that the account of anaphora as inheritance of substitutional commitments is not so restricted but applies to sentences, predicates, sortals, and so on, as discussed in Chapter 5.

73. The Appendix to this chapter mentions some of the otherwise important issues about anaphora that are for this reason *not* discussed here.

74. Indeed, the anaphoric antecedent may be a merely *possible*, rather than an actual, tokening. This possibility was adverted to in Chapter 5 in connection

with the way modal operators combine with anaphoric antecedents formed using the '. . . refers . . .' operator. The status of merely possible tokenings is discussed further at 8.4.2–3.

75. John Perry, "The Problem of the Essential Indexical," *Nous* 13 (1979): 3–21; David Lewis, "Attitudes *De Dicto* and *De Se*," *Philosophical Review* 88 (1979): 513–543; David Kaplan, "Demonstratives," in *Themes from Kaplan*, ed. J. Almog et al. (New York: Oxford University Press, 1989). Some of these issues are discussed below at 8.5.2–4.

76. Evans, *Varieties of Reference*, especially chap. 6.

77. For the sake of definiteness, the discussion here addresses the demonstrative use of singular terms. Similar remarks apply to demonstratives in other grammatical categories, with suitable adjustments for the differences in what sort of substitutional commitments govern the use of those expressions.

78. Insisting that demonstratives or indexicals put us in cognitive touch with the world only insofar as they can be hooked up to the rest of our conceptual apparatus in this way is the point of the discussion of "Sense Certainty" with which Hegel opens the *Phenomenology*.

79. This is to say that expressions of the form /The K S demonstrated while saying $\langle t \rangle$/ should be understood by analogy to the anaphorically indirect definite descriptions analyzed above in 5.4. They are not descriptions but anaphorically indirect demonstrations. Such expressions pass the iteration test proposed there for a suitably broad notion of demonstration—one that allows pointing at an image in a mirror, photograph, or video screen to count as a demonstration of the object imaged.

80. In a certain sense, however, it is not the case that *all* anaphoric token-recurrence structures exhibit modal rigidity. The telling exceptions might be called 'impure' or 'descriptive anaphoric structures'. Thus anaphoric dependents that are indirect definite descriptions do *not* behave rigidly (as is pointed out above in 5.4.5). Leibniz might not have been so-called: it is possible that the one referred to as 'Leibniz' not be Leibniz. So 'the one referred to as Leibniz', though in fact anaphorically dependent on a tokening of 'Leibniz', need not be coreferential with it. These cases are complicated, but they can be accounted for by assembling the raw materials presented here. The important thing is to see why they do not falsify the account offered in the text, for in the counterfactual situations being envisaged, various expressions *would* belong to *different* token-recurrence structures than they in fact do. Such cases are parasitic on the basic ones discussed here, and a full treatment of their intricacies would have to appeal to the phenomena described here in the base-level discussion of pronouns, descriptions, and rigidity. A complementary phenomenon is the anaphora that underlies the use of *reflexive* constructions. Here counterfactual mutation of the recurrence structures is forbidden; we cannot coherently specify situations in which Leibniz would not be *himself*.

81. For instance in "On the Logic of Demonstratives," *Journal of Philosophical Logic* 8 (1979): 81–98, reprinted in *Propositions and Attitudes* ed. N. Salmon and S. Soames (New York: Oxford University Press, 1988); and "Dthat," in *Syntax and Semantics*, ed. P. Cole (New York: Academic Press, 1978), reprinted in *Demonstratives*, ed. P. Yourgrau (New York: Oxford University Press, 1990).

82. So long as only pure or nondescriptive anaphoric links are involved.

83. "Naming and Necessity," in *Semantics of Natural Language*, ed. G. Harman and D. Davidson (Dordrecht: D. Reidel, 1972).

84. This suggestion was made by Chastain already in the late 1960s, at the outset of the contemporary discussion of causal-historical theories of reference (in his "Reference and Context," in *Language, Mind, and Knowledge*, Minnesota Studies in the Philosophy of Science, vol. 7, ed. Keith Gunderson [Minneapolis: University of Minnesota Press, 1975], pp. 194–269), but it has not been pursued.

85. A good introduction to the recent linguistics literature on anaphora is Joseph Aoun's *Grammar of Anaphora* (Cambridge: MIT Press, 1986). A very useful philosophical discussion of anaphora that (like the current approach) does not depend on a fundamental distinction between intrasentential and discourse anaphora is J. Hintikka and J. Kulas's *Anaphora and Definite Descriptions* (Dordrecht: D. Reidel, 1985).

86. Interpretation in this sense should be sharply distinguished from Davidsonian interpretation, which is allegedly the same within languages and across languages. Interpretation in that sense requires explanatory hypothesis formation and inferences starting from the *noises* another utters. Against this it has rightly been objected that sharing a language involves being able to "hear another's meanings" rather than another's noises. At issue here is the subcapacities involved in being able to do that. This is discussed further at 8.2.1–2.

87. Or one could conjoin with each of p and q a set of auxiliary hypotheses of the form $(p$ or $q) \rightarrow r$. Clearly versions of this strategy will apply to *theories* T_1 and T_2.

88. Frege had made the same point (distinguishing by their lack of inferential consequences cognitively trivial identities from the substantive ones that give us cognitive access to an object) already in the *Grundlagen*, in the passage from sec. 67 quoted above in Section I, Subsection 5.

89. Both quotations are from *Logic*, by Immanuel Kant, trans. R. S. Hartman and W. Schwarz (New York: Dover Publications, 1974), pp. 45 and 118 respectively.

90. I think that putting this word in is a slip on Kant's part (it is not a problem with the translation—the original does say that nothing further [*weiter*] is then known). The *Logik* consists of published lecture notes, and Kant did not revise or edit it as carefully as he did some of his other works. On Frege's view, as rehearsed in Section I, and also according to Kant's own official view, this sentence should end: ". . . I know nothing of him at all."

91. As long as the expressions occurring in the tautology also have nontautologous occurrences (that is, also occur in claims that *are* rich in inferential consequences), they have a sense—both for Kant and for Frege.

92. For present purposes it is possible to ignore the involvement of a claim in empirical entries and practical exits—perception and action—which add noninferential (but still norm-governed) circumstances and consequences of application and hence contribute to its content.

93. See for instance his *Against Method: Outline of an Anarchistic Theory of Knowledge* (London: Humanities Press, 1975).

94. Addressed in Israel Sheffler's *Science and Subjectivity* (Indianapolis: Bobbs-Merrill, 1967).

95. The difference between committive and permissive inference is being suppressed here.

96. Reprinted as Chapter 14 of David Lewis's *Philosophical Papers*, vol. 1 (New York: Oxford University Press, 1983); originally appeared in *Language, Mind, and Knowledge*, Minnesota Studies in the Philosophy of Science, vol. 7, ed. Keith Gunderson (Minneapolis: University of Minnesota Press, 1975), pp. 3–35.

97. Locus classicus is "Concepts As Involving Laws, and Inconceivable without Them," in *PPPW*, pp. 87–124.

98. For that reason it cannot be stated as a series of theses about what Fregean senses really are, and how they ought to be understood to be related to their referents. Compare the more orthodox multileveled schemes that Lewis considers in "Tensions," reprinted as Chapter 14 in his *Philosophical Papers*; originally in *Semantics and Philosophy*, ed. M. K. Munitz (New York: New York University Press, 1974).

99. So in spite of the absence of shared intensions, we are able to understand Rutherford, to extract information from his claims, and to agree or disagree with him by using the same interpretive strategies we use to cope with each other— by finding out what the other is talking about, how the other is representing things as being. Concrete scorekeeping practices do the duty of abstract intension-functions. The question of what it is for one rather than another set of finely individuated contents to be conferred on expressions and performances by their use is then traded for the question of what it is for a community to engage in one rather than another set of scorekeeping practices.

100. More precisely: it *cannot* arise for the case of intra*sentential* anaphora, for it can occur even in the intrapersonal case—for instance if there is sufficient lapse of time for the individual in question to have altered commitments.

101. K. Donnellan, "Reference and Definite Descriptions," *Philosophical Review* 77 (1966): 281–304; S. Kripke, "Speaker Reference and Semantic Reference," in *Contemporary Perspectives in the Philosophy of Language*, ed. P. A. French, T. E. Uehling, Jr., and H. K. Wettstein (Minneapolis: University of Minnesota Press, 1977), pp. 6–27.

102. Gareth Evans, "Pronouns, Quantifiers, and Relative Clauses," part 1, *Canadian Journal of Philosophy* 7 (1977): 467–536; reprinted in Evans, *The Collected Papers* (Oxford: Clarendon Press, 1985).

103. Stephen Neale, *Descriptions* (Cambridge: MIT Press, 1990), esp. pp. 170–183, whose account is followed in many respects in the discussion below.

104. Ibid., p. 171.

105. Neale draws a corresponding conclusion at ibid., pp. 182–183.

106. Evans, "Pronouns." Neale offers this formulation of the notion of c-command: "A phrase α c-commands a phrase β if and only if the first branching node dominating α also dominates β (and neither α nor β dominates the other)" (*Descriptions*, p. 173).

107. In Geach, *Reference and Generality*, p. 117; and in his "Referring Expressions Again," in *Logic Matters* (Berkeley: University of California Press, 1972).

108. L. Kartunnen, "Definite Descriptions and Crossing Coreference," *Foundations of Language* 7 (1971): 157–182.

109. Neale, *Descriptions*, p. 197.

8. Ascribing Propositional Attitudes

1. Jefferson is reported to have said about a reported sighting by a Yale geologist that it is easier to believe that Yankee professors lie than that stones fall from the sky.
2. See for instance E. Sosa, "Propositional Attitudes *De Dicto* and *De Re*," *Journal of Philosophy* 67 (1970): 883–896.
3. The account offered in what follows can be understood as adapting this standard analysis of the two readings (in terms of scope) to the methodological setting provided by deontic scorekeeping accounts of linguistic-social practice. The different repertoires of discursive commitments associated with the one under-taking a commitment and the one ascribing it provide the different contexts of evaluation.
4. The claim expressed by this sentence, or its French equivalent (though Voltaire spoke excellent English).
5. "Quantifiers and Propositional Attitudes," *Journal of Philosophy* 53, no. 5 (1956): 177–187; reprinted in Quine's *Ways of Paradox* (New York: Random House, 1966), pp. 183–194, and in *Reference and Modality*, ed. L. Linsky (New York: Oxford University Press, 1971).
6. Dennett argues for this way of thinking of the distinction in "Beyond Belief," in *Thought and Object*, ed. A. Woodfield (New York: Oxford University Press, 1982). Those who are inclined to think instead of *de dicto* and *de re* as two kinds of belief should be given pause by the fact that on this line it would seem there must be not just two sorts but a potentially infinite number of kinds of belief. For ascriptional operators iterate, so that S may believe of t that S' believes that $\Phi(it)$, or believe that S' believes of t that $\Phi(it)$, and these "mixed" ascriptions can themselves be embedded in turn in further ascriptions of either the *de re* or *de dicto* sort. These kinds of iterated ascriptions are discussed in more detail in the Appendix to this chapter.
7. Notice that 'that individual' here is an anaphorically dependent expression. The need to use anaphora in expressing this relation is reflected in the ascription-structural anaphora that connects an antecedent in the scope of the 'of' to a dependent in the scope of the 'that' in regimented *de re* ascriptions. This connection is a manifestation of the deep point that *de re* ascriptions make explicit the sort of *communication* that has already been picked out (in Chapter 7) as implicitly secured by interpersonal anaphora. This point is discussed further below.
8. From a technical point of view, as will become clear in the next section, this proposal just transposes into the current idiom a suggestion that Kaplan offers and then rejects early on in "On Quantifying In," in *Words and Objections: Essays on the Work of W. V. Quine*, ed. D. Davidson and J. Hintikka (Dordrecht: D. Reidel, 1969), also in *Reference and Modality*, ed. L. Linsky (New York: Oxford University Press, 1971), pp. 112–144. The immense influence of this essay, which secured the (from the point of view of this work misleading, indeed disastrous) emphasis on epistemically strong *de re* ascriptions as the phenomenon of most philosophical interest, is largely responsible for what is here taken to be the real philosophical significance of *de re* ascriptions having been over-looked.

9. Davidson, for instance, is criticized on these grounds in Ian Hacking's "Parody of Conversation," in *Philosophy of Donald Davidson: A Perspective on Inquiries into Truth and Interpretation,* ed. E. LePore (Oxford: Basil Blackwell, 1986), pp. 447–458, and in Stephen Mulhall's "Davidson on Interpretation and Understanding," *Philosophical Quarterly* 37, no. 148, (July 1987): 319–322.

10. *PI,* 201.

11. Only some claims are rules, for rules are a kind of normative claim (discussed in Chapter 4).

12. The phenomenology is not really to the point. The claim here is that consciousness in an important sense *consists in* the capacity to keep score, make substitutions, and so on. So on pain of an infinite regress, these must be things that can be done without conscious deliberation or rehearsal.

13. Iteration of ascription-forming operators produces further varieties as combinations of these basic ones, since a *de re* ascription can be embedded in a *de dicto* ascription ("*S'* believes that *S* believes of *t* that Φ{it}") and so on. The way in which the account of the basic forms is to be extended to these cases is discussed in the Appendix to this chapter.

14. Assuming that it is appropriate to treat intentions (or at least the theoretical version of them: practical commitments, commitments to act) as having *propositional* contents. This is not obviously correct. The primary locution for expressing and ascribing intentions in English is intending *to* do something rather than intending *that* something be true. The relations between them are complex, and pursuing them here would require more detail than is appropriate given the broad-brush strokes in which the rest of the treatment of intention is painted. Roughly, the idea is this. 'Intend' is a verb whose complement is restricted to agentives. It can be presented in a regimented language as taking exclusively propositional complements provided these are restricted to propositions whose outermost operator is an agentive-explicitating operator, of which the paradigm is Belnap and Perloff's STIT. (See M. Perloff and N. Belnap, "Seeing To It That: A Canonical Form for Agentives," *Theoria* 54 [1988]: 175–199; reprinted, with corrections, in *Knowledge, Representation, and Defeasible Reasoning,* Studies in Cognitive Systems, vol. 5, ed. H. E. Kyburg, Jr., R. T. Loui, and G. N. Carlson [Dordrecht: Kluwer Academic Press, 1990], pp. 167–190.) The canonical form of an ascription of intention is then something like:

> *i*: *S* intends that *a*(*S*) STIT(*p*),

where /*a*(S)/$_i$ is an ascription-structural anaphoric dependent whose antecedent is /S/$_i$. Indexicals and quasi-indexicals play a special role in such ascriptions, as discussed below in the section on strong *de re* ascriptions. The point being made in the text, which does not depend on these subtleties, emerges more clearly if they are suppressed in favor of a parallel treatment of ascriptions of doxastic and practical commitments as alike taking propositional content-specifying complements.

15. Requiring that the noninferential elicitation of the performance be the exercise of such a reliable differential responsive disposition (discussed for the case of perception in the first section of Chapter 4) obviates the difficulty otherwise raised by Davidson's nervous mountaineer. In offering this conceptual placeholder in lieu of Davidson's more general "caused in the right way" condition,

the theory-sketch presented here is more specific. In the absence of a discussion of the difficult problem of individuating skills and dispositions and sorting particular performances according to which are being exercised, this suggestion presents only the form of an account.

16. Compare the discussion of speaker-referring uses of definite descriptions—which pick out an object by a description that is not true of it.

17. This tendency is widespread. One relatively sophisticated representative of it is Colin McGinn's "Structure of Content," in *Thought and Object*, ed. A. Woodfield (New York: Oxford University Press, 1982).

18. In fact it is harder than generally acknowledged to tell a coherent story at all about states that are narrow in this sense. The canonical approach, through twin-earth examples, involves a relativity to admissible descriptive vocabulary whose significance is sometimes not understood. Narrow states are what would be shared by individuals as much alike as they could be, embedded in different total environments. But a difference in context is a difference in what is true of the individuals. They can be made out as twins only by restricting the vocabulary that is permitted to be used in distinguishing them. And there is no way to motivate privileging a vocabulary that will accomplish what is desired, without begging the question—the twins are not, for instance "atom for atom identical" when the vocabulary is restricted only to purely *physical* terms, for "being twenty feet away from 100 kilograms of H_2O," rather than XYZ, is a perfectly good physical predicate. Nor will it do to talk in terms of what the individuals can discriminate, for what they can respond differentially to depends on what counts as distinct responses, and the same problem arises for motivating restrictions on the vocabulary in which one specifies those responses. The challenge is to specify a vocabulary meeting the two constraints, first, that twins in different environments are indistinguishable by descriptions in that vocabulary (the language of physics or of purely physical properties of behavior will *not* do) and, second, that the underdetermination of semantic properties of their states by descriptions in that restricted vocabulary shows something interesting. The point is not that these constraints are impossible to satisfy but that doing so requires a more delicate touch than is often appreciated.

19. A worry one might have about this part of the program is that the account of the distinction between *de re* and *de dicto* in terms of social differences of inferential perspective corresponding to different repertoires of discursive commitments is tailored specifically to expressions of propositional attitudes, formed using some analog of a "believes that" sentential operator. But other sentential operators, most notably those expressing tense and alethic modality, also exhibit a distinction between *de re* and *de dicto* readings—a distinction that can usefully be construed as a difference in the scope of those operators. Since tensed constructions, necessity, and possibility are not among the locutions officially reconstructed as logical (explicative of discursive practice) in this work, there is a methodological excuse for not addressing this worry in detail. A word or two about how the present account might be extended to tense and modal operators is nonetheless in order. Each of these cases is like that of propositional attitudes in involving the adoption of different perspectives, specifying the content of a claim from the point of view of different repertoires of commitments. In the case of tense, these are commitments that have been or

will be adopted, rather than those adopted by different contemporary individuals. In the modal case, the repertoires are associated not with different interlocutors but with different possible situations. These can be construed as sets of facts—that is, true claims ('claims' in the sense not of claimings but of what is claimed). Modulo such differences in how the different repertoires of commitments that provide the different perspectives are conceived, the social-perspectival account of the distinction between *de re* and *de dicto* content-specifications offered here can be applied to the case of tense and modality, as well as to that of propositional attitudes, which provides the paradigm.

20. Of course it was claimed above that some of the components of propositional-attitude-ascribing constructions—paradigmatically the intentional 'of' or 'about' whose scope indicates *de re* uses of content-specifying vocabulary—play a *semantically* explicitating expressive role, making representational substitutional commitments explicit that would otherwise remain implicit in the interpersonal interpretive scorekeeping practices constitutive of communication. Here, as is the case with interpersonal anaphora and other phenomena involving hybrid deontic attitudes, it is impossible to keep semantic and pragmatic considerations distinct.

21. One need not have produced (nor have been disposed to produce) a performance explicitly acknowledging a commitment in order nonetheless to have undertaken it (a scorekeeper need not take one to have produced or been disposed to produce such a performance in order to attribute the corresponding commitment); some commitments are acquired consequentially in virtue of others one *is* prepared to acknowledge, and others may be acquired by default simply in virtue of one's *not* being prepared to disavow them. Yet such assertional performances are essential to discursive practice in the sense that the ideas of doxastic commitments undertaken consequentially and by default are unintelligible in abstraction from practices that include the possibility of avowing those commitments overtly.

22. Notice that even though one can find out about another's commitment-acknowledging performance by testimony rather than by perceiving it, doing so still requires perceiving the assertional performance that conveys the testimony.

23. This way of putting things is chosen to indicate that it is essential to ascriptional tropes that the one to whom a commitment is ascribed *may* be someone other than the ascriber. Apart from this possibility, the role of ascriptions as making deferrals explicit, for instance, is unintelligible. In the context of such an interpersonal practice, however, nothing rules out self-ascriptions. Section V (Subsection 2) discusses some aspects of self-ascriptions, in connection with the use of the word 'I', which makes explicit self-ascriptions when the self-attributions they express have the significance of *acknowledgments* of commitments. For not all attributions of a commitment to oneself are acknowledgments of it— they must, as it were, be attributions to oneself *as* oneself. So self-attribution cannot in general be identified with acknowledgment. Neither can it be identified with undertaking a commitment, for just as one can undertake a commitment without acknowledging it, one can undertake it without attributing it to oneself.

24. Quine's corner quotes are adopted here as a terminological convenience for expressing generalizations involving quotation. Thus $\ulcorner p \urcorner$ is to function as a

variable ranging over the results of applying quotation marks to the sentences that the variable p ranges over. (This fussy convention is needed because the other order of precedence in applying the two operations is standard for ordinary quotation marks: "p" is a quote-name of a letter of the alphabet.)

25. Davidson's piece was originally published in *Words and Objections*, ed. D. Davidson and J. Hintikka (Dordrecht: D. Reidel, 1969), pp. 158–174; reprinted in Davidson's *Inquiries into Truth and Interpretation* (New York: Oxford University Press, 1984), pp. 93–108. McDowell's essay is in *Reference, Truth, and Reality*, ed. Mark Platts (London: Routledge and Kegan Paul, 1980), pp. 206–237.

26. For some applications (see the next note), the account requires talk of merely *possible* tokenings. McDowell goes to some trouble to show that such an appeal should not be metaphysically or ontologically worrisome; possible tokenings can just be thought of as ordered triples of a lexical-syntactic sentence type, a speaker, and a time (and of course further indices, for instance specifying a possible world, could be added as needed).

27. "Would be" because the displayed sentence tokening is not used assertively and so is not a "saying" in the sense being reconstructed here. This formulation accordingly invokes a virtual tokening, of the same type as the displayed tokening in the 'that' clause, and in largely the same circumstances of evaluation as the ascribing tokening. This way of expressing the relation presents another avenue for avoiding the "subtle flaw" Davidson finds in more quotational treatments of indirect discourse, related to but distinct from the different strategies Davidson and McDowell put forward.

28. Which of course is not to say that it would be *successful*, for that would depend on whether the one to whom the claim is ascribed is entitled to it, and whether that entitlement is heritable by the ascriber.

29. "De Re Senses," *Philosophical Quarterly* 34, no. 136 (July 1984): 283–294.

30. Davidson does not consider the possibility of a variant that would assimilate the relation between 'that' and the sentence tokening that follows it to *anaphoric*, rather than *demonstrative*, constructions. Such a course is recommended by the fact that the relation involved is, like anaphora and unlike demonstration in general, intralinguistic. It may be noticed also that where saying that has the sense considered here, of saying *claimingly*, 'that' can be replaced by the properly prosentential 'that is true' in Davidson's account—Galileo said that is true: the earth moves. The upshot of the discussion here, though, is that neither of these heroic expedients is required by the essential insights of Davidson's approach.

31. Recall Field's approach, mentioned in Chapter 3, according to which S believes that p is analyzed into two parts: (i) S believes* σ (some sentence), and (ii) σ means that p.

32. Of course the sentential/propositional division is not exhaustive, and to mention these challenges is not to suggest that they cannot be met. A prominent approach, which Stalnaker calls "pragmatic," looks to intentional action to connect propositional contents with attitudes. In some forms, this approach appeals to decision-theoretic practical rationality in the individual case, and to conventions in the communal linguistic case. A primary source is David Lewis's "Languages and Language," in *Language, Mind, and Knowledge*, Minnesota Studies in the Philosophy of Science, vol. 7, ed. Keith Gunderson (Minneapolis:

University of Minnesota Press, 1975), pp. 3–35, reprinted in Lewis's *Philosophi-cal Papers*, vol. 1 (New York: Oxford University Press, 1983.)

33. Section V looks in more detail at communication that should be understood in terms of asymmetric constellations of tokenings bound together by anaphoric commitments. Section VI discusses in more general terms how propositional and conceptual contents should be thought of according to this model.

34. Recall the discussion above in 7.5 of the importance of anaphora in securing a common referent and so making communication possible.

35. Tyler Burge, "Belief *De Re*," *Journal of Philosophy* 74 (1977): 341, talking about a phenomenon previously noted by Castañeda and Loar.

36. Although for simplicity singular terms are used in all the examples here, it should be kept in mind that expressions of *any* grammatical category can be exported into the scope of the *de re* 'of'—not only predicates and sortals, but even whole sentences, as in "Senator McCarthy believed of the first sentence of the Communist Manifesto that it was true," in which the ascription-structural anaphoric dependent is a prosentence.

37. "Quantifiers and Propositional Attitudes."

38. The fact that what is reconstructed as corresponding to belief is the deontic status of doxastic commitment, construed as a normative social status conferred by linguistic scorekeeping practices, represents a difference from the traditional approach that can be largely ignored for the purposes of making the present point.

39. Sellars proposes a denotational view of what is expressed by *de re* belief ascriptions, in "Some Problems about Belief," in *Words and Objections*, ed. D. Davidson and J. Hintikka (Dordrecht: D. Reidel, 1969).

40. "On Quantifying In," first sentence of sec. 8.

41. The idea of *canonical designators* employed above in 7.2 in connection with the explication of existential commitments grows out of this idea of Kaplan's.

42. Sosa, "Propositional Attitudes *De Dicto* and *De Re*," p. 890.

43. Burge "Belief De Re"; John Perry, "The Problem of the Essential Indexical," *Nous* 13 (1979): 3–21, reprinted in *Propositions and Attitudes*, ed. N. Salmon and S. Soames (New York: Oxford University Press, 1988); David Lewis, "Attitudes *De Dicto* and *De Se*," *Philosophical Review* 88 (1979): 513–543; David Kaplan, "Dthat," in *Syntax and Semantics*, ed. P. Cole, (New York: Academic Press, 1978), pp. 221–243, reprinted in *Demonstratives*, ed. P. Yourgrau (New York: Oxford University Press, 1990), pp. 11–33; see also D. Kaplan, "Thoughts on Demonstratives," in ibid., pp. 34–49.

44. The essays in N. Salmon and S. Soames, eds., *Propositions and Attitudes*, present a good overview of this idea.

45. Dennett calls what is expressed by *de re* ascriptions on the denotational view "weak aboutness" and distinguishes it from "strong aboutness" in "Beyond Belief," p. 67.

46. One thinker who defends the not uncommon view that *de re* beliefs can be made sense of antecedently to any sort of *de dicto* beliefs is Fred Dretske, in *Knowl-edge and the Flow of Information* (Cambridge: MIT Press, 1981), esp. pt. 3.

47. Particularly helpful on this point is McDowell's criticism of Burge's formula-tions of the doctrine, in "De Re Senses." How one should think of the concep-tual in this connection is discussed further in the Conclusion of this work.

48. In "The Problem of the Essential Indexical" and "Frege on Demonstratives," *Philosophical Review* 86 (1977): 474–497. The earlier discussions are in Hector-Neri Castañeda's "'He': A Study in the Logic of Self-Consciousness," *Ratio* 8 (1966): 130–157, and "Indicators and Quasi-Indicators," *American Philosophical Quarterly* 4 (1967): 85–100, discussed below.

49. For prior intentions. In the case of intentions in action, the acknowledgment of a practical commitment that could be expressed by using 'I' (or 'now') is the nonlinguistic performance itself—fleeing the bear or heading for the meeting. See above at 4.4–6.

50. In *Mind and Language,* ed. S. Guttenplan (New York: Oxford University Press, 1975), reprinted in *Demonstratives,* ed. P. Yourgrau (New York: Oxford University Press, 1990), and in Anscombe's *Collected Philosophical Papers,* (Minneapolis: University of Minnesota Press, 1981).

51. Anscombe rejects the suggestion that 'I' is a pronoun, though not one that behaves just like any other, as "no good, because 'pronoun' is just a rag-bag category; one might as well say 'It is the word that it is.'" ("The First Person," p. 143; this and subsequent page references are to the Yourgrau collection). But this overstates the case, given that, as she acknowledges "'I' functions syntactically like a name" (p. 138)—that is, it plays the substitutional role characteristic of singular terms, including figuring in the identity claims that make the substitutional commitments that govern their use explicit. The challenge is, as she says, just to specify its meaning, how it *is* used.

52. Ibid., p. 151.

53. Ibid., p. 137.

54. Ibid., p. 138.

55. This discussion of Anscombe is much indebted to McDowell's lectures on her essay, but it should not be inferred that he would agree with what is said about it here.

56. See Gareth Evans's discussion, in Chapter 6 of his *Varieties of Reference,* ed. J. McDowell (New York: Oxford University Press, 1980).

57. McDowell, "De Re Senses," p. 105.

58. Hector-Neri Castañeda, "'He'"; "Indicators and Quasi-Indicators"; and "The Logic of Self-Knowledge," *Nous* 1 (1967): 9–22.

59. In "The First Person," p. 136.

60. It may be objected that in that case no definite belief has been ascribed because it has not been specified which painting is the subject of the belief. According to the anaphoric analysis of 'this*' offered below, the object has in fact been specified, though perhaps not in a form recoverable by the audience of the ascription. This can be remedied within the regimentation of ascriptions, however, as is discussed below.

61. The ascription is in the ascriber's language, though the tokening (perhaps merely virtual) that is reported need not be.

62. If one language is at issue—otherwise a tokening of the corresponding type in the language of the target of the ascription.

63. Note that doing without the quasi-indexicals would require a special convention concerning the understanding of 'as' clauses in the case where the ascription is in a different language than the believer would use to express the belief.

64. This formulation is incomplete, as appears below. The one responsible for a *de*

re ascription of a strong *de re* attitude also undertakes an *existential* commitment governing the use of the term *t*, in virtue of commitment to the identity. This point is discussed below.

65. For instance, it should now be clear how to extend the 'A' language discussed above so as to incorporate an anaphoric 'U' (corresponding as 'you', in some of its uses, does to 'I') that is used only as a dependent on some tokening of 'A'.

66. Bertrand Russell, "The Philosophy of Logical Atomism," reprinted in *Logic and Knowledge*, ed. R. C. Marsh (London: George Allen and Unwin, 1956), pp. 177–281.

67. David Kaplan introduces this notion; see for instance "Demonstratives," draft no. 2 (Mimeograph, UCLA Department of Philosophy, 1977). A general discussion can be found in Nathan Salmon, *Reference and Essence* (Princeton: Princeton University Press, 1981). Salmon puts the claim this way: "The contribution made by an ordinary proper name, demonstrative, or other simple singular term to securing the information content of, or the proposition expressed by, declarative sentences (with respect to a given possible context of use) in which the term occurs (outside the scope of nonextensional operators such as quotation marks) is just the referent of the term, or the bearer of the name (with respect to that context of use)" ("Reflexivity," reprinted in *Propositions and Attitudes*, ed. Nathan Salmon and Scott Soames [New York: Oxford University Press, 1988], p. 241). See also the essays by Mark Richard and Scott Soames in ibid.

68. John McDowell, "Singular Thought and the Extent of Inner Space," in *Subject, Thought, and Context*, ed. P. Pettit and J. McDowell (New York: Oxford University Press, 1986), pp. 137–168. See also McDowell, "De Re Senses."

69. One might be inclined to say that a properly produced tokening of 'this' always succeeds in picking out *something*, but this is not the case, for reasons that Anscombe admirably explains: "It used to be thought that a singular demonstrative, 'this' or 'that', if used correctly, could not lack a referent. But this is not so, as comes out if we consider the requirement for an answer to 'this what?' Someone comes with a box and says 'This is all that is left of poor Jones'. The answer to 'this what?' is 'this parcel of ashes'; but unknown to the speaker, the box is empty. What 'this' has to have, if used correctly, is something that it latches on to (as I will put it): in the present example it is the box . . . Thus I may ask 'What's that figure standing in front of the rock, a man or a post?' and there may be no such object at all; but there is an appearance, a stain perhaps, or other marking of the rock face, which my 'that' latches on to. The reference and what 'this' latches on to may coincide . . . But they do not have to coincide, and the referent is the object of which the predicate is predicated where 'this' or 'that' is a subject" ("The First Person," p. 143).

70. This is an additional condition on the use of expressions of the form *S* believes of$_{strong}$ *t* that Φ(it), beyond those indicated a few paragraphs back.

71. See McDowell's discussion of this sort of case in "Singular Thought and the Extent of Inner Space."

72. McDowell, following Evans, argues persuasively that he was not, in "De Re Senses" and "Singular Thought and the Extent of Inner Space." He would restrict the transparency claim to individuative omniscience concerning senses grasped *together*, in one thought. This does not exclude false beliefs of identity or nonidentity, but McDowell would insist that the man who does not know that Hesperus is Phosphorus is not confused about which object his Hesperus-

beliefs are about: he knows that full well—they are about Hesperus. The effect of this view is achieved here as well, but the interpretation of demands for transparency is different.

73. "Eines des folgenreichsten," in the first paragraph of "ÜSB." Recall the discussion below at 7.5.2.

74. Although as usual the discussion here is limited to singular terms, expressions in other grammatical categories can be used in the strong *de re* way—sortals and predicates prime among them. The account offered here generalizes straightforwardly to these cases.

75. Thus in accord with the phenomenalist methodology endorsed here, what matters is not who the name-user would treat as an *expert* with authority over the use of the name (one who knows who or what the name refers to), or even how it is determined who *ought* to exercise such authority, but rather what it is for a scorekeeper to *attribute* such authority to someone.

76. Saul Kripke, "A Puzzle about Belief," in *Meaning and Use*, ed. A. Margalit (Dordrecht: D. Reidel, 1979), pp. 239–283, reprinted in *Propositions and Attitudes*, ed. N. Salmon and S. Soames (New York: Oxford University Press, 1988), pp. 102–148 (page references are to this version).

77. Ibid., p. 130.

78. Ibid., p. 122.

79. Ibid., p. 132.

80. Ibid., p. 134.

81. Ibid., pp. 112–113.

82. Ibid., p. 113.

83. This issue is masked rhetorically by the structure of the piece. The question would arise evidently and urgently if the 'Paderewski' case were discussed immediately after the disclaimer about ambiguity in the statement of the disquotational principle. In fact the two are separated by a long detour through the (ultimately irrelevant) bilingual cases, in which translation is also involved.

84. Kripke, "A Puzzle about Belief," p. 111.

85. Ibid., p. 112.

86. Ibid., p. 111.

87. Ibid.

88. Any sort of belief, whether strong *de re* or not, can be ascribed in either the *de dicto* or the *de re* style. The *de re* form of ascription of strong *de re* beliefs is derived from the *de dicto* form of ascription of such beliefs in just the same way that *de re* specifications of the contents of ascriptions are formed from *de dicto* ones in the case of ordinary beliefs. (Which is not to say that on a case-by-case basis, the ascriber must be in a position to *produce* an underlying *de dicto* ascription in order to be entitled to the corresponding *de re* specification.)

89. Kripke, "A Puzzle about Belief," p. 111.

90. This social dimension may remain implicit, as when an interpreter looks at some nonscorekeeping system from the outside and takes or treats it as having representationally contentful states—that is treats it as having states that express a doxastic point of view that can in principle be contrasted and combined with our own in the way expressed by *de re* ascriptions. It is in this spirit that one can interpret the illumination of a region of the alarm panel as representing the front door as being open. See Section II, Subsection 4 above.

91. The Dummettian phrase "circumstances and consequences of application" used

in Chapter 2 has now been superseded by the full scorekeeping account of the force or pragmatic significance characteristic of various speech acts—paradigmatically assertion—that was offered in the following chapter. There the notion of *appropriateness* of circumstances and consequences of application is further articulated in terms of the deontic distinction between commitment and entitlement. A finer-grained account of the relation between circumstances and consequences is then offered in terms of the three notions of permissive, or entitlement-preserving, inference; commissive, or commitment-preserving, inference; and incompatibility, linking entitlement and commitment. These further developments do not make a difference to the point being pursued in the text here, however, so the original phrase can continue to be used as a sort of shorthand.

92. Recall the distinction offered above in 1.2–3 between *simple* regularity theories and those that invoke regularities concerning the practical *assessment* of performances by sanctions. The gerrymandering argument shows that only in the latter form can social regularity theories make sense of the notion of incorrect performances, even for individuals.

93. This is, attributing it to oneself *as* oneself; see Section V, Subsection 2.

94. "Thought and Talk," in *Inquiries into Truth and Interpretation* (New York: Oxford University Press, 1984), pp. 169–170.

95. Indeed, in the same essay Davidson implicitly connects this point about how the notion of beliefs as having objective truth conditions arises out of the contrast of doxastic perspective between interpreter and interpreted with what is expressed by ascriptions specifying the content of the attributed belief in the *de dicto* and *de re* styles. He says "the general and not very informative reason" why the attribution of any propositionally contentful intentional states (what he calls "thoughts") implicitly appeals to a situation in which *speech* acts are interpreted by another is that "without speech we cannot make the fine distinctions between thoughts that are essential to the explanations we can sometimes confidently supply." Those distinctions turn out to be just the ones expressed in the regimentation employed here by *de dicto* and weak *de re* ascriptional locutions: "Our manner of attributing attitudes ensures that all the expressive power of language can be used to make such distinctions. One can believe that Scott is not the author of *Waverly* while not doubting that Scott is Scott; one can want to be the discoverer of a creature with a heart without wanting to be the discoverer of a creature with a kidney; one can intend to bite into the apple in the hand without intending to bite into the only apple with a worm in it; and so forth. The intensionality we make so much of in the attribution of thoughts is very hard to make much of when speech is not present. The dog, we say, knows that its master is home. But does it know that Mr. Smith (who is his master), or that the president of the bank (who is that same master), is home?" (ibid., p. 163). These are, of course, just *de dicto* and *de re* specifications of the same belief. The dog knows *of* the president of the bank that he is home, he just does not know *that* the president is home. It was pointed out in Section III that one wants to appeal to the belief that his master is home to explain why the dog is so happy, and to its being a belief *of* the president of the bank (whether the dog knows that or not) in order to explain why one result of the dog's happiness is that he slobbers on the president of the bank. We want to be able to offer

intentional explanations of both sorts of things the dog does, and this requires content specifications that depend for their intelligibility on our mastery of the scorekeeping practices involved in interpreting the *utterances* of others.

96. Recall that 'claims' as reconstructed here is equivalent to 'is committed to the claim'; it does not require the performance of an overt speech act to that effect.

97. Notice that one reflection of the modal character of incompatibility entailments is that there is no requirement that a content that defeats a conditional claiming such an entailment (by being incompatible with the consequent but not with the antecedent) be *true*—or one the one assessing the conditional endorses. If in the definition

$$p \to q \text{ iff } \{x: x/q\} \text{ is a subset of } \{x: x/p\}$$

the range of x were restricted to those endorsed by the scorekeeper assessing the conditional, then conditionals could be nontrivially assessed only by those who endorsed something incompatible with their consequents, and hence someone who is precluded from detaching from those conditionals. The only restrictions on the range of x are those imposed by the order in which logically complex expressions must be introduced (so that if p and q are logically atomic, for instance, conditionals and negated claims are not yet available in assessing the conditional).

98. For present purposes, incompatibility may be assumed to be a symmetric relation; though that condition can be relaxed, doing so does not materially affect the demonstrations offered here.

99. For instance, if p^α is $(p)(S^\alpha$ claims that $p)$, then p^α *will* be incompatible with $\sim(S^\alpha$ claims that $p^\alpha)$ because the latter negation is incompatible with a substitution instance of the former quantificational claim. The recipe below will produce a q^α that defeats the conditional (i) for any *non*ascriptional p^α.

100. One might quibble here that whether or not the restricted quantifier over believers includes '!xDx' as a substituend is part of what is at issue in the case. The claim that $(S)(S$ claims that $p^\alpha)$ entails (!xDx claims that $p^\alpha)$ does turn on this issue, but the claim that these are not incompatible (which is all the proof appeals to) patently does not.

101. These issues are explored further in the author's "Truth and Assertibility," *Journal of Philosophy*, 73, no. 6 (March 1976): 137–149.

102. Nor, of course, have I undertaken incompatible commitments by these ascriptions, but that would be true even if I were ascribing to S commitment to incompatible claims. It remains true that incompatible claims are not attributed to S, even if I indicate that S attributes to me the use of the first-person pronoun:

S claims ⌜he$^{(I)}$ does not claim that p⌝ and S claims ⌜p⌝

are still not ascriptions to S of incompatible commitments.

103. Of course the original problem *does* arise if what I attribute to S is ascribed rather by my endorsement of:

S claims that he$^{(I)}$ does not claim that p, and S claims that p.

But given the account offered earlier of the use of 'I' and of quasi-indexicals anaphorically dependent on it, it is clear these ascriptions do not involve his attitudes toward claims with the same contents as mine. This can be seen as

well in the fact that S cannot vindicate his entitlement to the claim expressed in his mouth by ⌈I do not claim that p⌉ by deferring ascriptionally to *my* assertion of ⌈I do not claim that p.⌉

104. W. V. O. Quine, "Quantifiers and Propositional Attitudes"; McGinn, "The Structure of Content."

9. Conclusion

1. *Principia*, First Scholium to the Definitions, p. 11 in Motte's translation of 1729 (Berkeley: University of California Press, 1934, 1962). These words actually appear as the antecedent of a conditional, but since Newton immediately goes on to endorse its conclusion, it is clear that he endorses this claim. (Of course, for him, as for others of his time, this dictum applies only to *words*, not to the contents of the concepts they express.)

2. Cf. *Critique of Pure Reason*, A51/B75. What Kant actually says is that *thoughts* without *content* are empty, and intuitions without concepts are blind.

3. At least as far as the doctrines of the first two Critiques are concerned.

4. See John McDowell's "De Re Senses" (Philosophical Quarterly 34, no. 136 [July 1984]: 283–294), in which this strand of thought is discerned in Tyler Burge's "Belief De Re" (*Journal of Philosophy* 74 [1977]: 338–362), which presents very clearly views that are endorsed by many other thinkers.

5. Kant distinguishes the first two ways of characterizing the distinction between intuitions and concepts as "logical," in contrast to the "metaphysical" way of characterizing that distinction, which appeals to receptivity and spontaneity. (See for instance Section V of the Introduction to the *Logic*.)

6. As the tradition shows, use of this term is fraught with danger because of what Sellars (in "EPM") calls "the notorious 'ing'/'ed' ambiguity."

7. "On the Very Idea of a Conceptual Scheme," in *Inquiries into Truth and Interpretation* (New York: Oxford University Press, 1984), pp. 183–198.

8. This is the problem with which Hegel opens the Introduction to the *Phenomenology*.

9. For instance in "Concept and Object," in *Translations from the Philosophical Writings of Gottlob Frege*, ed. P. T. Geach, and M. Black (Oxford: Basil Blackwell, 1970), pp. 42–55. For many purposes it is more helpful to consult Frege's notion of *Sinn* than his notion of *Begriff* in order to learn about what is discussed here as conceptual articulation.

10. The point being made concerns contemporary ways of thinking about concepts that Kant inspired, rather than the details of his view. Nonetheless, particularly in the context of the present project, it would be unfair to accuse Kant of confusion on this point. As "double affection" readings of Kant make clear, a major interpretive challenge facing the readers of the first Critique is to make sense of talk of spontaneity and receptivity at the noumenal or transcendental level, even though the causal relations one naturally turns to in such a case are conceived as products of those two faculties, and so remain at the phenomenal level. The general picture presented here about the lessons to be learned from Kant suggests that the answer is not far to seek: spontaneity and receptivity are normative notions for Kant, not *causal* ones. Roughly, talk of spontaneity is talk about responsibility *for,* or authority *over,* something. Talk of receptivity is talk

about responsibility *to,* or being subject to the authority *of,* something. These notions accordingly line up with:

—The 'I think' that accompanies all representations as a purely *formal* feature—not part of what is represented or the content expressed by the representing, but of the form of the representing as such. The root notion is the taking of responsibility *for* a claim or judgment. It is because judgments are the minimal unit of responsibility or accountability that they are the minimal unit of experience for rational beings. This is also why spontaneity—that is, the understanding, the faculty of judgment—traces back to the Transcendental Unity of Apperception (the locus of coresponsibility for claims), which in turn underwrites the categories.

—The 'object = X' that implicitly accompanies all judgments of experience, as a purely formal feature of the representing. This is the dimension along which such judgments are responsible for their correctness *to* something represented—which is part of its being a representing at all: its representational purport.

That both of these purely formal features are necessary for representation at all is what makes the refutation of idealism work. Notice that in this Kantian dualism, the formal two-sides-of-one-coin relation of responsibility *for* and responsibility *to* as essential to discursive representation, *is* carried over into (indeed is essential to) the scheme presented in this work.

11. One modern version of Kant's contrast between spontaneity and receptivity takes the form of a contrast between the repeatable, purely qualitative or descriptive terms whose deployment is entirely within the "cognitive world" of the thinker, on the one hand, and the unrepeatable, contextually determined indexical expressions whose significance outruns what is within the grasp of the thinker and is the ultimate source of constraint on thought, on the other. Here are two representative passages: "A *de re* belief is a belief whose correct ascription places the believer in an appropriate nonconceptual, contextual relation to objects the belief is about"; and "A sufficient condition for a belief context to be *de re* . . . is for it to contain an indexical expression used deictically" (Burge, "Belief *De Re,*" pp. 346, 347).

Part of the responsibility for this line of thought lies in a common misreading of the nature of Kripke's project in "Naming and Necessity" (in *Semantics of Natural Language,* ed. G. Harman and D. Davidson [Dordrecht: D. Reidel, 1972]). While Kripke was addressing exclusively the semantics and epistemology of proper-name usage, people often derive consequences for semantics and epistemology generally. His distinction between descriptive and causal-historical approaches to proper names works well enough in the context of his actual project. But disaster results if one takes him as having shown, or even suggested, that there are two general styles of theory about how language or mind encompasses the world—namely descriptive and causal-historical or contextual ones. For that distinction patently does not extend to predicates. No one was ever a descriptivist about predicate denotation, because of the regress such a view would obviously involve. Everyone, even the ones who wanted to appeal to predicates via descriptions to account for proper names, always thought that at least the basic predicates achieved their reference in a different way, one that involved contact with the properties. Burge, for instance, says nothing about

how he conceives our relation to the semantic correlates of the predicates appealed to in the descriptions that underwrite *de dicto* (purely conceptual) attitudes for him. It is hard to see how his story can be extended to include them, for he seems to be working with a picture that allows only descriptive or demonstrative contact with the world, and he cannot treat predicates in either way without either falling into the regress or conceding that there cannot be attitudes that are *de dicto* in his sense.

12. Only "roughly" because at the next level of analysis, the work done by the undifferentiated notion of propriety (appealed to in talk of "circumstances of *appropriate* application" and "*appropriate* consequences of application") is taken over by a theoretical idiom that distinguishes the deontic statuses of commitment and entitlement. Then the scorekeeping significance that must (in context) be determined by conceptual content is articulated according to commitment-preserving inferences, entitlement-preserving ones, and incompatibility inferences (which involve both deontic statuses). See Chapter 3.

13. Notice that this definition says *what* it is to take one tokening to be anaphorically dependent on another, not *when* it is appropriate to do so. Thus it does not purport to answer the linguist's fundamental question about anaphora—namely what it is that determines what is the proper antecedent of any given anaphorically dependent expression. It purports to answer only the philosopher's fundamental question about anaphora—namely how to understand the relationship at issue in the linguist's question.

14. According to the theory expounded in Chapter 5, this eight-word phrase is itself a complex, anaphorically indirect definite description, which is anaphorically dependent on the original tokening of 'that'.

15. In the classical cases. In fact most of what appear syntactically as definite descriptions are not semantically definite descriptions, but function instead as anaphoric dependents.

16. 'Claim' in the sense of what is claimed, not the *claiming* of it. See 5.5 above.

17. To show why this *must* be the case, why the conceptual articulation of sentences must take a form analyzable in terms of the substitution-inferential roles characteristic of singular terms and predicates, is the burden of Chapter 6.

18. So it is here that the notion of *freedom* first finds application. The notions of will and choice are not available independently in order to make sense of this notion of freedom. Rather they are themselves to be understood in terms of the practical attitude of acknowledging norms as binding (acting according to a conception of a rule).

19. Though, as explained in Chapter 1, nondiscursive creatures can treat each other as bound by norms that are not inferentially articulated in the way constitutive of concepts.

20. Of course (as a helpful reader put it) it is ^snorms all the way down^s only if one digs in the right direction. For if what is to be explained is, say, the fact that the planets have elliptical orbits (the truth of what is claimed, rather than what we are doing in claiming it), the explanation does not need to appeal to anything normative. For that would have been a fact (what is claimed would have been true), even if norm-instituting, claim-making creatures had never evolved. Being entitled to say both of these sorts of things is a prime criterion of adequacy of the present account. See 1.5.3; 5.5.3–4; 8.6.3–5.

21. According to the account offered in Chapter 7, making explicit the sense in

which deontic statuses or discursive norms exist requires specifying the use of a class of *canonical designators* corresponding to this realm—expressions that can play for normative objects the role played for physical objects by spatiotemporally definite descriptions, for natural numbers by successor numerals, and so on. Although such a story is not filled in here, the raw materials for it are provided by the account of scorekeeping with deontic attitudes (in Chapter 3), of the way in which such attitudes can be made explicit using normative vocabulary (in Chapter 4), and of the use of singular terms and the formation of definite descriptions (in Chapters 6 and 7).

22. As a helpful reader points out, the terms in which one says that the umpire's taking the pitch to be a strike caused a fight are not purely causal, according to the theory put forward here. A more careful statement would be that some event that constituted the umpire's taking the pitch to be a strike caused a fight. The "constituted" here is normative talk about the pragmatic or scorekeeping significance of a performance. (Compare the way in which moving a pen in certain ways can constitute acknowledging a commitment to pay the bank a certain sum every month.)

23. Chapter 3 (Section IV) discusses how scorekeeping proprieties are codified in contents in the context of a theory of speech acts.

24. "Thought and Talk," in *Inquiries into Truth and Interpretation* (New York: Oxford University Press, 1984), p. 156.

25. This sense in which simple intentionality is derivative from the original intentionality attributable to language users should not be confused with the sense in which the intentionality of linguistic performances and expressions is derivative from that of the community of language users. See 1.6.2.

26. As indicated below, it should be clear by this point in the exposition that this is compatible with the community being ignorant or mistaken about what their practice has made their words mean (what concepts those words have come in practice to express).

27. Recall that it was applying this stance stance to internal intentional interpretation that makes possible the reconstruction of the hybrid deontic attitude of attributing knowledge, compounded in social-perspectival terms out of the *attribution*, to the one taken thereby as a knower, of commitment (belief) and entitlement (justification), and the *undertaking*, by the one attributing knowledge, of commitment to that same content—since undertaking a doxastic commitment of the sort acknowledged by asserting is taking-*true*. Again, it permits the epistemological via media between justificatory internalism and reliabilist externalism about entitlements—namely, a perspectival reliabilism that turns on an account of the inferences endorsed by scorekeepers who *attribute* reliability (and hence entitlement; see 4.3.3).

28. Giving and asking for reasons in discursive games of this sort belongs in this regard in a box with baseball, rather than with purely formal games such as chess. Baseball cannot be played except with a ball of a specified size and composition, on a field of a specified size and composition. By contrast, chess is an empirically and practically *hollow* or abstract game; it can be played with pieces of any size and composition, on a board of any size and composition, so long as the pieces are distinguishable and interpreted as standing and moving in appropriate formal relations to each other. Solid discursive practices incorporate nonlinguistic things in them (are *corporeal*). In the same way, the practice in

which a performance can have the significance of hitting a home run incorporates objects.

29. For instance in "The Conditions of Thought," in *Le Cahier du College International de Philosophie*, (Paris: Editions Osiris, 1989), pp. 165–171, reprinted in *The Mind of Donald Davidson*, ed. J. Brandl and W. Gombocz (*Grazer Philosophische Studien* 36 [1989]), pp. 193–200; "Meaning, Truth, and Evidence," in *Perspectives on Quine*, ed. R. Barrett and R. Gibson (Oxford: Basil Blackwell, 1990), pp. 68–79; and "The Myth of the Subjective," in *Bewusstsein, Sprache, und die Kunst*, ed. M. Benedikt and R. Burger (Edition S Verlag der Österreichischen Staatsdruckerei, 1988), pp. 45–54, reprinted in *Relativism: Interpretation and Confrontation*, ed. M. Krausz (Notre Dame, Ind.: University of Notre Dame Press, 1989).

30. "Two Dogmas of Empiricism," Essay 2 in *From a Logical Point of View* (Cambridge: Harvard University Press, 1953).

31. "Concepts As Involving Laws, and Inconceivable without Them," in *PPPW*, pp. 87–124. More recently, Dretske has taken a similar line, in his *Knowledge and the Flow of Information* (Cambridge: MIT Press, 1981).

32. Cf. Quine's remark in "Two Dogmas": "Meaning is what essence becomes when it is divorced from the object of reference and wedded to the word."

33. This notion is introduced at 8.1.1 and again at 8.1.3. Iteration of propositional attitude ascriptions is discussed in the Appendix to Chapter 8.

34. This use of 'should' has been explained (in 4.2.8) as making explicit the attribution of a practical discursive commitment, and so the endorsement or undertaking of commitment (on the part of the one employing 'should' or to whom its use is attributed) of a pattern of practical inference. 'Should' in its "all-in" sense, as opposed to its prima facie sense, is the third-person, or attributing, form whose first-person or acknowledging, form is 'shall'.

35. In the terms of the *Phenomenology*, the phenomenal attitudes whose development is being considered have been developed to the point where they coincide with the phenomenological view we have adopted all along in considering them. Hegel's alarming term for this sort of explicit interpretive equilibrium is "Absolute Knowledge."

36. The social 'we' is constructed out of perspectival *I-thou* scorekeeping or interpretive relations. At this point no distinction of kind remains separating the internal *I-thou* relations between the scorekeeping attributor of a commitment and the one to whom it is attributed, on the one hand, and those between the external attributor of scorekeeping practices and those to whom they are attributed, on the other.

37. Notice that to say this is to invert the Davidsonian order of explanation, which discusses conversation within a language as a special case to be modeled on interpretation across languages.

Index

< >   type designation, 314
//   token designation, 314
⌈ ⌉   corner quotes, 688n25, 704–705n24
ˢ ˢ   scare quotes, 545–547, 588–590, 672n19

and assertion, 157–159, 167–175, 190–191, 200–203
and commitments, 157–166, 178–180, 193–196, 237, 243, 263
and communication, 156, 174–176, 282, 478, 633
and content, 141–142, 145–147, 167–175, 186–191
and entitlement, 159–166, 176–180
and interpretation, 475, 508, 630, 644–645
model, 166, 168–172, 180–186, 190–191
and objectivity, 324, 529, 601–607, 627–628, 636
as perspectival, 185, 332, 488, 590, 602, 604, 627, 649
and practical reasoning, 230–233, 244, 256, 625, 640
and pragmatic significance, 142, 167–168, 182–190, 262, 284, 710n91
and speech acts, 142, 182–193
deontic status(es), 55, 165–166, 189–190, 201–205, 226, 237, 275, 290, 595–596, 636–637, 649. *See also* commitment(s); deontic attitudes; deontic scorekeeping; entitlement; intentional states; normative status(es); norms
and deontic attitudes, 161–162, 165–166, 334, 357, 593, 598–601, 623, 648, 687n11
and objectivity, 197, 201, 599–607
de re ascriptions, 391, 499–508, 511–515, 525, 542–548, 565, 566, 584, 589, 598, 601, 646, 648, 701n7, 703n19, 709n88, 710n95. *See also* ascription(s); *de dicto* ascriptions; information; representation
denotational sense of, 547–548, 566, 706nn39,45
expressive role of, 138, 502, 512, 516, 522, 525
and intentional explanation, 522–527
and objectivity, 528, 595, 598, 600–601
and 'of', 506
of propositional attitudes, 187, 279–281, 485, 490
and representation, xvii, 279, 499–503, 517, 519, 520, 586
strong, 503, 530, 566, 569–570, 581
weak, 503, 513, 529–530, 547–552, 566
Descartes, René, 6, 9–11, 31, 74, 93–94, 279, 614, 623, 655n2, 671n8, 681n5

designatedness, 340–350, 352, 353, 356–358, 687n8. *See also* multivalues; substitution
designation, 69, 84, 85, 665n32, 696n61. *See also* canonical designators; proper names; rigidity; semantics; singular terms
desires, 56, 58, 240, 246, 250, 256
Devitt, Michael, 481
Dewey, John, 289, 299
direct discourse ascriptions. *See* ascription(s), direct discourse
direct reference. *See* reference, direct
disavowals, 192–193, 670n6
discursive attitudes. *See* deontic attitudes
discursive commitments. *See* commitment(s)
discursive practice. *See* linguistic practice
discursive scorekeeping. *See* deontic scorekeeping
discursive statuses. *See* deontic status(es)
disjunction, 115, 429–430, 434–435, 438, 441–442
dispositions, 28–29, 35, 42, 45–46, 208, 625, 628, 629, 636, 638. *See also* causes; regularities; reliable differential responsive dispositions
disquotation, 300–301, 303, 323, 575, 577, 579, 581
distinction
  acknowledging/attributing, 55, 193–197, 259–262, 270, 554, 596, 633, 646, 649
  action/perception, 7–8, 119–120, 209–211, 233–238, 261, 335, 336
  analytic/synthetic, 145, 345, 358, 484
  anaphoric/causal-historical approaches to proper names, 308–309, 458, 459, 470–471, 572–574, 579–583, 585–586
  asserting/inferring, 158, 167–168, 171, 339, 347–348, 350, 351, 353, 355, 358, 359, 367
  asymmetric/symmetric, 372, 376–381, 384–385, 388, 391–395, 403, 455, 457, 490, 499, 564, 619–622
  attitude/status, 33, 37, 161–162, 165–166, 194, 197, 261, 290, 334, 357, 497, 593, 595–596, 598–601, 623, 648
  attributing/undertaking, 61, 62, 161–166, 196, 506–507, 525, 554, 596–601, 608
  authority/responsibility, xii, 161–165, 171, 174, 179, 229, 238, 242, 532