

Toward an Inferential Semantics

Sure, he that made us with such large discourse,
Looking before and after, gave us not
That capability and god-like reason
To fust in us unused.

SHAKESPEARE, Hamlet

I. CONTENT AND REPRESENTATION

1. *Intentionality: Propositional and Object-Representing Contentfulness*

Taking or treating someone as one of us may be called *recognizing* that individual.¹ According to the construal of recognition being developed here, taking or treating as one of us is adopting a certain kind of intentional stance. It requires first of all interpreting the one it is directed at as the subject of intentional states. But taking someone as one of us also requires, it was suggested, interpreting that individual as an intentional interpreter—as able to *attribute* intentional states, and so as able to adopt toward others just the same sort of attitudes out of which that very stance is constructed. The previous chapter assembled some raw materials for an account of the normative significance of the intentional states we attribute to each other—and take each other to attribute to each other—in adopting the attitudes of mutual recognition that institute the status of community membership, of being one of us.

Before such an account is presented, in the next chapter, it is necessary to look more closely at the sort of *content* that sets apart—as distinctively *intentional*—the states and statuses (and therefore the attitudes) that are attributed when we recognize someone. For intentional states are intention-

ally contentful states, and the theoretical job of the contents they are taken to have is precisely to determine, in context, the particular significance of being in or attributing the states those contents are associated with. As the terms are used here, *semantics* is the study of such contents, and *pragmatics* is the study of the force or significance of the states, attitudes, and performances that have those contents. Accordingly, to fill in the details of a story about the normative character of the *pragmatic significance* of intentionally contentful states, attitudes, and performances, an inquiry into the nature of their *semantic contents* is called for.

Brentano, who brought the term 'intentionality' back into modern usage, defines it this way: "Every mental phenomenon is characterized by what the scholastics of the Middle Ages called the intentional (also mental) inexistence of an object, and what we could call, although not in entirely unambiguous terms, the reference to a content, a direction upon an object (by which we are not to understand a reality in this case), or an immanent objectivity. Each one includes something as object within itself, although not always in the same way. In presentation, something is presented, in judgment something is affirmed or denied, in love loved, in hate hated, in desire desired, etc."² "Not always in the same way" indeed. "Intentional object" as used here involves assimilations along two dimensions. First is a dimension most clearly picked out in terms of grammatical categories: what is affirmed or denied in a judgment is something expressible by the use of a declarative sentence, while what is loved or hated is something referred to by the use of a singular term. The former may be called 'propositional contentfulness'. It is typically expressed by the use of a declarative sentence and is ascribed by the use of a 'that' clause appended to a specification of the contentful state or attitude, as in "the belief *that* Carlyle wrote *Sartor Resartus*" or "desiring *that* Pufendorf's reputation be rehabilitated." The latter may be called 'object-representing contentfulness'. It is typically expressed implicitly by the use of a singular term as a grammatical direct or indirect object, and it is attributed explicitly by using terms such as 'of' or 'about', as in "a belief *about* Carlyle" or "desiring something *of* Pufendorf's reputation" (for example that it be rehabilitated).

Putting these two sorts of contentfulness in a box together is not just an idiosyncrasy of Brentano's. Searle, for instance, offers this pretheoretical delineation of the subject matter of his book *Intentionality*: "If a state *S* is Intentional then there must be an answer to such questions as: What is *S* about? What is *S* of? What is it an *S* that?"³ To insist on distinguishing these sorts of content in the way indicated above is not yet to diagnose a confusion in remarks like this. There is no confusion insofar as propositional and object-representing contentfulness ought to be understood as species of a genus.

An approach to the characterization of that genus is not far to seek. Stalnaker speaks for the dominant tradition in offering this formulation:

"The problem of intentionality is a problem about the nature of representation. Some things in the world—for example pictures, names, maps, utterances, certain mental states—represent or stand for, or are about other things—for example people, towns, states of affairs."⁴ The genus, it is suggested, is *representational* content. Indeed, Stalnaker, like others, is comfortable talking interchangeably about "intentional or representational states." This basic insight should be accepted to this extent: it is clear that intentionality has a representational dimension and that to understand intentional contentfulness one must understand representation.

A common response to this insight is to envisage an explanatory strategy that starts with an understanding of representation and on that basis explains the practical proprieties that govern language use and rational action. It is not clear, however, that a suitable notion of representation can be made available in advance of thinking about the correct use of linguistic expressions and the role of intentional states in making behavior intelligible. The temptation to think otherwise is connived at by insufficient appreciation of some of the fundamental criteria of adequacy to which an account of the representational dimension of intentional contentfulness must answer. It is important to keep in mind the explanatory challenges faced by a semantic theory that appeals to representation as its basic concept, and some of the ways in which those explanatory obligations are liable to be unobtrusively shirked. To point these out is not to show that they cannot be satisfied—that representational explanatory strategies are in principle broken-backed. It is merely to guard against the danger that such an explanatory starting point may recommend itself in virtue of its apparent immunity to difficulties it has not squarely confronted.

A particularly unhelpful way of pursuing the representational semantic explanatory strategy is to model representation on *designation*. The designational model is objectionable on two grounds connected with the distinction of grammatical category between sentences and subsentential expressions such as singular terms. First, it assumes that the relation between a singular term and the object it picks out or refers to, for instance that between a name and its bearer, is antecedently intelligible—that the notion of tagging or labeling something can be made sense of before one considers the use of such tags or labels in *saying* something (paradigmatically, in making a claim). In this way, the strategy runs afoul of the principle of the pragmatic priority of the propositional, which is discussed further along.

Second, it assumes that the notion of representation as reference picked out in this way for the category of singular terms and predicates can be univocally and unproblematically extended to apply to the category of sentences. Sentences are understood as representing states of affairs, in the same sense that singular terms represent objects (and in the same sense that predicates represent properties or sets of objects). The notion of representation, conceived as designation, is then supposed to make the grammati-

cal distinction between singular terms and sentences intelligible by appealing to the ontological distinction between objects and states of affairs. Even if it is granted that there is a clear sense in which singular terms such as names and marks on maps represent particular objects, for instance individual people and cities, it does not follow that it is possible to introduce the category of states of affairs as what is in the same sense represented by declarative sentences and 'that' clauses. Nor ought it to be assumed that the ontological category of states of affairs can be made intelligible apart from and in advance of explaining the use of declarative sentences and the 'that' clauses used to report such uses in indirect discourse.

2. Two Senses of 'Represents'

Introducing the notion of states of affairs as the kind of thing represented by declarative sentences requires sensitivity to the second dimension of assimilation involved in Brentano's idiom. For one must be careful not to confuse what is *represented* by sentences with what is *expressed* by them. This is a familiar point, but it is worth emphasizing. As Brentano acknowledges by appending to his phrase "direction upon an object" the qualification "by which we are not to understand a reality in this case," 'represent' is ambiguous between two intimately related but importantly distinct senses. Searle puts the point this way: "'About' . . . has both an extensional and an intensional-with-an-s reading. In one sense (the intensional-with-an-s), the statement or belief that the King of France is bald is about the King of France, but in that sense it does not follow that there is some object which they are about. In another sense (the extensional) there is no object which they are about because there is no King of France. On my account it is crucial to distinguish between the content of a belief (i.e. a proposition) and the objects of belief (i.e. ordinary objects)."⁵ Thus as Searle sets things up, for a statement or belief to have content is for it to represent or be about something in the 'intensional' sense, while for it to have an object or objects is for it to represent something in the 'extensional' sense.⁶ The relation between the two senses emerges more clearly if one or the other is taken as primary and the remaining one specified in terms of it. Thus if 'represent' is reserved for the sense in which one can represent only what in fact exists, whether it be in the category of objects or of states of affairs—actual objects corresponding to singular terms and actual states of affairs corresponding to true claims—then the other sense can be picked out as *purporting* to represent. The other way to do things would be to use 'represent' even in the cases where nothing exists to be represented, where there need be no object or state of affairs as represented for there nonetheless to be a representing. When something does exist as represented, the representation might be called *successful* or correct.

An account of contentfulness in terms of representation needs to explain

both of these senses in which something can be a representing, and it needs to explain their relation to one another. It is clear that if contentfulness in general is to be identified with representational contentfulness, that is, with being a representing, then 'representing' should be understood as purported representing in a sense that contrasts with successful representing. For it makes sense to wonder whether, or to believe that, there is a present king of France or Schelling was the greatest German philosopher, even if it turns out that no object or state of affairs corresponds to that contentful state. A theoretical idiom that shrinks the scope of purported representing until it coincides with that of successful representing has no room for the notion of error, of representation that is incorrect or mistaken; and a notion of representation so thin as to preclude assessments of correctness provides no basis for any recognizable concept of intentional content.

A theoretical idiom that, on the contrary, expands the scope of successful representing until it coincides with that of purported representing is equally unpromising. The result of holding purported representing fixed and failing to distinguish successful representing from it is Meinongianism—commitment to a vast realm of entities, most of which do not exist, including many that *could* not exist. The trouble with taking it that there is something that is successfully represented by every purported representing is not just that it involves commitment to a luxuriant ontology; ontological self-indulgence is a comparatively harmless vice. But it can be symptomatic of a failure to shoulder an explanatory burden. In this case it evidently (and ultimately unhelpfully) transforms the demand for an account of the relation between correct and incorrect, unfulfilled or merely purported and actually successful representing, into a demand for an account of the relation between the statuses of what is represented in the two cases: between mere subsistence and robust existence. Ontological postulation can no more provide an explanation by itself in this case than it could in the one just considered, where the issue was an account of the relation between the sense in which singular terms are representationally contentful and the sense in which sentences are. (Of course, no more in this case than in that one does a commitment to taking the representational dimension of intentional content seriously entail going on to make such a mistake; it is important to recognize the temptation in order to resist it.)

Brentano, who did not make the mistake of his student Meinong, indicates some of the difficulties faced by such an attempt to ontologize the distinction between correct and incorrect representation by holding to a univocal sense of 'represents' and construing the distinction as a difference between two different sorts of representable: "It would be paradoxical to the highest degree to suppose that you could promise to marry an *ens rationis* and then keep the promise by marrying an actual, concrete particular."⁷ It is disastrous to put the notion of successful representing in place of that of purporting to represent, that is, to have it play the role of necessary condition for content-

fulness. But while the two senses of 'represent' or 'about' must not be run together (from either direction), there is also reason not to want them to be driven too far apart. Purporting to represent is intelligible only as purporting to represent *successfully* or *correctly*. If what would make the representings successful has no part to play in determining the purport or content of those representings, it is hard to see how assessments of correctness could even get a grip on them. The trouble then is not just that of skepticism about justification, in the Cartesian mode. If all our ideas could have just the content-as-representational-purport that they do, even though the rest of the world, the representeds those ideas purport to represent successfully, were entirely different from what it is represented (purported) to be, how could we ever be justified in taking ourselves to be correct? The difficulty that looms is more serious still, threatening not just the cogency but even the comprehensibility of the picture of states and attitudes as contentful in virtue of their representing or being about the way things are. For the very notion of representings so much as purporting to be about representeds becomes unintelligible.

Acknowledging this distinction between representational purport and representational success is one of the theoretical jobs Frege assigns to his paired semantic concepts *Sinn* and *Bedeutung*. A sign is contentful insofar as it *expresses* a sense. A thought is the sense, the propositional content expressed by a declarative sentence. To say that it is true or false—to assess it along the dimension of correctness semantically relevant to thoughts—is to classify it in terms of the result of applying a function to objects serving as the arguments of that function, where both the function and the objects are picked out as those referred to by components of that sense. The structure of the later Frege's semantic project accordingly encompasses accounts both of what it is to express a sense and of what it is for that sense to be *correct* in terms of how things are with what it represents. An utterance or inscription *expresses* a sense, for example a thought, and it is the sense expressed that then *refers* to objects, the thought that represents them as instantiating properties and standing in relations.⁸ This idiom avoids the dangerous ambiguity inherent in talking about propositions as *represented* by sentences. For that way of talking is liable to be misunderstood as involving the identification of propositions with the facts or states of affairs successfully represented by true claims (according to the representational model of contentfulness) rather than with the claims or purported representations expressed by sentences.

3. Representational Uptake

The notion of representational *purport* implicitly involves a notion of representational *uptake* on the part of some consumer or target of the purporting. It is only insofar as something can be *taken* to be a representation

that it can purport to be one. For purporting to be something is putting oneself forward as aptly or appropriately taken to be that. The purport is veridical or spurious (for instance the representation is successful or misleading) accordingly as the taking it invites is correct or incorrect. That grasp of something as a representation is coordinate with representational purport is the point Dennett is making when he says: "Something is a representation only for or to someone; any representation or system of representations thus requires at least one user or interpreter of the representation"⁹ (using "representation" to mean purported, not necessarily successful, representation). It was pointed out in the previous paragraph that according to the representational model of contentfulness being considered, representational purport is what is *expressed* by a representing, for instance a sign design, rather than what is *represented* by it if it is successful. The present point is then that talk of what is expressed is intelligible only in the context of talk of the activity of grasping what is expressed. By widening the focus a bit, this can be seen to be the manifestation (within the representational construal of contentfulness) of the general point that *meaning* and *understanding* are coordinate concepts. The notion of representational purport is one way of rendering what must be understood in grasping the content of an intentional state, attitude, or performance.¹⁰ Representational purport and the understanding that is its uptake must both be explained in order to make an account of intentional contentfulness in terms of representation work. As Kant says: "The understanding, as a faculty of knowledge which is meant to refer to objects, requires quite as much an explanation as the possibility of such a reference."¹¹

Looking back from the vantage point won for us by the later Wittgenstein, it is possible to see that one of the unfortunate emphases that Descartes imposed on the representationalist tradition is the privileging of *knowledge*, and therefore *successful* representation, as a topic of inquiry, over *understanding*, and therefore *purported* representation. For Descartes, representational purport, being "as if of" something, is an intrinsic and characteristic property of *pensées* (that is, specifically mental acts). He does not offer an account of what it is for a mind to grasp such purport, for it to take or treat an idea as being of or about something. He is concerned with how one might become entitled to a commitment to something that has objective (in his, neo-Scholastic sense) reality in one's thoughts having also formal reality outside them. He is not concerned with what the mind's taking one thing or sort of thing rather than another (or rather than nothing at all) as having objective reality in one's thoughts itself consists in. Representational purport, "the objective reality of things in thoughts," and its corresponding uptake by the mind whose thoughts they are serve Descartes as unexplained explainers. So the content of the representational commitments to which the mind's entitlement is at issue is never clarified. A representational model of contentfulness cannot rest with an account of

successful representation—not even if it is accompanied by a vindication of the right to believe that purported representation is often or even generally successful. It requires also an account of representational purport, and that requires an account of the uptake, grasp, or understanding of such purport.

It would of course be a blunder, of a familiar kind, to understand that uptake in general as consisting in *interpreting* something *as* a representation, in Wittgenstein's sense of 'interpreting'. Taking something as a representation must not be parsed in terms of the adoption of explicitly contentful attitudes or intentional states such as belief. If being a consumer of representational purport, taking something as a representation of something, is understood as believing of it that it correctly represents (or equally if the purport is understood as intending that it do so), then an infinite explanatory regress is generated by the possibility of querying the nature of the representational purport ('that . . .') and success ('of . . .') such a belief exhibits. There must be some way of understanding something as a representation that consists not in interpreting it (in terms of something else understood as a representation) but in taking, treating, or using it in practice *as* a representation. To understand what it is for red dots on a map to purport to represent cities and wavy blue lines to purport to represent rivers, the theorist must look to the practice of using a map to navigate. If such purport is to provide a model applicable to representational purport in general, that practice must admit of construals that do not appeal to the formation of propositionally contentful beliefs. The practice must be intelligible in terms of what counts as following it or going against it in what one actually does: the way it guides the behavior of those who can use maps.

The absence of a nonregressive account of what it is to take, treat, or use something *as* a representation of something else is the source of another traditional sort of dissatisfaction with the representationalist paradigm of contentfulness. It lies behind Rebecca West's irritated response to the "mind as the mirror of nature" model that it is hard to see why one would want a copy of the universe: "One of the damn things is enough." Progress in understanding intentional contentfulness is made by invoking representational relations only in the context of an explanation of what it is that makes representings graspable or intelligible as representings in a way in which what is represented is not. That is a matter of the uptake or consumption of representational contentfulness. Apart from the representational purport it expresses, and which is there to be grasped, a representing is just another bit of worldly furniture, like what it represents. Why is not confronting a map as well as terrain just adding one more thing to be baffled about? Invoking a relation (for instance some sort of isomorphism) between representing and represented does not by itself contribute to the task of explaining what the intelligibility of the representing consists in—why one of the damn things is not enough.

4. *Expression and Representation*

Restricting attention for the moment to the propositional contents characteristic of intentional states such as belief, it has been suggested that it is no use asking what a proposition (or propositional content) is, without asking what it is for a sentence to *express* a proposition, or for a state to have one as its content. Just so it is no use asking what it is for a proposition to be *true*, or a representation to be *successful* or correct, without asking what it is to express one—what *purporting* to represent consists in. And it is no use asking what it is to *express* a proposition or other content (to *purport* to represent), without asking what it is to *grasp* or *understand* such purport. An account is required not only of how representings are distinguished from and related to representeds (in successful representing). An account is needed also of the representational content representings express—their representational purport. And that requires an account of the *attitude* of grasping such purport: of taking, treating, or using a representing as a representing, of acknowledging or attributing to it in practice its representational purport.

The treatment of representational content in upcoming chapters centers on an account of this practical attitude. Becoming entitled to use a concept of intentional content involves a twofold explanatory task: to say what it is to *express* a *propositional* content in general, and then to say what more is required specifically for the content expressed to *represent* something *objective*, in the way that matters for empirical science. Furthermore, each of these must include an account of what those who exhibit and attribute states, attitudes, and performances with such contents must do in order to count as taking or treating them in practice as contentful in those ways.

This is a request that can sensibly be addressed to Wittgenstein, as well. Even his sustained, penetrating discussions do not offer an account of what distinguishes language games within which states and performances acquire specifically *propositional* significances (the only ones that, by the lights of this work, deserve the title 'Sprachspiel'), nor of what distinguishes those within which states and performances acquire specifically *representational* significances. He argues against understanding the contents determining the significances of all states and performances in terms of representational content. For one ought not simply to presume on syntactic grounds that terms are used to refer (or fail to refer), predicates are used to describe or characterize (or misdescribe), and sentences are used to claim (truly or falsely). Instead of asking what object is being referred to by the term, what property is being ascribed by the predicate, and what fact would make the sentence true, one ought first to look at the *use* of the expressions, to see if the putative referings, characterizings, and claimings in fact play a practical role that is best understood in terms of such contents. For many expressions

that might have been thought to be doing the jobs just mentioned ('sensation', 'intending', 'beetle', 'must', 'true', 'I am in pain', . . .), consideration of their use shows that another account of the content of the putative referings, characterizings, and claimings is more appropriate.

The idea that not all contentful expressions play a straightforwardly representational role is a development of a line of thought that is already important in the *Tractatus*. Some previous varieties of logical atomism had distinguished themselves by their insistence that the only way any expression, sentential or not, could have content or contribute to the content of an expression of which it is a part is by standing for or representing something.¹² Thus, not only did these views grasp the nettle of commitment to negative and conditional facts, they also were committed to 'not' and 'if . . . then . . .' standing for some element in a complex state of affairs. The undertakers of such commitments are admirable more for their conceptual heroism than for their good sense.

In the *Tractatus*, Wittgenstein showed that one could best treat logically compound propositional contents as representing states of affairs by not treating every contentful expression (every one whose occurrence is significant for determining the state of affairs represented by the whole) as itself having its content in a representational way, by standing for something. Purely formal vocabulary, paradigmatically logical vocabulary, is contentful but does not itself stand for anything. (Kant and Frege had of course earlier shown the possibility of this sort of approach.) The opening sections of the *Investigations* argue along just these lines: not every piece of a representation contributes to its content by itself representing, and not every move in a language game is a representing of something. But Wittgenstein does not explain what one must do to be using an expression to refer, characterize, or claim (the features of use he associates with representational content), nor does he explain what is required for something caught up in a language game to express a specifically propositional content.

The notion of expression—of making propositionally explicit—shows up at two different levels in what follows. First, one who adopts the pragmatist's approach to intentionality owes an account of the practices that ultimately confer explicitly *propositional* content on the states, attitudes, and performances that play appropriate roles in those practices. This is an account of the implicitly normative practices in virtue of which anything at all can be made explicit as the content of a possible claim or belief *that p*. Such a theory should explain what it is for a performance, paradigmatically but not exclusively the tokening (by speaking or writing) of some linguistic item, to express an intentional content. And it should explain the relation between such expressions and the possession of content by states or attitudes somehow related to them. That is, it should explain what it is to express or exhibit a specifically *propositional* content—intuitively, one that could be true or false. Furthermore, it should explain what it is to express or exhibit a content

that purports to *represent* something, and it should explain the relation between representing states of affairs and representing particular objects. As has been pointed out, a necessary part of explaining the expression of contents with representational purport is explaining the *grasping* of such contents, the uptake that is the other side of such purport.

The second level concerns not the making of ordinary claims but the formulation of rules or principles. The regress-of-interpretations argument shows that the intellectualist tradition erred in treating the explicit form of norms as fundamental. But once a notion of propositional explicitness has been brought onboard in terms respectable according to pragmatist scruples, the fact that the contentful norms implicit in practical doings can be expressed in rules, claims, and interpretations that say or state explicitly what is implicit in those practical proprieties itself still stands in need of explanation. An account is needed of what it is to make explicit in the form of something that can be *said* or *thought* what is otherwise merely implicit in what is *done*. At this level, the implicit proprieties of practice that make it possible to make propositionally explicit claims are themselves made propositionally explicit in the form of rules or principles. A theory of *expression* accordingly is to explain how what is *explicit* arises out of what is *implicit*. In the first instance, it must explain how propositional content (the form of the explicit) is conferred by norms that are implicit in discursive practice—that is, what proprieties of use having such a content consist in. Then it must show how those same implicit, content-conferring norms can themselves be made explicit in the form of rules or principles.

5. From Practice to Content

These two challenges are addressed in the rest of the work. First, what role must states, attitudes, and performances play in (as it turns out, social) practice for it to be correct to interpret them as being *propositionally* contentful? That is, how are propositional contents conferred by practice? What proprieties of practical employment does possession of such content consist in? As already suggested, any answer must specify what it is for the practitioners themselves practically to take or treat states, attitudes, and performances of others and of their own as having such contents, and thereby to confer those contents on them. Chapters 3 and 4 develop a response to these questions. Second, what must be true of such contentful states, attitudes, and performances for it to be correct to interpret them as *representing* objects and objective states of affairs? Again the answer must specify what it is for the practitioners themselves in practice to take or treat those states, attitudes, and performances as having such contents, and so by their practice to confer such contents on them. Part 2—particularly in Chapter 6 (on the representation of objects by singular terms) and Chapter 8 (on objective

representation)—presents an account of these phenomena, within the framework introduced in Chapter 3.

The practical uptake of specifically representational purport must include normative *assessment* of states, performances, and expressions—assessment of their specifically representational correctness. (Of course, on pain of the familiar regress, such assessment must not be understood as in every case consisting in judging *that* a representation is correct; besides such propositionally explicit attitudes there must be practically implicit ones.) Treating something as a representation involves acknowledging the possibility that it *misrepresents*—that the representational taking is a *mistaking* (the object represented does not exist, the state of affairs represented does not obtain). It is these attitudes of distinguishing in practice between representations that are taken to be correct and those taken to be incorrect that forge the connection between the notions of representational purport and representational success.

Practical representational uptake of representings—treating objects, states, or performances as purporting to be correct representations of objects and facts—consists in taking them to be takings: taking them to express attitudes concerning what there is and how things are. That they are accordingly essentially liable for assessment as to their representational success (that they in a characteristic way answer to how things actually are for their correctness) means that such uptake incorporates an implicit distinction between representational *attitude* (how things are taken to be by what is treated as a representation) and representational *status* (how things actually are, which determines the success or correctness of that attitude). Thus the normative pragmatic distinction between status and attitude is central to the intelligibility of fundamental semantic concepts. It is reflected in the distinction between representational purport and representational success.

The *objectivity* of representational content is a feature of the practices of assessing the correctness of representations. The status of representings as correct or incorrect, successful or unsuccessful, depends on how things are with what is represented, rather than on the attitudes of representers. What is distinctive of specifically *representational* correctness is this objectivity—the way in which assessments of representational correctness take representings to answer to what is represented, rather than to how what is represented is *taken* to be. It is the way in which the status being assessed outruns any particular attitude toward it. Understanding the objectivity of representational content requires understanding this particular structure of authority and its acknowledgment—what it is for those assessing the correctness of representings to cede authority over them to what is represented, to treat their correctness in practice as determined by those represented. Again, one lesson is that the representational dimension of semantic content cannot be understood apart from the normative pragmatic context in which it is embedded and in which it is accorded its characteristic significance.

It should be clear that the remarks in this section are not meant to have the force of arguments against treating representation as a central semantic category. Rather, they present some general criteria of adequacy for an account of this important semantic notion. By doing so, however, they do offer reasons not to treat representation as a semantic *primitive*, as an unexplained explainer. The next section shows why the role of semantic concepts in pragmatics (the proper use of language and the appropriate role of intentional states in rational action) dictates approaching semantics in the first instance through the notion of *propositional* contentfulness. The rest of the chapter then motivates an approach to propositional contentfulness that begins with the *inferential* articulation of the social practice of giving and asking for reasons. The following chapter presents a particular model of those social practices (in terms of deontic scorekeeping) and shows how they can be understood as at once instituting discursive commitments and conferring propositional contents on them. In Part 2, that framework is extended to include *representational* content, both of the sort expressed by sentences and that expressed by subsentential expressions. It concludes with a discussion of the social and inferential articulation of discursive practice in virtue of which the contents it confers are properly understood as involving an *objective* representational dimension.

II. THE PRIORITY OF THE PROPOSITIONAL

1. *Kant on Judgment as the Form of Awareness*

It is appropriate to begin by addressing propositional contents because of what can be called the *pragmatic priority of the propositional*. 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*. "As all acts of the understanding can be reduced to judgments, the understanding may be defined as the faculty of judging."¹³ For him, interpretations of something as classified or classifier make sense only as remarks about its role in judgment. A concept just is a predicate of a possible judgment,¹⁴ which is why "the only use which the understanding can make of

concepts is to form judgments by them."¹⁵ Thus for Kant, any discussion of content must start with the contents of judgments, since anything else only has content insofar as it contributes to the contents of judgments. This is why his transcendental logic can investigate the presuppositions of contentfulness in terms of the categories, that is, the "functions of unity in judgment."¹⁶

The understanding is the active cognitive faculty, the faculty of spontaneity—understanding is something *we do*. "We have before given various definitions of the understanding, by calling it the spontaneity of knowledge (as opposed to the receptivity of the senses), or the faculty of thinking, or the faculty of concepts or of judgments; all of these explanations, if more closely examined, coming to the same."¹⁷ What we do is synthesize, bring things into a unity—that is, subject them to rules or concepts. What we do, as opposed to what happens to us, is to judge. Although synthesis happens at other levels than that of judgment (there is synthesis in intuition and imagination also), that synthesizing activity is an aspect of judging. "The same function which imparts unity to various representations in one judgment imparts unity likewise to the mere synthesis of various representations in one intuition, which in a general way may be called the pure concept of the understanding. The same understanding, and by the same operations by which in concepts it achieves through analytical unity the logical form of a judgment, introduces also, through the synthetical unity of the manifold in intuition, a transcendental element into its representations."¹⁸ Thus all our cognitive activity consists of judgment and aspects of that activity. Any content that can be discerned in any category is derivative from the content of possible judgments, that is, from propositional content. Kant's *pragmatics*, or theory of cognitive *activity*, determines the fundamental unit of his *semantics*, or theory of the *contents* of cognitions.

2. Frege and Wittgenstein

This insight into the fundamental character of judgment and so of judgeable contents is lost sight of by Kant's successors (indeed it could be argued that appreciation of it is still missing from such broadly semantic traditions as semiotics and structuralism). It is next taken up by Frege. Looking back over his lifework in 1919, he picks out this point as basic to his orientation: "What is distinctive about my conception of logic is that I begin by giving pride of place to the content of the word 'true', and then immediately go on to introduce a thought as that to which the question 'Is it true?' is in principle applicable. So I do not begin with concepts and put them together to form a thought or judgment: I come by the parts of a thought by analysis [*Zerfällung*] of the thought."¹⁹ Already in 1870 in the *Begriffsschrift*, Frege introduces "contents of possible judgment" or "judgeable contents" in the second paragraph and subsequently defines other sorts

of contents in terms of them. In an essay explaining the *Begriffsschrift* he summarizes this approach: "I start out from judgments and their contents, and not from concepts . . . 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."²⁰ The concept of a function, which stands at the center of Frege's technical contribution to semantics, is introduced in the *Begriffsschrift* as an element in his substitutional methodology for decomposing contents of possible judgment.²¹ In the *Grundlagen* Frege continues to follow this Kantian line in insisting that "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 on its parts also their content."²² Frege holds this view because of the importance he assigns to the concept of *truth*; to talk about an expression as contentful is to talk about the contribution it makes to the truth-value of thoughts or propositions in which it occurs.

It is sometimes thought that Frege gave up his commitment to the primacy of the propositional by the late 1880s, when he began to assimilate sentences technically to singular terms under the heading *Eigennamen*, which includes everything except functional expressions. Such a view overlooks the very special role that sentences, as 'names' of truth-values, continue to play for him, even in the *Grundgesetze*. The importance of truth, and therefore of thoughts (the contents expressed by declarative sentences), continues to be emphasized at every stage in Frege's development. In his long 1914 essay entitled "Logic in Mathematics," he is still maintaining "that the name should designate something matters to us if and only if we are concerned with truth."²³ This is the same view that he had endorsed in his classic essay "Über Sinn und Bedeutung": "But now why do we want a proper name to have not only a sense, but also a reference [*Bedeutung*]? Why is the thought not enough for us? Because, and to the extent that, we are concerned with its truth value . . . It is the striving for truth that drives us always to advance from the sense to the reference."²⁴ In the context of such a view it is clear that the assimilation of sentences to singular terms as both having objects as *Bedeutungen* can in no way undercut the fundamental role played by truth-values, and so by the propositional contents that bear them. In that same essay he says that what is needed for a name to have content (express a sense) is that it "belong to a sufficiently complete totality of signs."²⁵ Given his views about identity, this means a system of signs that includes sentences in which the name occurs, and also further sentences that result from them by substituting other names for the ones in question. The totality of signs must include sentences, because to have a sense is to purport to have a *Bedeutung*, and as just indicated, such purport arises only in the context of concern with truth, because "anyone who seriously took the sentence to be true or false would ascribe to the name . . . a *Bedeutung*."²⁶ It is because the

point of deploying concepts in thought and talk is to judge, that is, take or treat judgeable contents as true, that such contents are given pride of place in Frege's scheme. As he says in the 1897 fragment on logic: "Every act of cognition is realized in judgments."²⁷

Indeed, it can be misleading to focus on the concept of *truth* as what enforces attention to sentences. 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. "'True' only makes an abortive attempt to indicate the essence of logic, since what logic is really concerned with is not contained in the word 'true' at all but in the assertoric force with which a sentence is uttered . . . the thing that indicates most clearly the essence of logic is the assertoric force with which a sentence is uttered."²⁸ Talk about the cardinal importance of concern with truth is a dispensable *façon de parler*. What actually matters is the pragmatic attitude of *taking-true* or putting forward *as true*, that is, judging or asserting. Semantic vocabulary is used merely as a convenient way of making explicit what is already implicit in the force or significance that attaches to the content of a speech act or attitude. (An account of just how this explicitation works is offered in Chapter 5, where specifically semantic vocabulary, paradigmatically 'true' and 'refers', is discussed.)

The point that the contents expressed by sentences must play a privileged explanatory role because it is to sentences that pragmatic force attaches has been brought home most forcefully by the later Wittgenstein. The use of sentences is prior in the order of explanation to the use of subsentential expressions because sentences are the only expressions whose utterance "makes a move in the language game." Sentences are expressions whose unembedded utterance performs a speech act such as making a claim, asking a question, or giving a command. That is why even when such a speech act is performed by an utterance that does not manifest the syntactic complexity typical of sentences (a shout of "Rabbit!" or "Fire!" for instance), the utterance should nonetheless be interpreted as a one-word sentence, as meaning what we might express by "Look at the rabbit!" or "There is a fire!"

Referring to something, indicating or naming it, is also something one can do with linguistic expressions; it is a speech act one can perform. But these belong to a class of speech acts that is in an important sense derivative from or parasitic on speech acts involving sentences, paradigmatically claiming, asserting, or putting forward as true. In order to use an expression as a name, to refer to or pick out an object with it, one must be able to use the name to *say* something (paradigmatically, to assert something) about the object referred to, indicated, or named. The significance of taking or treating something as a name, as purporting to refer to an object, consists in how one takes it to be proper to use the expression, and the use of expressions as names is unintelligible except in the context of using expressions containing them as sentences.²⁹

3. *Semantics Must Answer to Pragmatics*

The primacy of propositional intentional contents also shows up if one considers cases in which the use of language is not to the fore. Intentional interpretation of nonlinguistic organisms—intentional explanation of their behavior by attributing beliefs and desires that make what they do intelligible—also depends on attributing *propositionally* contentful states, attitudes, and performances. Behavior is made intelligible by exhibiting it as rational, given various beliefs and pro-attitudes, and to do that is to exhibit a piece of practical reasoning that is taken somehow to stand behind or be implicit in the behavior. The imputed reasoning shows why an organism with the states or attitudes that provide the premises *ought*, rationally, to behave in the way specified by the conclusion. But *what can serve as a premise in reasoning must have a propositional content*. This point is so important to the present project that the rest of this chapter is devoted to motivating the treatment of this feature, in the next chapter, as a defining characteristic distinctive of the propositional. The intentional interpreter attributes to the cat the belief *that* there is a mouse around the corner from it, and the desire *that* it catch the mouse, and so on. Attributing intentional states so as to render behavior intelligible in the light of them requires attributing propositional contents to them.³⁰ So propositional contents have a pragmatic priority, not only in the setting of assessments of the significance of speech acts, but also in the setting of attributions of intentional states that do not evidently depend on linguistic practices.

Semantics must answer to pragmatics. The theoretical point of attributing semantic content to intentional states, attitudes, and performances is to determine the pragmatic significance of their occurrence in various contexts. This means settling how linguistic expressions of those contents are properly or correctly used, under what circumstances it is appropriate to acquire states and attitudes with those contents, and how one then ought or is obliged to go on to behave. It is specifically *propositional* contents that determine these pragmatic significances, so it is specifically propositional contents that it is the task of semantic explanatory theories to attribute. Semantic contents corresponding to *subsential* expressions are significant only insofar as they contribute to the determination of the sorts of semantic contents expressed by full sentences. The pragmatic priority of sentence-use to name-use enforces a certain semantic explanatory priority of the contents expressed by sentences to those expressed by names. The task of the next chapter is to develop an account of the practices of using expressions as sentences—paradigmatically to make claims and so to confer specifically propositional contents on those expression uses and on the states and attitudes associated with them (to use them *as* having such contents).

What the theorist associates with states and expressions deserve to count as *semantic* contents only insofar as they play the right sort of role in determining the proprieties of practice governing those states and expres-

sions. It is possible to associate all sorts of abstract objects with strings of symbols in formalized languages, from sets of models to Gödel numbers. Such an association amounts to specifically *semantic* interpretation just insofar as it serves to determine how those strings are correctly used. For example, Tarski's mapping of well-formed formulas of the first-order predicate calculus onto topological domains qualifies as a semantic interpretation of them only because he can derive from it a notion of valid inference, a way of telling what follows from what—that is, a notion of their correct use. Apart from that, it would just be one more algebraic homomorphism.

4. *Two Mistakes the Designational Model Invites*

An account of content in terms of representation must satisfy the requirement that it must show how semantic content so construed matters for the pragmatic significance of what it is associated with. For the reasons indicated above, this demand focuses attention to begin with on the representational rendering of specifically propositional contents. Two difficulties arise at this point: it is not clear how to derive a notion of propositional contentfulness from the designational representational model, and construing content in representational terms requires supplementation by a further story to get to the proper use of contentful expressions and the correct circumstances and consequences of being in contentful states. (In contrast, the explanatory strategy pursued in Chapter 3 begins with an account of the practices within which producing a performance or altering an attitude can have the pragmatic force or significance of making a claim or judgment; the notion of propositional contentfulness is then understood as what is expressed by such acts.)

On the first point, the pre-Kantian representationalist tradition offers no useful account of what is represented by judgments. For this tradition, representational relations hold between *things*. This categorial nominalism of the designational model extends even to predicates, which are understood as 'general names,' standing for universals in the same sense in which singular terms stand for particular objects. Not until Frege's semantic interpretation of predicates as corresponding to functions—and hence as not being names of any sort—would the idea of semantic relations that are not assimilable to the name/named model enter the tradition.

Applied to propositional contents, the hegemony of the designational semantic model results in two characteristic mistakes: assimilating sentences to complex names, and assimilating judging to predicating. Kant provides the raw materials needed to move beyond these conceptions, but even he is not able to free himself entirely from them. That the first is a mistake becomes clear in the context of an attempt to explain the difference between *referring to a complex object*, for instance a squiggly blue line between a round dot and a square one, and *stating a fact* about its components, for instance *saying*

that the squiggly blue line is between a round dot and a square one.³¹ This crucial difference can be elided by an incautious assimilation of each to a generic notion of *representing*, for in each case the speaker can be said to be representing something.

This difficulty is merely relocated by the introduction of a notion of state of affairs defined as the sort of thing that is represented by utterances that purport to state facts. Making this move is a version of the attempt to solve ontologically the problem of distinguishing referring from saying or stating; the idea is that each is representing, and the specific differences between them are a matter of the kind of thing represented. At the least, such a strategy demands a careful account of the relation between complex objects and the corresponding states of affairs. Any account along these lines of discourse that purports to state facts by the assertive utterance of declarative sentences is also obliged to tell a story about the states of affairs corresponding to normative claims—for instance to the claim that Kant ought not to have written *Dreams of a Spirit-Seer*, or that anyone committed to the claim that snow is white is committed to the claim that snow is spatially extended.³²

The second mistake mentioned above as consequent on the unfortunate sway of designational semantic models corresponds to one way in which the representational tradition in semantics has attempted to acknowledge the special role played by the propositional. This strategy depends on a distinction on the side of the activity of representing, rather than (just) the category of thing represented, by distinguishing between representing as referring or naming, on the one hand, and representing as predicating, on the other. The notion that making a claim can be modeled on representing something (particular) *as* something (general), in the linguistic case picking out an object with a singular term and predicating something of it with a general term, has a distinguished history. The next chapter discusses Frege's decisive demonstration that this approach is a mistake. It is not simply a mistake, however. Looking more closely at what is right about this broadly *classificatory* model of consciousness and at how it can be fixed up in response to some fundamental difficulties provides a way into an idiom for talking about semantic content that does not employ representational vocabulary at the outset.

III. CONCEPTUAL CLASSIFICATION AND INFERENCE

1. *Classification*

An ancient tradition insists that cognition essentially involves generality or universality. Particulars are not directly intelligible as such. Knowing or understanding something particular requires assimilating it to others, taking it to be like them in some way, and so to be an instance of a kind. Kant's account of cognition as beginning with the classification of

intuitions under concepts is a particularly well-developed representative of this tradition. Believing or judging, taking-true in general—for Kant the central sort of cognitive intentional state or act—has as its most basic form subsuming something particular under a universal.³³ In conceiving judgment (the activity of the cognitive subject, the exercise of its faculty of spontaneity, namely understanding) in terms of the *classificatory* employment of concepts, Kant adopts a model that animates as well the thought of the pre-Kantian tradition he inherits—a tradition that had not yet achieved his insight into the privileged role of judgment as the preeminent form of cognitive activity.

That model evidently underlies the epistemologies of both his empiricist and his rationalist predecessors. It forms the common background of their dispute over the source of the universals or concepts by relation to which particulars become intelligible; it is what makes urgent the question whether those universals are formed by abstraction from a more primitive kind of nonconceptual awareness of particulars, or whether on the contrary a grasp of such concepts is a precondition of anything recognizable as awareness at all. Kant follows the rationalists in treating the classificatory account of cognition as a classificatory account of consciousness generally. All awareness is understood as exhibiting the classificatory structure of universal or repeatable concepts subsuming particulars. Where earlier empiricists admit varieties of conscious apprehension short of conceptual comprehension—immediate, nonclassificatory awareness of determinate sense repeatables, for instance—Kant denies apprehension without comprehension, insisting that there must be conceptual classification wherever there is any sort of awareness. Awareness of what is classified and of how things can be classified derives from awareness that consists in classifying.³⁴

A pragmatic version of this classificatory model results if it is de-intellectualized, stripped of residual commitments to understanding concepts as *explicit* to the mind—whether in the Kantian form of rules or recipes for it to follow in its synthesizing activity or in the pre-Kantian form as objects of its direct, nonclassificatory awareness. From such a perspective, the roots of conceptual classification are to be found in treating something in practice as being of a certain kind—taking something (particular) *as* something (universal), by behaving toward it in a way that assimilates it to others. Particular objects are classified as belonging together in some respect by being responded to alike in practice. A respect of similarity in what is responded to then corresponds to a repeatable response. Hegel develops such a pragmatic, indeed naturalized, version of Kant's account in the form of an *erotic* theory of the origins of awareness, an account of animal *desire* as the source of classification. As he puts it, an animal classifies some particular as food when it "falls to without further ado and eats it up."³⁵ Eating something is treating it, responding to it, classifying it in practice *as* food. It exhibits a kind of *practical*, pre-Cartesian awareness of it *as* being of a certain kind. That

repeatable activity on the part of the organism induces a repeatable respect of similarity among the things that tend to elicit that activity.

On this account, classification of particular stimuli as instances of a general kind is *implicit* in what the responding organism *does*. So to each sort of thing that it does, there corresponds a different sort of repeatable proto-concept under which things can be classified: as food, sexual partner, prey, or predator, and so on. That no sort of explicit awareness is presupposed by this sort of implicit practical classificatory awareness or understanding is clear from the fact that all that the latter requires is a reliable differential responsive disposition. For any concrete object displays such dispositions. A chunk of iron reliably responds to some environments by melting, to others by rusting, to still others by falling. In each case it can be understood as classifying that environment, treating it in practice as being of a certain kind, assimilating it to some other possible environments and distinguishing it from others, by responding to it in a certain way.³⁶

The Kantian rationalistic strategy of demarcation by *sapience*, awareness, and consciousness in a sense that requires the application of *concepts* would be trivialized by a classificatory model of the use of concepts that indiscriminately discerns classification according to concepts in the responsive regularities exhibited by the antics of every physical system whatsoever. Classification by the exercise of regular differential responsive dispositions may be a necessary condition of concept use, but it is clearly not a sufficient one. Such classification may underlie the use of concepts, but it cannot by itself constitute discursiveness. The chunk of iron is not conceiving its world as wet when it responds by rusting. Why not? What else must be added to responsive classification to get to an activity recognizable as the application of concepts? What else must an organism be able to do, what else must be true of it, for performances that it is differentially disposed to produce responsively to count as applications of concepts to the stimuli that evoke those responses? One dimension of a reply was indicated in the previous chapter—a normative dimension is required, which can underwrite a distinction between correct and incorrect applications of concepts. But many things can be done correctly or incorrectly. The question being asked now is what it is for what is subject to such assessment to be concept use (rather than, say, hammer use, or tooth use).

2. Inferential Demarcation of the Conceptual

An easy answer is that the response must be the forming of a belief or the making of a claim, acquiring a state or attitude or producing a performance that has an intentional content. This is of course correct, but unhelpful in the current setting. For the question is precisely what is required for a response to count as contentful in this sense. What is wanted is a characterization that does not appeal to semantic concepts such as *content*

and *concept*. If the issue is put in terms of the semantic concept of representation, it takes the form of inquiring as to what more is needed, beyond being a representation in the responsive-classificatory sense, to be a discursive or intentional representation, one that is conceptually contentful. (According to the idiom being employed, implicit grasp of such contents, of the representational purport they consist in—a grasp to be conceived of as some sort of practical mastery, as a kind of know-how—would then in favored cases count as sapient consciousness or awareness of what is represented as exhibiting a certain character.)

A more concrete way to put the question is to ask, What are the salient differences between a measuring instrument, such as a thermometer or spectrophotometer, and an observer who noninferentially acquires beliefs or makes claims about environing temperatures and colors? Artificial instruments differ from other physical systems, such as chunks of iron, only in having been constructed so that some subset of the partition of possible stimuli into equivalence classes according to the distinguishable responses the instruments are disposed to produce corresponds to some distinction of practical or theoretical significance to the user, who thereby attaches some significance to them. Suppose a spectrophotometer is hooked up to a tape recorder in such a way that it produces a noise of the acoustic type "That's red" when and only when it is irradiated with light of the proper frequency. And suppose that a fanatical human red-reporter nearby has just the same responsive dispositions to produce those noises. That is, the two systems are disposed to respond in the same way to the same stimuli, exhibiting the same noninferential circumstances of application for their responsive classifications of things as red. What makes the noise the one produces merely a signal on the basis of which someone else might conclude that something red is present, while the very same noise, reliably elicited under just the same circumstances from the other, counts as a noninferential report, expressive of a perceptually acquired belief, with an intentional content that includes the concept *red*? To vary the case, suppose the reporter's differential responsive dispositions to call things red are matched by those of a parrot trained to utter the same noises under the same stimulation. What practical capacities of the human distinguish the reporter from the instrument or the parrot? What, besides exercise regular differential responsive dispositions, must one be able to do, in order to count as having or grasping concepts, and so as able to perform not only classification but specifically *conceptual* classification?

Putting things this way makes it clear that what is at issue is a kind of *understanding*. The reporter's response is meaningful—not just, as in the case of the measuring instrument or the parrot, to others, but to the responding reporter personally. The spectrophotometer and the parrot do not understand their responses; those responses mean nothing to them, though they can mean something to us. The reporter understands the response he or she makes, attributes to it a kind of significance that the measuring instrument

and the parrot are oblivious to. The challenge is to explain what sort of practical capacity the relevant kind of understanding consists in, without an ultimately circular appeal to semantic concepts such as intentional content, concept-use, or the uptake of representational purport (treated as an explanatory primitive).

The leading idea of the approach to content and understanding to be developed here is due to Sellars. Sellars's suggestion is that the key element missing from the parrot and the measuring instrument—the difference between merely *responsive* classification and *conceptual* classification—is their mastery of the practices of giving and asking for *reasons*, in which their responses can play a role as *justifying* beliefs and claims. To grasp or understand a concept is, according to Sellars, to have practical mastery over the *inferences* it is involved in—to know, in the practical sense of being able to distinguish, what follows from the applicability of a concept, and what it follows from. The parrot does not treat “That’s red” as incompatible with “That’s green,” nor as following from “That’s scarlet” and entailing “That’s colored.” Insofar as the repeatable response is not, for the parrot, caught up in practical proprieties of inference and justification, and so of the making of further judgments, it is not a conceptual or a cognitive matter at all. What the parrot and the measuring instrument lack is an appreciation of the significance their response has as a reason for making further claims and acquiring further beliefs, its role in justifying some further attitudes and performances and ruling out others. Concepts are essentially inferentially articulated. Grasping them in practice is knowing one’s way around the proprieties of inference and incompatibility they are caught up in. What makes a classification deserve to be called *conceptual* classification is its *inferential* role. It is practical mastery of the inferential involvements of a response, the responder’s understanding it in this sense, that makes the response an intentional state or performance—one having a content for the one whose state or performance it is, and not merely for those using it as an indicator.

3. *Holistic Consequences of Inferential Approach to Concepts*

One immediate consequence of such an inferential demarcation of the conceptual is that one must have many concepts in order to have any. For grasping a concept involves mastering the proprieties of inferential moves that connect it to many other concepts: those whose applicability follows from the applicability of the concept in question, those from whose applicability the applicability of the target concept follows, those whose applicability precludes or is precluded by it. One cannot have just one concept. This holism about concepts contrasts with the atomism that would result if one identified concepts with differential responsive dispositions. The capacity to

treat some things as food by eating them need have no particular connection to the capacity to treat other things as dangerous by fleeing them. To treat states or performances as intentionally contentful in the sense of being conceptually articulated involves treating them as situated in a web of proprieties of inferential transitions from one content to another. Knowing one's way around the bit of the web centered on one conceptual content, being able to tell in practice which moves to it and from it are permitted or required and which forbidden, accordingly requires mastery of the proprieties of inference that govern the use of other concepts and contents as well.

By contrast, there is *prima facie* no reason why the fact that some object or property is represented by one simple idea, term, or predicate should be relevant to what is represented by others. Representational relations between nonintentional objects or properties and the intentional representings of them might be treated (as the empiricists in fact treat them) as separate building blocks that, when properly put together, determine what inferences are good in the sense of preserving accuracy of representation. Serving this role seems compatible with these representational relations being quite independent of one another. Knowing what one state or expression represents need convey no information at all about what anything else might represent.

But the inferential notion of semantic content is essentially holistic. Inferences involve both premises and conclusions. The inferential role of one of the premises essentially depends on that of the conclusions, and vice versa. One could not know something about the inferential role of one content without knowing at least something about the inferential roles of others that could be inferred from it, or from which it could be inferred. Contents understood in terms of inferential roles are evidently interdefined in a way in which contents understood in terms of representational purport need not be.

In his masterwork, "Empiricism and the Philosophy of Mind," Sellars exploits this consequence of his insight into the significance of inferential connections to concept-use, even in cases of responsive classification. He argues there that noninferential reports, by which perceptual states are made explicit, cannot constitute an autonomous fragment of a language—one that might be understood though no others are. Observation reports do indeed have a certain priority in the order of *justification* of empirical claims. But they cannot be accorded a similar priority in the order of *understanding* of those claims. Since knowledge requires not only justification but grasp or understanding of the content being justified, there can be no observational knowledge without inference. There can be no purely observational language or set of concepts with respect to which one could then ask whether the decision to append an inferential superstructure is a rational or justifiable one. The rock on which foundationalism founders is accordingly its incapacity to explain what it is to *understand* the significances of elements in the observational justificatory basis. For in order to be able to apply one concept

noninferentially, exercising a disposition to respond differentially to nonlinguistic stimuli, one must be able to apply others inferentially. Unless the response has such an inferential significance, it is not a conceptually contentful response. So the idea of an autonomous language game (or set of practices of applying concepts) consisting entirely of noninferential reports (even of purely mental happenings) is a radical mistake.

The argument does not rule out the possibility of languages or conceptual schemes that are devoid of *theoretical* claims and concepts—that is, that lack concepts that are applicable only as conclusions of inferences. One can have a scheme in which all the concepts have reporting uses and so are in this sense concepts of observables. But they must also have inferential uses. *Red* can be applied either noninferentially, as a response the reporter has been trained to make to a certain kind of visual stimulus, or inferentially, on the basis of entitlement to a prior application of the concept *scarlet*. The conclusion that there can be no conceptually articulated observation apart from inferential capacities holds equally whether what is being reported consists of external observable situations or internal, purely mental happenings. It is this argument that lies at the base of Sellars's critique of broadly Cartesian philosophies of mind.

4. *Inference and Practice*

As ought to be expected from his discussion of the regress-of-rules argument, it is important to Sellars that the inferential conception of concepts connects the grasp or understanding of concepts (the uptake of conceptual content) with a certain kind of practical activity. Inferring is a kind of doing. Acknowledgment of inferential proprieties need not be explicit in the endorsement of rules or principles of inference but may remain implicit in the capacity to take or treat inferential transitions as correct or incorrect in practice. Inferential relations among concepts are implicit in the practice of giving and asking for reasons. The norms that govern these justificatory practices can be understood to confer inferentially articulated contents on the states, attitudes, and performances subject to them: for something to have such content just is for such norms to determine how it is correctly used or manipulated. The status of inference as something that can be *done* accordingly holds out the promise of securing an appropriate relation between pragmatics, the study of the practices, and semantics, the study of the corresponding contents.

Furthermore, because the activity through which the norms get their grip on conceptual contents is construed as *inference*, it is specifically *propositional* contents that in the first instance count as conceptually articulated. Inferential relations hold, in the paradigm case, between contents that are expressed explicitly by declarative sentences. The premises of inferences, and in the central cases their conclusions as well, must be understood to have a

propositional form. Insofar as an independent theoretical grip is possible on the notion of inference, propositional contents can be picked out by appeal to this property. (This is the strategy pursued in the next chapter.) So on an inferential rendering of the conceptual, the sort of doing that inferring yields in a natural way the priority of propositional conceptual contents.

The pragmatic turn aside, this view too is due to Kant. As Sellars puts it: "Kant was on the right track when he insisted that just as concepts are essentially (and not accidentally) items which can occur in judgments, so judgments (and therefore, indirectly concepts) are essentially (and not accidentally) items which can occur in reasonings or arguments."³⁷ The subtlety and sophistication of Kant's concept of representation is due in large part to the way in which it is integrated into his account of the inferential relations among judgments. It remained for Hegel, however, to complete the inversion of the traditional order of semantic explanation by beginning with a concept of experience as inferential activity and discussing the making of judgments and the development of concepts entirely in terms of the roles they play in that inferential activity. Although something like this point had been implicit in Kant's notion of reason as systematicity, it was the young Hegel who first appreciated the line of reasoning, made familiar to us by Quine in "Two Dogmas"—namely, that if the content of a claim must at least determine what follows from it (what else it commits one to), then since what a claim commits one to depends on what collateral commitments are available to serve as additional premises ("auxiliary hypotheses"), the significance of undertaking any particular commitment cannot be determined without appeal to the contents of all those collateral commitments.

Hegel's two central semantic concepts in the *Phenomenology* are both inferential notions. "Mediation," his term for inferential articulation, is derived from the role of the middle term in syllogistic inference. "Determinate negation" is his term for material incompatibility, from which, he takes it, the notion of formal negation is abstracted. The contents of concepts are identified and individuated by the functional roles they play in historically evolving webs constituted by relations of mediation and determinate negation, that is, by their material inferential and incompatibility relations to each other. Hegel's interest in the significance of inference in semantics does not (as with Kant) arise primarily in the investigation of how it might be combined with representationalist insights (although he has something to say about that too). It arises rather in the investigation of how this rationalist insight might be combined with the insights of the Romantic expressivists.

The Romantics are perhaps best known for their rejection, not just of the Enlightenment's representationalism, discussed above, but also for their rejection of the significance it assigns to reason. They sought to displace the general demarcational emphasis on giving and asking for reasons or inquiring after truth, not just the specific version that sought to understand these matters in representational terms. The Romantic recoil from understanding us as representers overshoot that mark and came to rest in an esteem for

feeling and inarticulate empathy and enthusiasm. Hegel saw in inferential notions of content a way to join the Romantics in rejecting representationalism, while parting company with them in their hostility to reason. The result is a synthesis of Enlightenment inferentialism and Romantic expressivism.³⁸

5. *Inferentialism and Representationalism*

Kant, however, did not originate the inferentialist line of semantic thought that Sellars appropriates from him, and that Hegel develops. In a discussion of his break with traditional empiricism, prompted by the issue of the sort of content that ought to be associated with logical, causal, and deontological modalities, Sellars puts the idea that stands at the center of his systematic thought in the form in which it originally occurred to him in the 1930s: "What was needed was a functional theory of concepts which would make their role in reasoning, rather than supposed origin in experience, their primary feature."³⁹ Put this way, the idea forms one of the mainstays of classical rationalism, even in the absence of Kant's insight about the privileged role that must be assigned to judgments on such an inferential-functional approach. Pre-Kantian empiricists and rationalists alike were notoriously disposed to run together causal and conceptual issues, largely through insufficient appreciation of the normative character of the "order and connection of ideas" that matters for the latter. But there is another, perhaps less appreciated, contrast at work here, besides that of the causal and the conceptual. Enlightenment epistemology was always the home for two somewhat uneasily coexisting conceptions of the conceptual. The fundamental concept of the dominant and characteristic understanding of cognitive contentfulness in the period initiated by Descartes is of course *representation*. However there is a minority semantic tradition that takes *inference* rather than representation as its master concept.

Rationalists such as Spinoza and Leibniz accept the central role of the concept of representation in explaining human cognitive activity, but they are not prepared to accept Descartes's strategy of treating the possession of representational content as an unexplained explainer. Each of them develops instead an account of what it is for one thing to represent another, in terms of the *inferential* significance of the representing. They are explicitly concerned (as Descartes is not) to be able to explain what it is for something to be understood, treated, or employed in practice *as* a representing *by* the subject—what it is for it to be a representing *to* or *for* that subject. Their idea is that the way in which representings point beyond themselves to something represented is to be understood in terms of *inferential* relations among representings. States and acts acquire content by being caught up in inferences, as premises and conclusions.⁴⁰ Thus a big divide within Enlightenment epistemology concerns the relative explanatory priority accorded to the concepts of representation and inference.

The British empiricists are more puzzled than Descartes about representational purport, the property of seeming to be about something. But they are clear in seeking to derive inferential relations from the contents of representings, rather than the other way around. In this regard they belong to the still-dominant tradition that reads inferential correctnesses off from representational correctnesses, which are assumed to be antecedently intelligible. The post-Cartesian rationalists, the claim is, give rise to a tradition based on a complementary, semantically reductive order of explanation. These *inferentialists* seek to define representational properties in terms of inferential ones, which must accordingly be capable of being understood antecedently. They start with a notion of content as a matter of what is a reason for what and understand truth and representation as features of ideas that are not only manifested in, but conferred by their role in reasoning. This is the tradition that Sellars inherits and builds on by developing a notion of conceptual content that starts with inferential roles.

IV. MATERIAL INFERENCE, CONCEPTUAL CONTENT, AND EXPRESSION

1. Frege on Conceptual Content

The rationalists' inferential understanding of conceptual content, which Kant inherits and which remains one of the strands from which his systematic semantic tapestry is woven, provides the starting point as well for Frege's semantic investigations. Frege may seem an unlikely heir to this inferentialist tradition. After all, he is usually thought of as the father of the contemporary way of working out the representationalist order of explanation. Its strategy is to start with an independent notion of relations of reference or denotation obtaining between mental or linguistic items and objects and sets of objects in the largely nonmental, nonlinguistic environment. Then it determines from these in the familiar fashion: first truth conditions for the sentential representings built out of the subsentential ones, and then, from these, a notion of goodness of inference understood in terms of set-theoretic inclusions among the associated sets of truth conditions. But insofar as it is appropriate to read this twentieth-century story back into Frege at all (a dangerous and potentially misleading enterprise), it would be possible only beginning with the Frege of the 1890s. He starts his semantic investigations, not with the idea of *reference*, but with that of *inference*. His seminal first work, the *Begriffsschrift* of 1870, takes as its task the explicit expression of *inferential* roles: "In my formalized language [*Begriffsschrift*] . . . only that part of judgments which affects the possible inferences is taken into consideration. Whatever is needed for a correct [*richtig*] inference is fully expressed; what is not needed is . . . not."⁴¹

These inferential roles form the basis of his notion of *content*. It is because the sorts of contents that are associated with expressions are to be defined

in the first place in terms of inference that Frege must insist on the distinction between the sorts of contents that can, and those that cannot, serve as premises and conclusions of inference, and so play the basic sort of inferential roles. "We distinguish contents that *are*, and contents that *are not, possible contents of judgment*."⁴² Frege's Kantian insistence on the priority of the propositional, of judgeable contents, is an aspect of his pursuit of the rationalists' inferentialist order of semantic explanation. He embraces Kant's insight that the notion of content must be made intelligible first for judgments, which alone can figure as premises and conclusions of inference, and only then extended to the contents expressed by fragments of declarative sentences. Recall the passage (already quoted in Section II of this chapter) in which he contrasts his procedure with that pursued by others in the tradition: "In Aristotle, as in Boole, the logically primitive activity is the formation of concepts by abstraction, and judgment and inference enter in through an immediate or indirect comparison of concepts via their extensions . . . I start out from judgments and their contents, and not from concepts . . . I only allow the formation of concepts to proceed from judgments . . . 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."⁴³ It is for this reason that the fundamental definition introducing the notion of "conceptual content" (*begriffliche Inhalt*) (for which, as its name implies, the *Begriffsschrift* is supposed to supply a means of explicit expression) applies only to the contents of possible judgments. It will have to be extended later, by Frege's substitutional methodology, to allow the assignment of *indirectly* inferential roles to *subsential* expressions, according to the contribution their occurrence makes to the directly inferential role (as premise or conclusion) of judgment-expressing sentences in which they occur. The substitutional strategy that Frege devised for quarrying subsentially expressed contents from sententially expressed ones is of the first importance for carrying out the inferentialist semantic explanatory program. Much is made of it in subsequent chapters of this work. Before Frege, one could only hope that there was some way of bridging this gap.

That the target notion of content is specifically *conceptual* content is accordingly not to be understood in terms of some antecedent notion of concepts. Rather, the conceptual is explicitly construed in *inferential* terms: "There are two ways in which the content of two judgments may differ; it may, or it may not, be the case that 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 two propositions 'the Greeks defeated the Persians at Plataea' and 'the Persians were defeated by the Greeks at Plataea' differ in the former way; even if a slight difference of sense is discernible, the agreement in sense is preponderant. Now I call that part of the content that is the same in both the

conceptual content. Only this has significance for our symbolic language [*Begriffsschrift*]."⁴⁴ Two claims have the same conceptual content if and only if they have the same inferential role: a good inference is never turned into a bad one by substituting one for the other. The fundamental semantic assignment of conceptual content to judgments is derived from the ultimately pragmatic notion of *correctness* of *inference*. This derivation is the first application of the substitutional methodology: semantically assimilating expressions accordingly as substitution of one for another preserves some semantically relevant property. In this case (prior to the others in the order of explanation), the semantically relevant invariant is propriety of inference.

This way of specifying the explanatory target to which semantic theories, including referential ones, are directed is picked up by Frege's student Carnap, who in *The Logical Syntax of Language* defines the content of a sentence as the class of nonvalid sentences that are its consequences (that is, can be inferred from it). Sellars in turn picks up the idea from him, as his references to this definition indicate.⁴⁵ As will emerge, an important feature of Carnap's definition is the appeal to *nonvalid* consequences. In this way what pertains to the *content* of a claim is distinguished from what pertains to its *form*.

This distinction is operative in the *Begriffsschrift* as well. Yet when Frege wants to be clear about what is expressed by even the purely formal assertions appearing in proofs about the expressive capacity of the *Begriffsschrift* itself, he does so by specifying their inferential role, restricting himself in this case to inferences whose propriety is underwritten by their form alone. So each assertion is introduced by displaying a proof of it from already-established assertions, thereby exhibiting the premises from which it follows as conclusion. Showing what a claim follows from is not sufficient to specify its inferential role, however. It matters as well what follows from it. Indeed, Frege often complains (for instance in the *Grundlagen*) about systems that introduce definitions that are never then employed in subsequent demonstrations. These provide a case where looking at inferential consequences is particularly important; since definitions do not have inferential antecedents, if their inferential consequents are not specified, their content is left entirely indeterminate. In order to complete the specification of the inferential roles of the assertions of the system he presents, Frege appends to the *Begriffsschrift* a list indicating for each assertion all of the subsequent assertions in whose proof it is used as premise. That is, he specifies for each assertion what follows from it (together with other assertions, of course) as well as what it follows from. In this way he makes explicit the inferential roles, and so the conceptual contents, conferred on the judgments he puts forward by the purely formal reasoning involving them that is displayed in his book.

In contrast to his original procedure, the tradition Frege initiated in the 1890s makes truth, rather than inference, primary in the order of semantic explanation. Dummett says of this shift: "In this respect (and in this respect

alone) Frege's new approach to logic was retrograde. He characterized logic by saying that, while all sciences have truth as their goal, in logic truth is not merely the goal, but the object of study. The traditional answer to the question what is the subject-matter of logic is, however, that it is, not truth, but inference, or, more properly, the relation of logical consequence. This was the received opinion all through the doldrums of logic, until the subject was revitalized by Frege; and it is, surely, the correct view."⁴⁶ And again: "It remains that the representation of logic as concerned with a characteristic of sentences, truth, rather than of transitions from sentences to sentences, had highly deleterious effects both in logic and in philosophy. In philosophy it led to a concentration on logical truth and its generalization, analytic truth, as the problematic notions, rather than on the notion of a statement's being a deductive consequence of other statements, and hence to solutions involving a distinction between two supposedly utterly different kinds of truth, analytic truth and contingent truth, which would have appeared preposterous and irrelevant if the central problem had from the start been taken to be that of the character of the relation of deductive consequence."⁴⁷ The important thing to realize is that the Frege of the *Begriffsschrift* has not yet made this false step. Of course, adopting a semantic order of explanation that begins with proprieties of inference requires both an account of those proprieties (that is, an account of its raw materials) and an account of how talk about truth is eventually to be construed in these terms (that is, an account of its consequences). This is the strategy pursued in this work. The first of these challenges is responded to in Chapters 3 and 4, and the second in Chapter 5.

There are two further points to keep in mind regarding this passage of Dummett's. First, shifting from concern with *inference* to concern with *truth* is one move; understanding *truth* in terms of prior primitive *reference* relations involving objects and properties is another. Since the mature Frege treats truth as indefinable and primitive, the extraction of a representationalist commitment even from the texts of the 1890s requires further showing (compare Davidson's truth-without-reference view in our own day). Second, understanding the topic of logic in terms of inference is not the same as seeing it in terms of *logical* inference, or of "deductive consequence," as Dummett puts it (see the discussion of "formalism" about inference, below). The view propounded and attributed to Frege below is a different one—and from the contemporary vantage point it is a more surprising one than the one that Dummett endorses here.

2. *Material Proprieties of Inference and the Dogma of Formalism*

The kind of inference whose correctnesses essentially involve the conceptual contents of its premises and conclusions may be called, following Sellars, "*material* inference."⁴⁸ As examples, consider the inference from

"Pittsburgh is to the West of Philadelphia" to "Philadelphia is to the East of Pittsburgh," the inference from "Today is Wednesday" to "Tomorrow will be Thursday," and that from "Lightning is seen now" to "Thunder will be heard soon." It is the contents of the concepts *West* and *East* that make the first a good inference, the contents of the concepts *Wednesday*, *Thursday*, *today*, and *tomorrow* that make the second inference correct, and the contents of the concepts *lightning* and *thunder*, as well as the temporal concepts, that underwrite the third. Endorsing these inferences is part of grasping or mastering those concepts, quite apart from any specifically logical competence. From the point of view of a familiar sort of semantics (different from that to be explored here), one could say that the set of possible worlds in which the premises of these inferences are true is a subset of the set of possible worlds in which their conclusions are true. Since neither the premises nor the conclusions of such inferences employ logical concepts, it seems appropriate to distinguish them from inferences whose correctness depends only on logical form.

Often, however, *inferential* articulation is identified with *logical* articulation. Material inferences are then treated as a derivative category. The idea is that being *rational*—mastering proprieties of inference and so being subject to the force of the better reason—can be understood as a purely *logical* capacity. In part this tendency is encouraged by merely verbally sloppy formulations of the crucial difference between the inferential force of reasons and the physically efficacious force of causes: formulations that render it as the difference between 'logical' and 'natural' compulsion. Mistakes ensue, however, if the concept *logical* is employed with these circumstances of application conjoined with consequences of application that restrict the notion of the logical force of reasons to *formally* valid inferences. The substantial commitment that is fundamental to this sort of approach is what Sellars calls "the received dogma . . . that the inference which finds its expression in 'It is raining, therefore the streets will be wet' is an enthymeme."⁴⁹

According to this line of thought, wherever an inference is endorsed, it is because of belief in a conditional. Then the instanced inference is understood as implicitly involving the conditional "If it is raining, then the streets will be wet." With that "suppressed" premise supplied, the inference is an instance of the formally valid scheme of conditional detachment. The "dogma" expresses a commitment to an order of explanation that treats all inferences as good or bad solely in virtue of their form, with the contents of the claims they involve mattering only for the truth of the (implicit) premises. According to this way of setting things out, there is no such thing as material inference. This view—which understands "good inference" to mean "formally valid inference," postulating implicit premises as needed—might be called a *formalist* approach to inference. It trades primitive goodnesses of inference for the truth of conditionals. Doing so is taking a retrograde step that corresponds to the one Dummett complains about. The grasp of logic

that is attributed must be an *implicit* grasp, since it need be manifested only in distinguishing material inferences as good and bad, not in any further capacity to manipulate logical vocabulary or endorse tautologies involving them. But what then is the explanatory payoff from attributing such an implicit logical ability rather than just the capacity to assess proprieties of material inference?

It is worth considering an example of how formalist presuppositions can be embodied misleadingly in vocabulary. Here is Dennett in "Intentional Systems":

Earlier I alleged that even creatures from another planet ["in virtue of their rationality"] would share with us our beliefs in logical truths; light can be shed on this claim by asking whether mice and other animals, in virtue of being intentional systems, also believe the truths of logic. There is something bizarre in the picture of a dog or mouse cogitating a list of tautologies, but we can avoid that picture. The assumption that something is an intentional system is the assumption that it is rational; that is, one gets nowhere with the assumption that entity x has beliefs $p, q, r \dots$ unless one also supposes that x believes what follows from $p, q, r \dots$; otherwise there is no way of ruling out the prediction that x will, in the face of beliefs $p, q, r \dots$ do something utterly stupid, and, if we cannot rule out that prediction, we will have acquired no predictive power at all. So whether or not the animal is said to believe the truths of logic, it must be supposed to follow the rules of logic.⁵⁰

Dennett understands intentionality in terms of rationality (as the view being developed here does), and understands rationality in terms of the discrimination in practice of good inferences ("what follows") from bad ones (as the view being developed here does). But there is a slide here from "follows" to "logically follows." No justification is offered for the move, first, from discriminating good from bad inferences to the need for any specifically logical capacity or, second, for the move from logical capacity to belief in logical truths. Perhaps appropriate (even logically valid) inferences can be endorsed without commitment to the corresponding (logical) conditional truths.

On the first point: perhaps there are good nonlogical inferences, and rationality consists in the way discriminating them matters to one's deliberations and assessments. Why should "following the rules of logic" be either necessary or sufficient for this discrimination? In any case, it was argued in Chapter 1 that one ought to distinguish both exhibiting a regularity and acknowledging a norm implicitly in one's practice (two construals of discriminating good from bad inferences) from following a rule. On the second point, Dummett was cited above as pointing out that defining logical consequence in terms of logical truth is neither a trivial nor a harmless move.

In fact Dennett (and in this regard he is typical) thinks of this way of putting things as a harmless *façon de parler*, warranted by a general inter-

changeability of talk of endorsing inferences and talk of believing conditionals. The "belief in logical truths," or even, less committally, endorsement of logically good inferences, that he has in mind is implicit in practical discriminations. The passage continues: "Surely our mouse follows or believes in *modus ponens*, for we ascribed to it the beliefs: (a) there is a cat to the left, and (b) if there is a cat to the left, I had better not go left, and our prediction relied on the mouse's ability to get to the conclusion." What was actually attributed to the mouse is a belief with content (a) and a desire to avoid the cat. Citing its intelligent behavior licenses the attribution of a practical inference. It does not, by itself, tell for or against expressing that inference as a material inference or as detachment from an endorsed conditional. Why should all inferences be assimilated to detachments, or other formal logical rules of inference? Dennett's justification is that "in general there is a trade-off between rules and truths; we can suppose *x* to have an inference rule taking *A* to *B* or we can give *x* the belief in the 'theorem': if *A* then *B*. As far as our predictions are concerned, we are free to ascribe to the mouse either a few inference rules and belief in many logical propositions, or many inference rules and few if any logical beliefs."⁵¹

The conditional beliefs that can be traded off for endorsements of inferences should not be called "logical" beliefs simply because they concern inferences. Though they involve logical concepts, namely the conditional, they are not in general logically true. Indeed, Dennett continues: "We can even take a patently nonlogical belief like (b) and recast it as an inference rule taking (a) to the desired conclusion." To do so would be to establish or endorse a *material* correctness of inference, what Dennett calls "a set of nonlogical inference rules." Once the possibility of this sort of inference is acknowledged, inferential formalism surrenders a priori privileges and must contend with inferential materialism for privileges of explanatory priority. According to the famous argument of Lewis Carroll in "What the Tortoise Said to Achilles," as Dennett acknowledges, some inferential commitments ("rules of inference") must be attributed if any consequences are to be licensed by the attribution of beliefs, even conditional beliefs. So there must be "rules" as well as "truths." However, once the purely formal-logical inferences are allowed (paradigmatically detachment inferences licensed by conditionals), accounts of rational performance can take the form either of attributions of endorsements of material inferences or of conditional propositions, as might be theoretically convenient for other reasons. Either decision ought to be justified.

What considerations ought to persuade a theorist to accord explanatory priority to the attribution of material inferential commitments or to the attribution of conditional propositional commitments, and so to treat material or formal inference as fundamental? Dennett's answer is: "If we found an imperfectly rational creature whose allegiance to *modus ponens*, say, varied with the subject matter, we could characterize that by excluding *modus*

ponens as a rule and ascribing in its stead a set of nonlogical inference rules covering the modus ponens step for each subject matter where the rule was followed."⁵² This is a formalist position, in that all inferences are assimilated to detachments and are understood as involving, at least implicitly, endorsements of conditionals whose logical content explicitly relates premises and conclusions. The only concession to material inferences arises in the possibility of licensing detachment in a retail, content-respecting fashion, rather than wholesale, in a purely formal logical way. But why should this model be employed? Why should all goodness of inference be seen as logical goodness, even at the cost of postulating "implicit" premises involving logical concepts?

What is at issue is two different ways of understanding the relation between something implicit and an explicit expression of it. It is possible to agree with the formalist in understanding conditionals as inference licenses, which make explicit in the content of a claim what is implicit in the endorsement of an inference, without going on to construe all inferences as involving the use of conditionals. The question is how one ought to construe the relation between what is explicit in the form of a rule or principle (in this case a conditional claim) and what is implicit in proprieties of practice (in this case in the endorsement of an inference). The formalist line of thought begins with explicit propositional licenses that license inferences in virtue of their logical form. Material inferences (say from rain to wet streets or vice versa) are understood privatively: as enthymemes resulting from the suppression or hiding of one of the premises required for a proper warrant. Opposed to this might be a pragmatist line of thought, beginning with material inferences—that is, nonlogical, content-based reasoning. It would then be necessary to explain how logical vocabulary such as the conditional is to be understood as permitting the expression of those implicit inferential commitments in an explicit fashion—that is, as judgeable, claimable, believable contents, as the contents of potential propositional commitments.

There are general reasons to prefer an order of explanation that begins with what is implicit in practice (what people do) and proceeds to an account of what they explicitly believe or say, over one taking the opposite tack. Only in this way can one hope to understand believing or saying in terms of more primitive capacities (knowing-that in terms of knowing how). That asymmetry manifests itself in this case in the question of how one understands logical concepts or the use of logical vocabulary. On the formalist line, anything that has any inferential capacities at all is credited with mastery of a battery of logical concepts and the corresponding inference rules without which they would be without content. These can be thought of as introduction and elimination rules, of which detachment is a cardinal example. Logical concepts are quite different from others in being presupposed by all contentful concepts and inferences. It is a short step from treating mastery of these concepts as implicit in inferential abilities to treating it as an innate

presupposition of them. This sort of thing gave the classical rationalists a bad name. Kant rescued them by insisting that it is the formality of logical (and, more controversially, transcendental) concepts that entitles them to a special status that would indeed be absurd for ordinary contentful concepts. Assessing the Kantian formalist move requires looking more closely at what is being said when an inference is described as being valid in virtue of its logical form.

3. *Conceptual Content and Material Inference*

Before looking at how logical concepts might function to make explicit conceptual contents that are implicit in practical proprieties of inference, however, it is worth looking more closely at the relation between inference and content. The picture being developed is one according to which materially good inferences correspond to the conceptual content of nonlogical expressions, while inferences valid in virtue of their logical form alone correspond to the conceptual content of purely logical expressions. This can be approached by considering, to begin with, the notion of *material* inferences: inferences whose propriety essentially involves the nonlogical conceptual content of the premises and conclusions. The approach Sellars endorses is best understood by reference to the full list of alternatives he considers:

We have been led to distinguish the following six conceptions of the status of material rules of inference:

(1) Material rules are as essential to meaning (and hence to language and thought) as formal rules, contributing to the architectural detail of its structure within the flying buttresses of logical form.

(2) While not essential to meaning, material rules of inference have an original authority not derived from formal rules, and play an indispensable role in our thinking on matters of fact.

(3) Same as (2) save that the acknowledgment of material rules of inference is held to be a dispensable feature of thought, at best a matter of convenience.

(4) Material rules of inference have a purely derivative authority, though they are genuinely rules of inference.

(5) The sentences which raise these puzzles about material rules of inference are merely abridged formulations of logically valid inferences. (Clearly the distinction between an inference and the formulation of an inference would have to be explored.)

(6) Trains of thought which are said to be governed by "material rules of inference" are actually not inferences at all, but rather activated associations which mimic inference, concealing their intellectual nudity with stolen "therefores."⁵³

His own position is that an expression has conceptual content conferred on it by being caught up in, playing a certain role in, material inferences: "It is

the first (or 'rationalistic') alternative to which we are committed. According to it, material transformation rules determine the descriptive meaning of the expressions of a language within the framework provided by its logical transformation rules . . . In traditional language, the 'content' of concepts as well as their logical 'form' is determined by the rules of the Understanding."⁵⁴

Sellars, in arguing that material inferences are essential to the meaning (content) of nonlogical locutions, cites a phenomenon that is as important to the expressivist picture of logical concepts as it is to the materialist conception of inference presupposed by inferentialist approaches to conceptual content. Sellars's argument that material inferences are essential to the meaning (content) of nonlogical locutions depends on a central conceptual phenomenon. He argues for the theoretical indispensability of a conception of material inferences in terms of the practical indispensability of what is made explicit by a certain familiar kind of vocabulary. His argument is attributed to an interlocutor who maintains that:

such subjunctive conditionals as "If I had released this piece of chalk, it would have fallen," and "If there were to be a flash of lightning, there would be thunder" . . . [must be interpreted] as expressions of material rules of inference . . . He therefore claims to have shown beyond reasonable doubt not only that there are such things as material rules of inference, but, which is far more important, that they are essential to any conceptual frame which permits the formulation of such subjunctive conditionals as do not give expression to logical principles of inference. Since we are all conscious of the key role played in the sciences, both formal and empirical, in detective work and in the ordinary course of living by subjunctive conditionals, this claim, if substantiated, would indeed give a distinguished status to material rules of inference.⁵⁵

He concludes: "Now, unless some other way can be found of interpreting such subjunctive conditionals in terms of logical principles of inference, we have established not only that they are the expression of material rules of inference, but that the authority of these rules is not derivative from formal rules. In other words, we have shown that material rules of inference are essential to the language we speak, for we make constant use of subjunctive conditionals."⁵⁶

The point is not the indispensability of the vocabulary of conditionals that permit detachment inferences even with counterfactual premises. It is the indispensability of what those conditionals express: the implicit proprieties of material inference that they help make explicit. "Even though material subjunctive conditionals may be dispensable, permitting the language to be extensional, it may nevertheless be the case that the function performed in natural languages by material subjunctive conditionals is indispensable."⁵⁷

The material inferences codified in subjunctive conditionals are inferential involvements that are essential to the contents of the concepts used in science and everyday life. These are not logically valid inferences. But logical

vocabulary, subjunctive conditionals, can be used to express these material inferential relations. Without such vocabulary, the inferences can still be endorsed. With it, those content-generating inferential endorsements can be made explicit as the content of a claim or propositional endorsement.

4. *From Material to Formal Proprieties of Inference*

Should inferentialist explanations begin with inferences pertaining to propositional *form*, or those pertaining to propositional *content*? One important consideration is that the notion of formally valid inferences is definable in a natural way from that of materially correct ones, while there is no converse route. For given a subset of vocabulary that is privileged or distinguished somehow, an inference can be treated as good in virtue of its form, with respect to that vocabulary, just in case it is a materially good inference and cannot be turned into a materially bad one by substituting nonprivileged for privileged vocabulary, in its premises and conclusions. This is another application of the substitutional methodology Frege employs in individuating the conceptual contents of judgments, and again in discerning indirectly conceptually contentful components within them. All it requires is a partition of vocabulary into two kinds: those that are to be held fixed and those that are to be regarded as replaceable. Call the kind of vocabulary that is to be held fixed the *K-vocabulary*. The general structure of formality definitions is then that the set of *K-valid* inferences (those that will be understood as good in virtue of their *K-form* alone) comprises those that meet the two conditions of being inferences that (1) are good inferences and (2) cannot be turned into bad inferences by substituting non-*K* for *K* vocabulary.

Clearly, what inferences are treated as valid in virtue of their form by such a procedure depends on how the vocabulary is divided into the two kinds. In the limit, if *all* the vocabulary were treated as irreplaceable, no substitutions of non-*K* for *K* vocabulary would be possible, and a fortiori none could turn a correct inference into one that is not correct. So *all* materially good inferences would count as good in virtue of their *K-form*, in the case where *K* comprises the whole vocabulary over which the field of inferences is defined. At the opposite end of the spectrum, if *no* vocabulary is treated as irreplaceable, then if there were any bad inferences at all, none of the good inferences would count as good in virtue of their *K-form*. For all could be turned into bad inferences by some substitution or other.

If the *K-vocabulary* (that which is not substituted for) is *logical* vocabulary, then the good inferences whose correctness is invariant under substitution of non-*K* for *K* vocabulary (nonlogical for logical vocabulary) are the *logically* valid inferences—namely those that are good in virtue of their *logical* form. (Quine recommends this Fregean substitutional way of thinking about logical form, although he appeals to *truth* rather than propriety of

inference as the semantically relevant invariant whose preservation is at issue.) But this substitutional conception of what it is for an inference to be good in virtue of its form is not essentially restricted to a notion of *logical* form. If one picks out specifically zoological vocabulary or moral vocabulary or theological vocabulary to play the role of the distinguished *K*-vocabulary, the substitutional mechanism will take as its input a practical classification of inferences into good or bad, correct or incorrect, and yield as its output a distinguished set of inferences that are not just good, but are good in virtue of their zoological, moral, or theological form. The mechanism is perfectly general.

It follows that on this way of thinking about things, logical vocabulary cannot be picked out by appeal to its formality or by its involvement in formal proprieties of inference. If it is specifically logical form that is of interest, then one must antecedently be able to distinguish some vocabulary as peculiarly logical. That done, the Fregean semantic strategy of looking for inferential features that are invariant under substitution yields a notion of logically valid inferences. So the formal goodness of inferences derives from and is explained in terms of the material goodness of inferences, and so ought not to be appealed to in explaining it. And logical vocabulary must be picked out in some way that does not appeal to inferences that are formally valid or good in virtue of their form. Frege's way of specifying the characteristic linguistic role in virtue of which vocabulary qualifies as logical is discussed below.

5. Sellars on Expressive Rationality

So far two related claims have been introduced: that conceptual contents are inferential roles, and that the inferences that matter for such contents in general must be conceived to include those that are in some sense materially correct, not just those that are formally valid.⁵⁸ It will be argued in a moment that a commitment to the second of these, no less than the first, is to be found already in Frege's early writings, though not in the developed form to which Sellars brings it. But in both thinkers these ideas are combined with a third, which makes this line of thought especially attractive. In one of his early papers, Sellars introduces the idea this way: "In dealing with such situations [attempts to justify acceptance of a law by means of an argument from instances], philosophers usually speak of inductive arguments, of establishing laws by induction from instances . . . I am highly dubious of this conception. I should be inclined to say that the use Jones will make of instances is rather in the nature of Socratic method. For Socratic method serves the purpose of making explicit the rules we have adopted for thought and action, and I shall be interpreting our judgments to the effect that *A* causally necessitates *B* as the expression of a rule governing our use of the terms '*A*' and '*B*'."⁵⁹ Sellars understands such modal statements as inference

licenses, which formulate as the content of a claim the appropriateness of inferential transitions. More than this, he understands the function of such statements to be making explicit, in the form of assertible rules, commitments that had hitherto remained implicit in inferential practices. Socratic method is a way of bringing our practices under rational control, by expressing them explicitly in a form in which they can be confronted with objections and alternatives, a form in which they can be exhibited as the conclusions of inferences seeking to justify them on the basis of premises advanced as reasons, and as premises in further inferences exploring the consequences of accepting them.

In the passage just cited, Sellars tells us that the enterprise within which we ought to understand the characteristic function of inductive inference is a form of rationality that centers on the notion of expression: making explicit, in a form that can be thought or said, what is implicit in what is done. This is a dark and pregnant claim, but it epitomizes a radical and distinctive insight. What follows is intended to shed some light on it and its role in an inferentialist vision of things. The general idea is that the paradigmatically rational process that Sellars invokes under the heading of "Socratic method" depends upon the possibility of making *implicit* commitments *explicit* in the form of claims. Expressing them in this sense is bringing them into the game of giving and asking for reasons as playing the special sort of role in virtue of which something has a conceptual content at all—namely an inferential role, as premise and conclusion of inferences.

This is distinct from (but obviously related to) the sort of rationality that then consists in making the appropriate inferential moves. Even totalitarian versions of the latter—for instance those that would assimilate all goodness of inference to logical validity, or to instrumental prudence (that is, efficiency at getting what one wants)—depend upon the possibility of expressing considerations in a form in which they can be given as reasons, and reasons demanded for them. All the more does Socratic reflection on our practices, particularly on those material-inferential practices that determine the conceptual contents of thoughts and beliefs, depend on the possibility of their explicit expression. Here is another early (perhaps equally dark) statement of this important Sellarsian theme:

Now, among the linguistic activities which can be discriminated are the 'explicative' or 'analytic' which, to use Ayer's phrase 'elucidate the proper use' of linguistic expressions. Furthermore the anthropologist . . . can distinguish within language activity between that which "deals directly with the environment" and that which attempts to mirror, within language itself, the relation of language to the world. In connection with this Fichtean self-diremption, the language user makes use of such words as 'means', 'true', 'verified' and so on. This is linguistic activity as semantic and pragmatic metalanguage. But the language

activity of human organisms can achieve an even greater degree of internal complexity, such as comes out most clearly in the 'explicative' metalinguistic activity of the logician and epistemologist, but is also to be found, highly confused, in more practical beings.⁶⁰

6 . *The Expressive Project of the Begriffsschrift*

To begin to explicate this notion of explication, it is helpful to return to the consideration of the young Frege's inferentialist program. Frege's *Begriffsschrift* is remarkable not only for the inferential idiom in which it specifies its topic but equally for how it conceives its relation to that topic. The task of the work is officially an *expressive* one—not to *prove* something, but to *say* something. Frege's logical notation is designed for *expressing* conceptual contents, making *explicit* the inferential involvements that are *implicit* in anything that possesses such content. As the passage quoted above puts it: "Whatever is needed for a correct inference is fully expressed." Talking about this project, Frege says: "Right from the start I had in mind the expression of a content . . . But the content is to be rendered more exactly than is done by verbal language . . . Speech often only indicates by inessential marks or by imagery what a concept-script should spell out in full."⁶¹ The concept-script is a formal language for the explicit codification of conceptual contents. In the preface to the *Begriffsschrift*, Frege laments that even in science, concepts are formed haphazardly, so that the ones employing them are scarcely aware of what they mean, of what their content really is. When the correctness of particular inferences is at issue, this sort of unclarity may preclude rational settlement of the issue. What is needed, he thinks, is a notation within which the rough-and-ready conceptual contents of the sciences, beginning with mathematics, can be reformulated so as to wear their contents on their sleeves. His explanatory target avowedly concerns a sort of inference, not a sort of truth, and the sort of inference involved must be content-conferring material inferences, whose propriety is determined before logical vocabulary comes on the scene, not the derivative formal ones whose propriety is underwritten by the use of that vocabulary.

Frege explicitly contrasts his approach with that of those, such as Boole, who conceive their formal language only in terms of formal inference, and so express no material contents: "The reason for this inability to form concepts in a scientific manner lies in the lack of one of the two components of which every highly developed language must consist. That is, we may distinguish the formal part . . . from the material part proper. The signs of arithmetic correspond to the latter . . . In contrast, Boole's symbolic logic only represents the formal part of the language."⁶² Frege's own project is to express the contents that make up the material part of the language, not just the "formal cement that can bind these stones together": "My concept-script has a more far-reaching aim than Boolean logic, in that it strives to make it

possible to present a content when combined with arithmetical and geometrical signs . . . It is in a position to represent the formation of the concepts actually needed in science."⁶³ It is the wider domain to which his expressive ambition extends that Frege sees as characteristic of his approach.

Since contents are determined by inferences, expressing inferences explicitly will permit the expression of any sort of content at all: "It seems to me to be easier still to extend the domain of this formula language to include geometry. We would only have to add a few signs for the intuitive relations that occur there . . . The transition to the pure theory of motion and then to mechanics and physics could follow at this point."⁶⁴ Indeed, he goes on to suggest that for this reason, "by laying bare the misconceptions that through the use of language often almost unavoidably arise concerning the relations between concepts and by freeing thought from that with which only the means of expression of ordinary language, constituted as they are, saddle it . . . my ideography [*Begriffsschrift*], further developed for these purposes, can become a useful tool for the philosopher."⁶⁵

7. Frege's Expressive Conception of Logic

Frege's early understanding of logic offers some specific content to the notion of explicitly expressing what is implicit in a conceptual content. That is what is required to fill in a notion of expressive or elucidating rationality that might be laid alongside (and perhaps even be discovered to be presupposed by) notions of rationality as accurate representation, as logically valid inference, and as instrumental practical reasoning. Before he makes the fateful step from seeing logic as an attempt to codify inferences to seeing it as the search for a special kind of truth (which Dummett bemoans, and to which we owe much of contemporary logic), Frege's aim is to introduce vocabulary that will let one say (explicitly) what otherwise one can only do (implicitly). Consider the conditional, with which the *Begriffsschrift* begins. Frege says of it: "The precisely defined hypothetical relation between contents of possible judgments has a similar significance for the foundation of my concept-script to that which identity of extensions has for Boolean logic."⁶⁶ Prior to the introduction of such a locution, one could *do* something, one could treat a judgment as having a certain content (implicitly attribute that content to it), by endorsing various inferences involving it and rejecting others. After conditional locutions have been introduced, one can *say*, as part of the content of a claim, that a certain inference is acceptable. One is able to make explicit material inferential relations between an antecedent or premise and a consequent or conclusion. Since according to the inferentialist view of conceptual contents, it is these implicitly recognized material inferential relations that conceptual contents consist in, the conditional permits such contents to be explicitly expressed. If there is a disagreement about the goodness of an inference, it is possible to say what the dispute

is about and to offer reasons one way or the other. The conditional is the paradigm of a locution that permits one to make inferential commitments explicit as the contents of judgments.

The conditional ("the precisely defined hypothetical relation between contents of possible judgments"), rather than inclusion relations among extensions of concepts, plays the central role in Frege's logic because of two cardinal features of his view that distinguish it from the modern set-theoretic interpretations that develop from Boole's approach. First, he understands the content of nonlogical concepts in terms of their inferential role, rather than in terms of their extensions. Second, he understands the task of logical vocabulary to be expressing explicitly what is implicit in those material conceptual contents. What is implicit in those contents, according to the first or inferentialist commitment, is proprieties of inference. Making what follows from what explicit, as itself a *judgeable* content, one that can itself appear as a premise or conclusion in inference, is exactly the job of the conditional.

Frege's overall project for his *Begriffsschrift* is to use conditionals to make it possible to *say* explicitly what the inferential role of ordinary, nonlogical concepts is. Where, as he thinks is often the case in natural language, the content expressed by words is unclear, the project of expressing them explicitly will show where they need or can use clarification. The project is the rectification of concepts: clarifying them by explicating their contents. It is *saying* what their inferential role is: what follows from the applicability of each concept and what its applicability follows from. Employing the explicating logical locutions of which the conditional is the paradigm is to enable what Frege calls "the scientific formation of concepts." Such concepts will wear their contents on their sleeves; the inferential proprieties in virtue of which they mean what they mean are written down for all to read. The particular sciences can then proceed with their reasonings according to the same standards of rigor in the definition and use of their concepts that nineteenth-century mathematics finally came to aspire to. Although the application of this expressive methodology to the special case of mathematics always was closest to Frege's heart and occupied the greatest part of his energies, right from the beginning he had wider expressive ambitions. So the later writings on the sorts of content to be associated with nonmathematical concepts ("On Sense and Reference" prime among them) ought not to be seen to represent any change of interest or detour from his primary project.

Frege is not as explicit about the role of materially correct inferences as Sellars is, but his commitment to the notion is clear from the relation between two of the views that have been extracted from the *Begriffsschrift*: semantic expressivism about logic and inferentialism about content. Expressivism about logic means that Frege treats logical vocabulary as having a distinctive expressive role—namely making explicit the inferences whose goodness is implicit in the conceptual contents of nonlogical concepts. Infer-

entialism about those conceptual contents is taking them to be identified and individuated by their inferential roles. Together these views require that it be coherent to talk about inference prior to the introduction of specifically logical vocabulary, and so prior to the identification of any inferences as good in virtue of their logical form.⁶⁷ In the context of an inferential understanding of conceptual contents, an expressivist approach presupposes a notion of nonlogical inference, the inferences in virtue of which concepts have nonlogical content. So the early Frege envisages a field of material inferences that confer conceptual content on sentences caught up in them. Although Frege does not offer an explanation of the concept, in the *Begriffsschrift* his expressive, explicating project commits him to something playing the role Sellars later picks out by the phrase “material inference.”

There is a sense, then, in which the early Frege does see endorsement of conditional judgments as *implicit* in endorsement of the correctness of inferences. It is implicit in exactly the sense that what one is committed to by endorsing an inference as correct, and so by associating a certain conceptual content (that is inferential role) with the premises and conclusion, can be *made* explicit by expressing it in the form of a conditional judgment. The point of introducing logical vocabulary is precisely to make it possible to trade hitherto merely implicit inferential commitments for explicit assertional commitments to conditionals. And the payoff from expressing explicitly (in the form of judgments) the content-constitutive commitments that were implicit in prior inferential practice is the clarification and rectification of those conceptual contents. Formalism about inference—denying the existence of materially good inferences by assimilating all good inferences to logically good inferences, understanding all proprieties of inference as always already underwritten by logical form—turns things on their head. It misses the point of the process of explicitation that Frege puts at the center of the logical enterprise. It is a form of intellectualism, platonism, or *regulism* in the sense defined in the previous chapter. For it sees *rules* or *principles* as already standing behind every propriety of (in this case inferential) practice.

Frege’s primary interest is in the process of explicitation: of expressing what is implicit in a practice, formulating it as an explicit rule or principle. This pragmatist project of explaining how knowing-that is founded on knowing-how, of explaining the codification in (conditional) principles of (inferential) practice, is unintelligible from any theoretical standpoint that acknowledges only the explicit form of propriety. Frege’s fundamental insight into the expressive role of logical vocabulary (above all the conditional) is not incompatible with claiming that commitment to a conditional is implicit in endorsement of an inference, provided one is careful about what is meant by ‘implicit’—provided, that is, that it is understood as making reference to the possibility of engaging in the substantive activity of making it explicit in the form of a claim or principle. But if one goes on to treat all reasoning as explicitly involving *detachment* from conditionals, and there-

fore implicit endorsement of logical truths involving conditionals (including tautologies involving nested or iterated conditionals), then the line has been crossed and Frege's expressive insight has been lost. As should become clearer from the discussion of Section V below, one of the casualties of the inferential formalist's inversion of the significance of the role of conditionals in making explicit in the form of a principle what is implicit in an inferential practice is a proper understanding of the way in which the contents of inferentially articulated concepts evolve and are clarified as they are expressed with the help of logical locutions.

8. Expressive Completeness and the Two-Valued Conditional

Various special features of Frege's presentation of his conditional, and of the use he goes on to make of it in the *Begriffsschrift*, tend to obscure the crucial expressive role in explicating inferences (and therefore conceptual contents) that he assigns to it. These are picked up and emphasized, to some extent even in his own later work, but especially in the subsequent logistical tradition to which he gave birth, and make it difficult to work back to an appreciation of his original logical project. The difficulties stem from his use of the now-classical two-valued conditional.

It is difficult now to read the definition by which he introduced his conditional (in the fifth paragraph of the *Begriffsschrift*) without being blinded by hindsight—in particular by the glare of the truth-tabular tautology formulation presented by Wittgenstein in the *Tractatus*. Frege does define what has come to be called (by the lights of this work ludicrously inappropriately) the 'material' conditional. He does so, however, not in terms of a semantic distinction between judgeable contents that are true and those that are false, but rather in terms of a pragmatic distinction between those that are affirmed (*bejaht*) and those that are denied (*verneint*). This is his invariable practice in the *Begriffsschrift*, although in later years he is happy enough to recast these claims in terms of truth (as part of the reorientation of his thought toward logical truth that Dummett rightly complains about). Putting things in terms of truth rather than affirmation pushes into the background (though it does not abolish) the way in which the semantic notion of content is beholden to the pragmatic notion of force, in the explanation of which it serves, and which is the source of the priority of judgeable contents and so, even in the later work, of the special central and ineliminable role played by the True as *Bedeutung*. It is worth recalling in this connection Frege's formulation of his view in 1915, already quoted in Section II above: "'True' only makes an abortive attempt to indicate the essence of logic, since what logic is really concerned with is not contained in the word 'true' at all but in the assertoric force with which a sentence is uttered . . . the thing that indicates most clearly the essence of logic is the assertoric force with which a sentence is uttered."⁶⁸

In fact Frege's view is that 'true' is a bit of logical vocabulary, which serves to express explicitly what is done implicitly in asserting. This is why by the time of the *Grundgesetze* (1893) he has adopted a regimentation in which all claims are expressed explicitly in the form of identities that have a sentence on one side and the canonical name 'the True' on the other. It is with that regimentation in mind, in turn, that he claims that the True is an object that must be recognized, at least implicitly, by anyone who makes judgments at all. (Identity claims are explained as the explicit expression of "recognition judgments" in the *Grundlagen*; see 7.1 below.) Assimilating all assertions to assertions of identities permits the use of his (ultimately substitutional) semantics for identity statements (forwarded in "Über Sinn und Bedeutung") in general application to all claims, which is his strategy in the *Grundgesetze*.

The key point is that explicitation is not explanation. Proprieties of inference are not *explained* in terms of something more primitive by being expressed in the explicit form of claims by the use of conditionals; the force of asserting or judging is not *explained* by expressing it explicitly as a saying of a sentence that it is (a name of the) true. This is why Frege always insists that truth is indefinable, something the understanding of which is always already implicit in claiming. "The True" is not a name whose sense one can grasp first, and then appeal to in explaining what it is to make a claim; its use merely makes explicit what is implicit in claiming. It has an expressive, not an explanatory role. Thinking of it the other way around is making a mistake with respect to 'true' and claiming that is strictly analogous to the inferential formalist's mistake regarding the conditional and inferring. A version of the preferred pragmatic direction of explanation is presented below, where Chapter 5 discusses the role of 'true' in terms of the expressive, explicitating function it performs with respect to claiming, according to the account of that practice offered in Chapter 3.

Bracketing subtleties regarding the relation between truth and commitment or affirmation, the fact remains that the conditional Frege actually defines and employs rules out only the case in which the consequent is denied or taken to be false while the antecedent is affirmed or taken to be true. This form of conditional, whatever its compositional virtues, is an extremely impoverished resource for the expression of proprieties of inference. The job that has been attributed to the conditional for Frege is that a conditional be affirmable or taken to be true just in case the inference from its antecedent to its consequent is endorsed or taken to be correct. Using the two-valued conditional to establish the connection between the correctness of an inference and the truth or endorsement of the claims that are its premises and conclusions has unpalatable results. Frege clearly has in mind a fundamental semantic principle regarding this connection: a good inference never takes one from premises that are true to a conclusion that is not true. This is a way of thinking about inferences as commitment-preserving: if one is committed to the premises of a good inference, in the sense of taking them

to be true (the sense that matters for assertion and judgment), then one is committed in the same sense to the conclusion.⁶⁹

Such a principle could be agreed to both by those who adopt the traditional order of semantic explanation—by understanding the principle as explicating the correctness of inference in terms of a prior notion of truth (or taking-true)—and by those who adopt the converse order of semantic explanation (pursued in this work)—by taking truth or the sort of commitment involved in taking-true (and hence in asserting and judgment) to be explicated as what is preserved by good inferences.⁷⁰ But in either case, that truth or commitment is preserved by an inference ought to be taken to be a *necessary* condition of its being a good inference, not a *sufficient* condition. Affirming or taking-true both the claim that Hegel was Hölderlin's roommate and the claim that 43 is prime, and so being committed to the inference from the one claim to the other preserving truth and commitment, does not involve endorsing the propriety of that inference.

The two-valued conditional is subject to this familiar sort of complaint about fallacies of irrelevance precisely because the inferences it codifies explicitly are those that result from implausibly treating the plausible semantic preservation principle as, not only a necessary condition of good inference, but also as a sufficient one. It follows that the two-valued conditional Frege actually defines is an alarmingly bad choice for making explicit actual proprieties of inference. That fact in turn seems to cast doubt on the expressive understanding of his project. If he really wants logical vocabulary to make inferences explicit—because he wants to make nonlogical contents explicit and understands them in inferential terms—why does he employ the blunt, crude tool that is the two-valued conditional, whose expressive powers are hopelessly inadequate for the task of expressing the material inferences that might plausibly be identified with conceptual contents?

The answer is that although he hopes eventually to be able to use logical vocabulary to make explicit the inferential involvements in virtue of which nonlogical claims have the conceptual contents they do, the task Frege actually undertakes in the text of the *Begriffsschrift* is much less ambitious. The only concepts whose inferential role he actually makes explicit there are the logical concepts themselves, and those mathematical concepts that turn out to be definable from them. The concepts of geometry and mechanics—and indeed, the rest of the nonlogical concepts that philosophers might be interested in clarifying by expressing them explicitly—are to be *expressible* by means of logical vocabulary, together with other primitive signs. They are not understood to be, as some of the mathematical concepts (but not, for instance, those of geometry) are, *definable* by means of the logical vocabulary *alone*.

The first stage of Frege's grand project of clarification of nonlogical concepts through their explicitation in logical terms is to make explicit the conceptual contents of the logical expressions that are to be employed in that project. These concepts must themselves be "formed scientifically." This is

why he is proud to display, for each of the official propositions of the *Begriffsschrift* (couched entirely in logical vocabulary), what it follows from (in the proof of the proposition) and what follows from it (in the appendix). Doing so specifies the inferential role of those propositions, and so, indirectly, the conceptual content of the subsentential logical vocabulary that occurs in them.

Thus the only inferences Frege makes explicit in the *Begriffsschrift* are the inferences that are good in virtue of their logical form—for these determine the conceptual content (in his sense) of his logical vocabulary. He finds, in the two-valued conditional, an expressive equilibrium: the inferences in virtue of which that conditional means what it means can themselves be expressed and codified by the use of that conditional. Frege's logical vocabulary is potentially (and he makes it actually) self-explicating.⁷¹ The official propositions of the *Begriffsschrift* explicitly specify the inferential roles of the logical vocabulary, and the inferential roles of those propositions can be expressed explicitly in terms of that vocabulary. Fascinated by how much of mathematical vocabulary turns out to be logical vocabulary in this sense, Frege does not in this work pursue the question of the expressive adequacy of his conditional for material, nonlogical, conceptual contents. He devotes most of the rest of his life to exploring the conceptual contents that can be made explicit by the use of this extensional conditional.⁷²

The results he achieves with the poor expressive resources of the two-valued conditional deserve our awe and admiration. Nevertheless, the motivations remain for the grander semantic expressive aspirations that the young Frege contributes to the inferentialist tradition. It was pointed out above that distinguishing a privileged class of good inferences as good in virtue of their logical form, that is, as logically valid inferences, requires being able to pick out some vocabulary as distinctively *logical* vocabulary. Then the logically valid inferences are just those good inferences that remain good on all substitutions of nonlogical for nonlogical vocabulary. The demarcational question of how logical locutions ought to be identified has received various influential answers. The current suggestion is that Frege's early work is predicated on the idea that what distinguishes vocabulary as specifically *logical* is its expressive role in making conceptual content explicit. Vocabulary deserves the appellation 'logical' just in case it serves to make *explicit*, as the content of a claim, proprieties concerning the use of the expression that otherwise remain *implicit* in practice, specifically the proprieties in virtue of which it has the conceptual content that it does. It is because Frege understands those content-conferring practical proprieties to be in the first instance proprieties of *inferential* practice that the paradigmatic sentential logical locution for him is the conditional. One of the central tasks of the rest of this work is to show how this semantic expressive paradigm can be extended to other logical and semantic locutions.

In the next chapter it is argued that a key link connecting the implicit

norms governing the use of expressions with the conceptual content those practices confer on them is provided by the notion of the *incompatibility* of commitments. In practical terms of normative status, to treat p and q as incompatible claims is to take it that commitment to one precludes entitlement to the other. Practices properly articulated to be interpretable as instituting the normative statuses of commitment and entitlement required for incompatibility relations are thereby interpretable as conferring semantic content on the states, attitudes, and performances that stand in incompatibility relations. The content of a claim can be represented by the set of claims that are incompatible with it. For instance, a relation of entailment, required for an inferential semantics, can be derived according to the principle that p entails q just in case everything incompatible with q is incompatible with p . The formal semantics generated by such incompatibility interpretations is quite rich. It has been shown, for instance, how to represent classical logic, relevance logic, and various systems of orthologic (or quantum logic) by constraints on incompatibility relations.⁷³

Negation, as a logical connective supporting formally valid inferences, plays the same explicating role with respect to material *incompatibility* relations among judgeable (that is propositional) contents that the *conditional* plays with respect to material *inferential* relations among such contents. The formal negation of a claim is constructed as its minimal incompatible, the claim that is entailed by each one of the claims incompatible with the claim of which it is the negation. Thus in the context of a conditional that makes entailment relations explicit, the introduction of a locution playing the inferential role of negation makes it possible to make explicit the relation of material incompatibility between claims. To assert that p is incompatible with q , one asserts the conditional whose antecedent is p and whose consequent is the negation of q . Conjunction and disjunction can be handled straightforwardly as corresponding to Boolean operations on the sets of incompatibles that represent conceptual contents according to this sort of semantic model.

Chapter 7 below discusses Frege's treatment of identity locutions as making explicit the substitution-inferential commitments that are implicit in the use of singular terms. It also shows how that idea can be extended to an account of the use of quantifiers as making explicit the different sort of substitution-inferential commitment that is implicit in the use of predicates. The job of the next chapter is to offer an account of the normative practices of claiming and judging, and of the propositional contents conferred on states, attitudes, performances, and expressions by their playing appropriate roles in those practices. This account gives a definite sense to the notion of explicit *sayings*, in terms of norms implicit in *doings*. What is explicit is then the propositional content that is said or believed. In this fundamental sense, H. L. Mencken makes the content of his thought explicit, and expresses it fully, by asserting the declarative sentence: "Natives of Appalachia are clay-eating

sub-humans." But it is also possible to use logical vocabulary to make explicit expressively essential inferential involvements that remain implicit in the concepts employed in making this claim.

In subsequent chapters various other locutions are introduced as being used so as to make explicit, in this sense, some feature of the practices that originally confer propositional content (so that having such contents can be understood to consist in how it is correct for those locutions to be used, according to the practices in question). Not only the standard logical vocabulary, but also traditional semantic vocabulary such as 'true', 'refers', and the 'of' of intentional aboutness, should be understood as semantically explicating. The point of using these sorts of expressions is to make explicit as the contents of claims (whose consequences can be explored and which can be justified and disputed) some critical element of the practices of talking and believing in virtue of which it is possible to interpret anything as propositionally contentful in the first place. Furthermore, another range of expressions, including such locutions as 'claims that', 'believes that', 'intends that', and normative talk of commitments and entitlements, is interpreted as *pragmatically* explicating. The point of using these sorts of expressions is to make explicit as the contents of claims some of the pragmatic elements of the practices of talking, believing, and acting that confer propositional contents. One thread running through the later chapters of this work is the attempt to achieve an analog of the expressive equilibrium Frege achieves in the propositional fragment of the *Begriffsschrift*. The challenge is to show how not only the semantics, but the pragmatics outlined in the first four chapters can be made explicit, in terms of vocabulary that is introduced by specifying practices of using it that are sufficient to confer on it the content that is then employed in making explicit precisely those practices and that content. The ideal is that the theory should specify practices sufficient to confer on the various locutions considered all the kinds of content required to state the theory itself.

V. CIRCUMSTANCES AND CONSEQUENCES OF APPLICATION

1. Dummett's Model

The previous section of this chapter introduced three themes: that conceptual content is to be understood in terms of role in reasoning rather than exclusively in terms of representation, that the capacity for such reasoning is not to be identified exclusively with mastery of a logical calculus, and that besides theoretical and practical reasoning using contents constituted by their role in material inferences, there is a kind of *expressive* rationality that consists in making implicit, content-conferring inferential commitments explicit as the contents of assertional commitments. Being

rational in the primary sense is having states and attitudes and producing performances that have propositional contents. The next chapter discusses how the inferential articulation essential to such contents is conferred on them by the way in which the states, attitudes, and performances exhibiting those contents are caught up in the game of giving and asking for reasons. Rationality consists in mastery of those practices. It is not to be understood as a logical capacity. Rather, specifically logical capacities presuppose and are built upon underlying rational capacities. The fundamental characteristic role of logical vocabulary is to make it possible to talk and think *explicitly* about the inferentially articulated semantic contents *implicitly* conferred on expressions (among other things) by their role in rational practice. The optional introduction of sophisticated logical explicating vocabulary has an expressive point and payoff. By its means the material inferential practices, which govern and make possible the game of giving and asking for reasons, are brought into that game (and so into consciousness) as explicit topics of discussion and justification. In this way, in the context of the three basic themes mentioned above, an *expressive* understanding of logic was introduced—according to which formal validity of inferences is substitutionally defined in terms of material correctness of inferences together with the discrimination of some privileged vocabulary; that privileged vocabulary is identified as logical vocabulary; and what it is for something to be a bit of logical vocabulary is explained in terms of its semantically expressive role.

These ideas, to be found in the early works of Frege and Sellars, provide the beginnings of the structure within which modern inferentialism develops. The approach they suggest can be made more definite by considering a general model of conceptual contents as inferential roles that has been recommended (in somewhat different terms) by Dummett. According to that model, the use of any linguistic expression or concept has two aspects: the circumstances under which it is correctly applied, uttered, or used, and the appropriate consequences of its application, utterance, or use. Though Dummett does not put the point this way, this model connects to inferentialism of the Sellarsian sort via the principle that part of the content to which one is committed by using the concept or expression may be represented by the material inference one implicitly endorses by such use: the inference from the circumstances of appropriate employment to the appropriate consequences of such employment.

The original source for the model lies in a treatment of the grammatical category of sentential connectives. Dummett's two-aspect model is a generalization of a standard way of specifying the inferential roles of logical connectives, due ultimately to Gentzen. Gentzen defined connectives by specifying *introduction rules*, or inferentially sufficient conditions for the employment of the connective, and *elimination rules*, or inferentially necessary consequences of the employment of the connective. So, to define the

inferential role of an expression '&' of Boolean conjunction, one specifies that anyone who is committed to p , and committed to q , is thereby to count also as committed to $p \& q$, and that anyone who is committed to $p \& q$ is thereby committed both to p and to q . The first schema specifies, by means of expressions that do not contain the connective, the *circumstances* under which one is committed to claims expressed by sentences that contain (as principal connective)⁷⁴ the connective whose inferential role is being defined, that is, the sets of premises that entail them. The second schema specifies, by means of expressions that do not contain the connective, the *consequences* of being committed to claims expressed by sentences that contain (as principal connective) the connective whose inferential role is being defined, that is, the sets of consequences that they entail.

Dummett makes a remarkable contribution to inferentialist approaches to conceptual content by showing how this model can be generalized from the case of logical connectives to provide a uniform treatment of the meanings of expressions of other important grammatical categories: sentences, predicates and common nouns, and singular terms. The application to the *propositional* contents expressed by whole declarative sentences is straightforward. What corresponds to an introduction rule for a propositional content is the set of inferentially sufficient conditions for asserting it, and what corresponds to an elimination rule is the set of inferentially necessary consequences of asserting it, that is, what follows from doing so. Dummett says: "Learning to use a statement of a given form involves, then, learning two things: the conditions under which one is justified in making the statement; and what constitutes acceptance of it, i.e., the consequences of accepting it."⁷⁵ Dummett presents his model as specifying two fundamental features of the use of linguistic expressions. In what follows, it is applied in the context of the previous ideas, to bring into relief the implicit material inferential content a concept or expression acquires in virtue of being used in the ways specified by these two "aspects." The link between *pragmatic significance* and *inferential content* is supplied by the fact that asserting a sentence is (among other things) implicitly undertaking a commitment to the correctness of the material inference from its circumstances to its consequences of application.

Dummett applies his model exclusively to conceptually contentful *linguistic* expressions. But it is clear that the model has a wider application—to intentional states and attitudes generally. Dummett's model just provides some structure to the representation of the functional roles of intentional states. For instance, one could think about the functional role played by a belief with a particular propositional content in terms of the circumstances in which it is appropriate to acquire a belief with that content, and the appropriate consequences of such acquisition. In the next two chapters this basic Dummettian structure is further articulated, to allow a richer representation of the functional roles of states, performances, and expressions.

2. *Inferential Connection between Even Noninferential Circumstances and Consequences of Application*

The concepts least easily assimilated to an inferential model are the empirical concepts whose core employment is in perception and the formulation of observation reports. For such reports are essentially *noninferential*—in the sense that they are elicited as responses to features of the largely nonlinguistic environment, rather than as conclusions drawn from other claims. Their content accordingly derives (at least in large part) from the *reliable differential responsive dispositions* that those who have mastered the concepts exhibit with respect to their application. Such concepts can be assimilated to the inferentialist understanding of conceptual contents by adapting Dummett's idea of distinguishing two crucial features of the use of linguistic expressions: their *circumstances* of appropriate application, and the appropriate *consequences* of such application. In terms of this model, it is possible to understand the use of *any* expression as implicitly involving an inferential commitment. In particular, by using the expression, one is (among other things) committed to the propriety of the inference from its circumstances to its consequences of application. The consequences of application are always themselves inferentially related to the concept in question (although the inference involved may include practical inferences, whose conclusions are commitments to act). The circumstances of application need not themselves be linguistic. For the concept *red*, for instance, they include the presence of visibly red things. Nonetheless, even the use of concepts of this sort can be seen to embody inferential commitments, to the propriety of applying inferential consequences of *red*—for instance, *colored*—to anything that *red* is properly applied to.⁷⁶

It is in this way that a broadly inferential approach can incorporate into its conception of the contents of empirical concepts the nonlinguistic circumstances in which they are correctly noninferentially applied. Thus the concepts *water* and its twin-earth analog *twater*, which are by hypothesis alike except that one is appropriately applied to H₂O and the other to XYZ, count as involving different *inferential* contents. This is so even though they fund inferential moves involving the same sorts of *noises*, from saying "That's water" to "That's liquid," for instance. For they involve different circumstances of appropriate application, and hence different inferential transitions from those circumstances to their consequences. So even though it is the practices of those whose concepts they are that confer on them their contents, the earthlings and twin-earthlings need not be able to tell that they have different concepts, if water and *twater* are indistinguishable to them. They are not omniscient about the inferential commitments implicit in their own concepts. For the interpreter who is making sense of their practices, and who is able (not perceptually, but conceptually) to distinguish H₂O and XYZ, can understand transported earthlings as *mistaking* the XYZ they look at for

water, as *inappropriately* applying the concept they express with their word 'water' to that unearthly stuff. The circumstances of appropriate noninferential application of the concept expressed by the English word 'water' require that it be applied in response to a sample of H_2O .

In this way the circumstances of appropriate application of a claim can include not only other claims (from which the one in question could be inferred) but also perceptual circumstances (to which one has been trained to respond *noninferentially* by endorsing the target claim). The appropriate consequences of application of a claim can include not only the inferential acquisition of further beliefs whose contents follow from the contents of the belief in question but also, in the context of further contentful intentional states, the noninferential responsive performance of *actions*, under the descriptions by which they can be exhibited as the conclusions of practical inferences. This is explicitly acknowledged in the continuation of the passage from Dummett quoted above: "Here 'consequences' must be taken to include both the inferential powers of the statement and anything that counts as acting on the truth of the statement." So the circumstances and consequences of application Dummett is talking about should not be identified with inferentially necessary and sufficient conditions, where this means identifying them with sets of claims or beliefs that are conclusions or premises of theoretical inferences involving the content in question. It turns out, however, that the circumstances and consequences model can be understood as an inferential model, regardless of whether the circumstances and consequences are themselves already thought of in inferential terms. The inferential element enters this picture in the commitment undertaken by one who employs a given content to the propriety of the transition from the circumstances of appropriate application to the appropriate consequences of application of a conceptual content. This will be construed as a broadly inferential commitment, though the detailed justification for this characterization will not emerge until the next chapter.⁷⁷

One advantage of thinking about conceptual content as determining functional role specified in terms of proprieties governing circumstances and consequences of application is the room it makes for a *pragmatic* picture of understanding or grasping such a content. Understanding can be understood, not as the turning on of a Cartesian light, but as practical mastery of a certain kind of inferentially articulated *doing*: responding differentially according to the circumstances of proper application of a concept, and distinguishing the proper inferential consequences of such application. This is not an all-or-none affair; metallurgists understand the concept *tellurium* better than most of us do, for training has made them master of the inferential intricacies of its employment in a way that we can only crudely approximate. On this inferentialist rendering, thinking clearly is a matter of knowing what one is committing oneself to by a certain claim, and what would entitle one to that commitment. Writing clearly is providing enough clues for a reader to infer

what one intends to be committed to by each claim, and what one takes it would entitle one to that commitment. Failure to grasp either of these components is failure to grasp the inferential commitment that use of the concept involves, and so failure to grasp its conceptual content.

3. *One-Sided Theories of Meaning*

Verificationists, assertibilists, and reliabilists make the mistake of treating the first aspect as exhausting content. Understanding or grasping a content is taken to consist in practically mastering the circumstances under which one becomes entitled or committed to endorse a claim or acquire a belief, quite apart from any grasp of what one becomes entitled or committed to by such endorsement or acquisition. But claims can have the same circumstances of application and different consequences of application, as for instance "I foresee (or predict) that I will write a book about Hegel" and "I will write a book about Hegel" do. Any circumstances under which one is entitled to one of these claims (or to acquire the belief it expresses) are circumstances under which one is entitled to the other. (If this does not seem right for the actual concepts expressed by 'foresee' and 'predict', artificial variants clearly can be constructed for which it is.) Yet what follows from the claims is quite different. If I will write a book about Hegel, then I will write a book about Hegel. Yet that I will write a book about Hegel does not follow from my foreseeing or predicting that I will (as the sad history of orphaned first volumes of ambitious projects attests). The consequences of these claims are quite different. Examples meeting the conditions required for this point are forthcoming in any idiom expressively rich enough to contain *pragmatically* explicitating locutions, which permit one to *say* what one is *doing* in performing a certain speech act or acquiring a certain state or attitude—for instance "I claim that *p*" or "I believe that *p*." (These locutions are discussed in Chapter 8.)

In any idiom expressively rich enough to contain *semantically* explicitating locutions, whose paradigm is the conditional, the difference in inferential consequences of application between the sentence whose utterance performs a speech act (a doing in which the force is left implicit) and the sentence whose utterance explicitly *says* that that is what one is doing (so that force becomes part of the content) itself becomes explicit in the use of conditionals with those sentences as antecedents. The circumstances of appropriate application or assertibility conditions of the conditionals "If I will write a book about Hegel, then I will write a book about Hegel" and "If I foresee (or predict) that I will write a book about Hegel, then I will write a book about Hegel" are quite different. The assertibility of the second conditional, but not the first, depends on auxiliary hypotheses about how good at foreseeing or predicting I am. So for idioms that contain both the pragmatically explicitating locutions that permit the construction of pairs of sentences with identical

circumstances of application and different consequences of application, and the semantically explicating locutions that permit the construction of conditionals whose circumstances of application differ depending on the consequences of application of their antecedents, it is possible to show the inadequacy of a semantics that avails itself only of assertibility conditions or circumstances of appropriate application. For examples of the sort just considered show that substituting another sentence with the same assertibility conditions for a sentence that is the antecedent of a conditional can alter the assertibility conditions of the compound. In such an expressively rich environment, then, assertibility conditions cannot provide an adequate model of what Dummett calls "*ingredient content*," the contribution the occurrence of a sentence makes to the use of sentences in which it appears as a component. But this fact simply reflects the inadequacy of the model for the expression of conceptual content as inferential role, even in the more expressively impoverished idioms in which the pragmatically and semantically explicating locutions are not available.⁷⁸

The inadequacy of a notion of semantic content that is restricted to circumstances of application to the exclusion of consequences of application has already appeared in another guise above. The point of the discussion of Sellars's application of inferentialist ideas to the understanding of noninferential reports, in Section III, was that parrots and photocells and so on might reliably discriminate the circumstances in which the concept *red* should be applied, without thereby grasping that concept. This would happen precisely in the case where they have no mastery of the consequences of such application—when they cannot tell that it follows from something being red that it is colored, that it is not a prime number, and so on. You do not convey to me the content of the concept *gleeb* by supplying me with an infallible gleebsness tester that lights up when and only when exposed to gleeb things. I would in that case know what things were gleeb, without knowing what I was saying about them when I called them that, what I had found out about them or committed myself to. Dummett offers two examples of philosophically important concepts where it is useful to be reminded of this point: "An account, however accurate, of the conditions under which some predicate is rightly applied may thus miss important intuitive features of its meaning; in particular, it may leave out what we take to be the point of our use of the predicate. A philosophical account of the notion of truth can thus not necessarily be attained by a definition of the predicate 'true', even if one is possible, since such a definition may be correct only in the sense that it specifies correctly the application of the predicate, while leaving the connections between this predicate and other notions quite obscure."⁷⁹ Even more clearly: "A good example would be the word 'valid' as applied to various forms of argument. We might reckon the syntactic characterization of validity as giving the criterion for applying the predicate 'valid' to an argument, and the semantic characterization of validity as giving the consequences of such an application . . . If [a student] is taught in a very unimaginative way,

he may see the classification of arguments into valid and invalid ones as resembling the classification of poems into sonnets and non-sonnets, and so fail to grasp that the fact that an argument is valid provides any grounds for accepting the conclusion if one accepts the premises. We should naturally say that he had missed the point of the distinction."

Pragmatists of the classical sort, in contrast, make the converse mistake of identifying propositional contents exclusively with the consequences of endorsing a claim: looking downstream to the claim's role as a premise in practical reasoning and ignoring its proper antecedents upstream. The fact that the pragmatist's emphasis is on practical consequences is not relevant to this complaint. The problem is that one can know what follows from the claim that someone is responsible for a particular action, that an action is immoral or sinful, that a remark is true or in bad taste, without for that reason counting as understanding the claims involved, if one has no idea when it is appropriate to make those claims or apply those concepts. Being classified as AWOL does have the consequence that one is liable to be arrested, but the specific circumstances under which one acquires that liability are equally essential to the concept.

It was pointed out that Frege's practice in the *Begriffsschrift* is to specify both the circumstances and the consequences of application of his claims, which in the context of that project (excluding as it does concepts with empirical and practical content deriving from their relation to perception and action) can be identified with the inferentially sufficient premises from which they follow and the inferentially necessary conclusions they lead to. Yet his official definition of conceptual content refers only to consequences, and Carnap follows him in this regard. For the special sort of concepts they are concerned with, where only inferential circumstances and consequences are in play, this restriction does not amount to ignoring circumstances of application. Restricting consideration for the sake of an example to one-premise inferences, associating with each sentence the set of sentences that follow from it determines for each sentence which sentences it follows from. So at the global level, nothing is lost by officially defining content in terms of inferential consequences alone. As will emerge below, it is quite otherwise when one is concerned locally with the content associated with each sentence—for instance in asking what it is to understand the content expressed by one sentence (but perhaps not another), or to alter the content expressed by one sentence, or to introduce a new content. Nor will the technical dodge of restriction of content to consequences be adequate when attention is turned to the sort of empirical and practical content concepts get from their involvement in perception and action.

4. *Conservativeness and the Coherence of Logical Concepts*

Of course, such one-component theories do not simply ignore the aspects of content they do not treat as central. Dummett says:

Most philosophical observations about meaning embody a claim to perceive . . . a simple pattern: the meaning of a sentence consists in the conditions for its truth and falsity, or in the method of its verification, or in the practical consequences of accepting it. Such dicta cannot be taken to be so naive as to involve overlooking the fact that there are many other features of the use of a sentence than the one singled out as being that in which its meaning consists: rather, the hope is that we shall be able to give an account of the connection that exists between the different aspects of meaning. One particular aspect will be taken as central, as constitutive of the meaning of any given sentence . . . ; all other features of the use of the sentence will then be explained by a uniform account of their derivation from that feature taken as central.⁸⁰

Pursuing this notion of derivation provides a helpful perspective on the idea of conceptual contents articulated according to material inferences, and on the role of explicit inference licenses such as conditional statements in expressing and elucidating such inferences and so such contents.

The strategy of attempting to *derive* one aspect of the use of an expression (or the significance of an intentional state) from another—in particular to derive appropriate consequences of application from circumstances of appropriate application, or vice versa—expresses Dummett's appreciation of the need for the semantic theorist to be able to explain two crucial features of our practices regarding conceptual contents. Concept-users are often confronted with decisions regarding *alternative* concepts and so are obliged to decide not only that certain uses of a given concept should be rejected as incorrect but also that certain concepts should themselves be rejected as inadequate or incorrect. We criticize our concepts and sometimes reject them. Furthermore, doing so is not simply a matter of free or arbitrary stipulation. Criticism of our concepts is constrained and sometimes compelled. These are important phenomena—an attempt to take proper account of them guides the discussion below. Dummett acknowledges them as motivating the theoretical acknowledgment of a need for *harmony* between the circumstances and consequences of application: "A naive view of language regards assertibility-conditions for a statement as exhausting its meaning: the result is to make it impossible to see how a meaning can ever be criticized, revised, or rejected; it was just such a naive view which led to the use of the notorious 'paradigm-case argument'. An almost equally naive view is that which distinguishes the assertibility-conditions of a statement as its 'descriptive meaning' and its consequences as its 'evaluative meaning', dispensing with any requirement of harmony between them, but holding that we have the right to attach whatever evaluative meaning we choose to a form of statement irrespective of its descriptive meaning."⁸¹

For the special case of defining the inferential roles of logical connectives by pairs of sets of rules for their introduction and for their elimination, which

motivates Dummett's broader model, there is a special condition it is appropriate to impose on the relation between the two sorts of rules. "In the case of a logical constant, we may regard the introduction rules governing it as giving conditions for the assertion of a statement of which it is the main operator, and the elimination rules as giving the consequences of such a statement: the demand for harmony between them is then expressible as the requirement that the addition of the constant to a language produces a conservative extension of that language."⁸² Recognition of the appropriateness of such a requirement arises from consideration of connectives with "inconsistent" contents. As Prior pointed out, if a connective, which after Belnap may be called 'tonk', is defined as having the introduction rule proper to disjunction and the elimination rule proper to conjunction, then the first rule licenses the transition from p to $p \text{ tonk } q$, for arbitrary q , and the second licenses the transition from $p \text{ tonk } q$ to q . The result is what he called a "runabout inference ticket," which permits any arbitrary inference.

Prior thought that this possibility shows the bankruptcy of Gentzen-style definitions of inferential roles. Belnap shows rather that when logical vocabulary is being introduced, one must constrain such definitions by the condition that the rule not license any inferences involving only old vocabulary that were not already licensed before the logical vocabulary was introduced.⁸³ That is, it must be ensured that the new rules provide an *inferentially conservative extension* of the original field of inferences. From the point of view of the joint commitments to understanding conceptual content in terms of material inference and conceiving the distinctive role of logical vocabulary as making those content-conferring inferential connections explicit in the form of claims, this constraint on the definition of logical particles by introduction and elimination rules makes perfect sense. For if those rules are not inferentially conservative, the introduction of the new vocabulary licenses new material inferences and so alters the contents associated with the old vocabulary. The expressive approach to logic motivates a criterion of adequacy for introducing logical vocabulary to the effect that no *new* inferences involving only the *old* vocabulary be made appropriate thereby. Only in this way can logical vocabulary play the expressive role of making explicit the original material inferences and so nonlogical conceptual contents.

5. *Nonlogical Concepts Can Incorporate Materially Bad Inferences*

The problem of what Dummett calls a lack of "harmony" between the circumstances and the consequences of application of a concept can arise, however, not only for logical vocabulary but also for concepts with material contents. Seeing how it does provides further help in understanding the notion of expressive rationality and the way in which the explicating role of logical vocabulary contributes to the clarification of concepts. For conceptual change can be:

motivated by the desire to attain or preserve a harmony between the two aspects of an expression's meaning. A simple case would be that of a pejorative term, e.g. 'Boche'. The condition for applying the term to someone is that he is of German nationality; the consequences of its application are that he is barbarous and more prone to cruelty than other Europeans. We should envisage the connections in both directions as sufficiently tight as to be involved in the very meaning of the word: neither could be severed without altering its meaning. Someone who rejects the word does so because he does not want to permit a transition from the grounds for applying the term to the consequences of doing so. The addition of the term 'Boche' to a language which did not previously contain it would produce a non-conservative extension, i.e. one in which certain other statements which did not contain the term were inferable from other statements not containing it which were not previously inferable.⁸⁴

This crucial passage makes a number of points that are worth untangling.

First of all, it shows how concepts can be criticized on the basis of substantive beliefs. If one does not believe that the inference from German nationality to cruelty is a good one, one must eschew the concept *Boche*. For one cannot deny that there are any Boche—that is just denying that anyone is German, which is patently false. One cannot admit that there are Boche and deny that they are cruel—that is just attempting to take back with one claim what one has committed oneself to with another. One can only refuse to employ the concept, on the grounds that it embodies an inference one does not endorse. (When the prosecutor at Oscar Wilde's trial asked him to say under oath whether a particular passage in one of his works did or did not constitute blasphemy, Wilde replied, "Blasphemy is not one of my words."⁸⁵) Highly charged words like 'nigger', 'whore', 'Republican', and 'Christian' have seemed a special case to some because they couple "descriptive" circumstances of application to "evaluative" consequences. But this is not the only sort of expression embodying inferences that requires close scrutiny. The use of *any* concept or expression involves commitment to an inference from its grounds to its consequences of application. Critical thinkers, or merely fastidious ones, must examine their idioms to be sure that they are prepared to endorse and so defend the appropriateness of the material inferential transitions implicit in the concepts they employ. In Reason's fight against thought debased by prejudice and propaganda, the first rule is that material inferential commitments that are potentially controversial should be made explicit as claims, exposing them both as vulnerable to reasoned challenge and as in need of reasoned defense.

It is in this process that formal logical vocabulary such as the conditional plays its explicating role. It permits the formulation, as explicit claims, of the inferential commitments that otherwise remain implicit and unexam-

ined in the contents of material concepts. Logical locutions make it possible to display the relevant grounds and consequences and to assert their inferential relation. Formulating as an explicit claim the inferential commitment implicit in the content brings it out into the open as liable to challenges and demands for justification, just as with any assertion. In this way explicit expression plays an elucidating role, functioning to groom and improve our inferential commitments and so our conceptual contents—a role, in short, in the practices of reflective rationality that Sellars talks about under the heading of “Socratic method.”

But if Dummett is suggesting that what is wrong with the concept *Boche* is that its addition represents a nonconservative extension of the rest of the language, he is mistaken. Its nonconservativeness just shows that it has a substantive content, in that it implicitly involves a material inference that is not already implicit in the contents of other concepts being employed. This is no bad thing. Conceptual progress in science often consists in introducing just such novel contents. The concept *temperature* was introduced with certain criteria or circumstances of appropriate application and with certain consequences of application. As new ways of measuring temperature are introduced, and new consequences of temperature measurements adopted, the complex inferential commitment that determines the significance of using the concept of temperature evolves.

The proper question to ask in evaluating the introduction and evolution of a concept is not whether the inference embodied is one that is already endorsed (so that no new content is really involved) but whether that inference is one that ought to be endorsed. The problem with ‘Boche’ or ‘nigger’ is not that once we explicitly confront the material inferential commitment that gives them their content, it turns out to be *novel*, but that it can then be seen to be indefensible and inappropriate. We want to be aware of the inferential commitments our concepts involve, to be able to make them explicit, and to be able to justify them. But there are other ways of justifying them than showing that we were already implicitly committed to them, before introducing or altering the concept in question. Making implicit commitments explicit is only a necessary condition of justifying them.

Even in the cases where it does make sense to identify harmony of circumstances and consequences with inferential conservativeness, the attribution of conservativeness is always relative to a background set of material inferential practices, the ones that are conservatively extended by the vocabulary in question. Conservativeness is a property of the conceptual content only in the context of other contents, not something it has by itself. There can be pairs of logical connectives, either of which is all right by itself, but both of which cannot be included in a consistent system. It is a peculiar ideal of harmony that would be realized by a system of conceptual contents such that the material inferences implicit in every subset of concepts represented a conservative extension of the remaining concepts, in that no infer-

ences involving only the remaining ones are licensed that are not licensed already by the contents associated just with those remaining concepts. Such a system is an idealization because all of its concepts would already be out in the open, with none remaining hidden, to be revealed only by drawing conclusions from premises that have never been conjoined before, following out unexplored lines of reasoning, drawing consequences one was not previously aware one would be entitled or committed to by some set of premises. In short, this would be a case where Socratic reflection—making implicit commitments explicit and examining their consequences and possible justifications—would never motivate one to alter contents or commitments. Such complete transparency of commitment and entitlement is in some sense an ideal projected by the sort of Socratic practice that finds current contents and commitments wanting by confronting them with each other, pointing out inferential features of each of which we were unaware. But as Wittgenstein teaches in general, it should not be assumed that our scheme is like this, or depends upon an underlying set of contents like this, just because we are obliged to remove any particular ways in which we discover it to fall short.

These are reasons to part company with the suggestion, forwarded in the passage above, that inferential conservatism is a *necessary* condition of a “harmonious” concept—one that won’t “tonk up” a conceptual scheme. In a footnote, Dummett explicitly denies that conservativeness can in general be treated as a *sufficient* condition of harmony: “This is not to say that the character of the harmony demanded is always easy to explain, or that it can always be accounted for in terms of the notion of a conservative extension . . . The most difficult case is probably the vexed problem of personal identity.”⁸⁶ In another place, this remark about personal identity is laid out in more detail:

We have reasonably sharp criteria which we apply in ordinary cases for deciding questions of personal identity: and there are also fairly clear consequences attaching to the settlement of such a question one way or the other, namely those relating to ascriptions of responsibility, both moral and legal, to the rights and obligations which a person has . . . What is much harder is to give an account of the connection between the criteria for the truth of a statement of personal identity and the consequences of accepting it. We can easily imagine people who use different criteria from ours . . . Precisely what would make the criteria they used criteria for *personal identity* would lie in their attaching the same consequence, in regard to responsibility, motivation, etc., to their statements of personal identity as we do to ours. If there existed a clear method for deriving, as it were, the consequences of a statement from the criteria for its truth, then the difference between such people and ourselves would have the character of a factual disagreement, and one

side would be able to show the other to be wrong. If there were no connection between truth-grounds and consequences, then the disagreement between us would lie merely in a preference for different concepts, and there would be no right or wrong in the matter at all.⁸⁷

Dummett thinks that there is a general problem concerning the way in which the circumstances and consequences of application of expressions or concepts ought to fit together. Some sort of "harmony" seems to be required between these two aspects of the use. The puzzling thing, he seems to be saying, is that the harmony required cannot happily be assimilated either to compulsion by facts or to the dictates of freely chosen meanings. But the options—matter of fact or relation of ideas, expression of commitment as belief or expression of commitment as meaning—are not ones that readers of "Two Dogmas of Empiricism" and its heirs ought to be tempted to treat as exhaustive.⁸⁸

As already pointed out, talk of derivability is strictly stronger than talk of conservativeness. On the other side of the divide, the notion of a completely factual issue that Dummett appeals to in this passage is one in which the applicability of a concept is settled straightforwardly by the application of other concepts: the concepts that specify the necessary and sufficient conditions that determine the truth conditions of claims involving the original concept.⁸⁹ This conception, envisaged by a model of conceptual content as necessary and sufficient conditions, seems to require a conceptual scheme that is ideally transparent in the way mentioned above, in that it is immune to Socratic criticism. For that conception insists that these *coincide*—in that the individually sufficient conditions *already* entail the jointly necessary ones. Only then is it attractive to talk about content as truth conditions, rather than focusing on the substantive inferential commitments that relate the sufficient to the distinct necessary conditions, as recommended here. By contrast to the either/or that Dummett presents, in a picture according to which conceptual contents are conferred by being caught up in a social practical structure of inferentially articulated commitments and entitlements, material inferential commitments are a necessary part of any package of practices that includes material assertional or doxastic commitments. From this point of view, rendering conceptual content as truth conditions and thinking of them as necessary and sufficient conditions leaves out precisely the material content of concepts.

For the circumstances and consequences of application of a concept may stand in a substantive material-inferential relation. To ask what sort of "harmony" they should exhibit is to ask what material inferences we ought to endorse, and so what conceptual contents we ought to employ. This is not the sort of a question to which one ought to expect or even welcome a general or wholesale answer. Grooming our concepts and material inferential commitments in the light of our assertional commitments (including those we

find ourselves with noninferentially through observation) and the latter in the light of the former is a messy, retail business. Dummett thinks that a theory of meaning should take the form of an account of the nature of the "harmony" that ought to obtain between the circumstances and the consequences of application of the concepts we ought to employ. The present point is that one should not expect a theory of that sort to take the form of a specification of necessary and sufficient conditions for the circumstances and consequences of application of a concept to be harmonious. Rather, insofar as the idea of such a theory makes sense at all, it must take the form of an investigation of the ongoing elucidative process, of the "Socratic method" of discovering and repairing discordant concepts, which alone gives the notion of harmony any content. It is given content only by the process of *harmonizing* commitments, from which it is abstracted.

In Sellars's characterization of induction, introduced above, inductive inference is assigned an expressive role insofar as its conclusion is understood as being an inference license making explicit a commitment that is implicit in the use of conceptual contents antecedently in play. Rules of this sort assert an authority over future practice and for their entitlement answer both to the prior practice being codified and to concomitant inferential and doxastic commitments. In this way they may be likened to the principles formulated by judges at common law, intended both to codify prior practice, as represented by precedent, expressing explicitly as a rule what was implicit therein, and to have regulative authority for subsequent practice. The expressive task of making material inferential commitments explicit plays an essential role in the reflectively rational Socratic practice of harmonizing our commitments. For a commitment to become explicit is for it to be thrown into the game of giving and asking for reasons as something whose justification, in terms of other commitments and entitlements, is liable to question. Any theory of the sort of inferential *harmony* of commitments we are aiming at by engaging in this reflective, rational process must derive its credentials from its expressive adequacy to that practice, before it should be accorded any authority over it.

6. Varieties of Inferentialism

Section IV of this chapter introduced three related ideas:

1. the inferential understanding of conceptual content,
2. the idea of materially good inferences, and
3. the idea of expressive rationality.

These contrast, respectively, with

- 1'. an understanding of content exclusively according to the model of the representation of states of affairs,

- 2'. an understanding of the goodness of inference exclusively on the model of formal validity, and
- 3'. an understanding of rationality exclusively on the model of instrumental or means-end reasoning.

In this section these ideas were considered in relation to the representation of inferential role suggested by Dummett, in terms of the circumstances of appropriate application of an expression or concept and the appropriate consequences of such application. Both sections sought to introduce an *expressive* view of the characteristic role played by logical vocabulary and to indicate its relation to the practices constitutive of rationality.

One of the important benefits afforded by the emphasis in this section on understanding the Dummettian model of the use of linguistic expressions in terms of appropriate circumstances and consequences of application linked by an inferential commitment is the clarification it offers concerning the options that are open in working out an inferentialist approach to semantics. There are three different ways in which one might take inference to be of particular significance for understanding conceptual content. The weak inferentialist thesis is that inferential articulation is *necessary* for specifically *conceptual* contentfulness. The strong inferentialist thesis is that broadly inferential articulation is *sufficient* for specifically conceptual contentfulness—that is, that there is nothing more to conceptual content than its broadly inferential articulation. Dummett's model is particularly helpful for focusing attention on how important the qualification 'broadly' is in this formulation. For strong inferentialism as it is worked out in the rest of this project is *not* committed to the hyperinferentialist thesis, which maintains that *narrowly* inferential articulation is sufficient for conceptual contentfulness of all sorts.

The difference between the broad and the narrow conception of inferential articulation has three dimensions. First, and most important, the broad conception includes the possibility of *noninferential* circumstances and consequences of application. In this way (discussed in Chapter 4) the specifically *empirical* conceptual content that concepts exhibit in virtue of their connection to language entries in *perception* and the specifically *practical* conceptual content that concepts exhibit in virtue of their connection to language exits in *action* are incorporated into the inferentialist picture. The use of concepts with contents of these sorts can still be understood in terms of the material inferential commitment one who uses them undertakes: the commitment to the propriety or correctness of the inference from their circumstances to their consequences of application. Conceiving such inferences broadly means conceiving them as involving those circumstances and consequences, as well as the connection between them. The hyperinferentialist about conceptual content (adopting a position not endorsed here) would allow only *inferential* circumstances and consequences of application. Under

such a restriction, it is impossible to reconstruct the contents of actual concepts, except perhaps in some regions of mathematics.

Second, relations of *incompatibility* among claims and (so) concepts are considered broadly inferential relations, on grounds of their antecedents and their consequences. On the side of consequences, incompatibility relations underwrite the modal inferences codified by strict implication. For p entails q in this sense just in case everything incompatible with q is incompatible with p . So being a square entails being a rectangle, because everything incompatible with being a rectangle is incompatible with being a square. On the side of antecedents, the semantic relation of incompatibility will be understood (in the next chapter) in terms of the very same normative statuses of doxastic commitment and entitlement to such commitments, in terms of which inferences are construed (with commitment-preserving inferences corresponding roughly to deductive inferences, and entitlement-preserving inferences corresponding roughly to inductive inferences).

Finally, the notion of broadly inferential articulation is extended in subsequent chapters to include the crucial inferential substructures of *substitution* and *anaphora*. Substitutional commitments are defined as a species of inferential commitments (in Chapter 6) by distinguishing a class of substitution inferences. In this way the inferentialist paradigm can be extended so as to apply to the conceptual contents of subsentential expressions such as singular terms and predicates. Then anaphoric commitments are defined in terms of the *inheritance* of substitution-inferential commitments (in Chapter 7). In this way the inferentialist paradigm can be extended so as to apply to unrepeatable or token-reflexive expressions, such as demonstratives, indexicals, and pronouns.

It is important to keep in mind in reading what follows that the inferentialist project pursued here is a defense of the strong, not only the weak, inferentialist thesis. But it is not a form of hyperinferentialism. And while it eschews representational semantic primitives in favor of others more easily grounded in pragmatics, this is not because of a denial of the importance of the representational dimension of discursive practice. On the contrary, that choice serves rather an aspiration to make intelligible in a new way just what that representational dimension consists in.

VI. CONCLUSION

1. *Grounding an Inferential Semantics on a Normative Pragmatics*

Inferentialism about conceptual content is not an explanatory strategy that can be pursued in complete abstraction from pragmatism about the norms implicit in the practical application of concepts. The considerations assembled here to motivate and recommend an inferentialist order of

semantic explanation appeal to a notion of *materially correct inferences*. In this chapter material proprieties of inference have been treated as primitives, playing the role of unexplained explainers. A critical criterion of adequacy by which such an approach should be assessed is clearly the extent to which a philosophically satisfying story can be told about these primitive proprieties of nonformal inference. The semantic theorist's entitlement to explanatory use of such primitives must be vindicated by situating the project of semantic theory in a broader context. Conceptual contents, paradigmatically propositional ones, are associated with linguistic expressions as part of an attempt to specify, systematically and explicitly, the correct *use* of those expressions. Such contents are associated with intentional states such as belief as part of a corresponding attempt to specify their behavioral significance—the difference those states make to what it is appropriate for the one to whom they are attributed to *do*.

The study of the practical significance of intentional states, attitudes, and performances (including speech acts) is *pragmatics*, as that term is used here. The projects of semantic theory and of pragmatic theory are intricately interrelated. If the semantic content and pragmatic context of a linguistic performance of a particular kind (paradigmatically assertion) are specified, a general theory of speech acts seeks to determine in a systematic way the pragmatic significance of that contentful performance in that context. But besides the direction of explanation involved in the local determination of pragmatic significance by semantic content, there is also a converse direction of explanation involved in the global conferral of semantic content by pragmatic significance. It must be explained how expressions can be used so as to confer on them the contents they have—what functional role the states they manifest must play in practice for them to be correctly interpreted as having certain intentional contents. Such an explanation amounts to an account of what it is for a state, attitude, performance, or expression to be *propositionally* contentful. Once a general notion of content has been made sense of in this way, particular attributions of contentfulness can then be offered as part of explanations or explicit specifications of the pragmatic significance of a state, attitude, performance, or expression.

The discussion of the next chapter should begin to make clearer just how a story about the conferral of content by practice is envisaged as relating to the use of attributions of content in the determination of pragmatic significances. One aspect of the situation of the semantic concept of content in a wider pragmatic context, however, is of particular relevance to the issue of entitlement to appeal to material proprieties of inference as semantic primitives. For the inferential proprieties that from the point of view of semantic theory are treated as primitive can be explained in the pragmatic theory as implicit in discursive practice (which includes intentional agency). An inferential move's normative status as correct or incorrect can be construed as instituted in the first instance by practical attitudes of taking or treating it

as correct or incorrect. The inferential norms that govern the use of expressions (or the significance of states, attitudes, and performances) are then understood as instituted by practical attitudes toward what the content is attributed to; they in turn confer that content on what it is attributed to.

Expressions come to mean what they mean by being used as they are in practice, and intentional states and attitudes have the contents they do in virtue of the role they play in the behavioral economy of those to whom they are attributed. Content is understood in terms of proprieties of inference, and those are understood in terms of the norm-instituting attitudes of taking or treating moves as appropriate or inappropriate in practice. A theoretical route is accordingly made available from what people *do* to what they *mean*, from their *practice* to the *contents* of their states and expressions. In this way a suitable pragmatic theory can ground an inferentialist semantic theory; its explanations of what it is in practice to treat inferences as correct are what ultimately license appeal to material proprieties of inference, which can then function as semantic primitives.

Sketching the possibility of such an explanatory path from attributions of practical attitudes to attributions of semantic content should help alleviate one sort of worry that might be elicited by the inferentialist invocation of materially correct inferences in explaining conceptual contentfulness. For otherwise the employment of a notion of *material* proprieties of inference in explaining *content* might seem blatantly circular. After all, are not materially good inferences just those that are good in virtue of the contents of the nonlogical concepts applied in their premises and conclusions, by contrast to the logically valid inferences, which are good in virtue of the logical form of those premises and conclusions? Presystematically, this is indeed how they should be thought of. But officially, the strategy is to start with proprieties of inference and to elucidate the notion of conceptual content in terms of those proprieties.

Talk of materially correct inferences is indeed intended to enforce a contrast with those that are formally correct (in the sense of logically valid). But the force of this contrast is just that the validity of inferences in virtue of their logical form is to be understood as a sophisticated, late-coming sort of propriety of inference, founded and conceptually parasitic on a more primitive sort of propriety of inference. This is the repudiation of the formalist approach to inference, for which the correctness of inference is intelligible only as formal logical validity, correctness in virtue of logical form. Calling the more primitive sort of propriety of inference *materially* correct simply registers the rejection of this order of explanation. It does not involve commitment to a prior notion of nonlogical content. If what it means to call an inference correct in the relevant sense can be explained without appeal to the use of logical concepts—for instance in terms of conduct interpretable as a practical taking or treating of an inference as correct—then there need be no circularity in appeal to such inferential proprieties in elaborating a notion of conceptual content.

2. *Knowing-That in Terms of Knowing-How, Formal Proprieties of Inference in Terms of Material Ones, Representational Content in Terms of Inferential Content*

A story that begins with inferring as a kind of practical *doing* and that leads to an account of the specifically *propositional* contentfulness of speech acts and intentional states holds out the promise of yielding an account of propositionally *explicit* saying, judging, or knowing-*that*, in terms of practically *implicit* capacities, abilities, or knowing-*how*. This would discharge one of the primary explanatory obligations of the pragmatist foe of the intellectualist understanding of norms. For if practical knowing-how is taken as prior in the order of explanation to theoretical knowing-that, one must not only offer an independent account of the practically implicit grasp or mastery of norms. One must also explain how the propositionally explicit grasp of norms expressed in the form of rules, principles, or claims can be understood as arising out of those practical capacities.

In the same way, the inferentialist approach to content treats material proprieties of inferences as prior in the order of explanation to formal logical proprieties of inference. It is accordingly obliged not only to offer an independent account of those material proprieties but also to offer an account of how logical goodness of inference can be explained in terms of that primitive sort of goodness of inference. One who denies that logic is to be understood as underlying (and so presupposed by) rationality in the sense involved in the inferential articulation of conceptual contents (and so in any exercise of the capacity to give and ask for reasons) is obliged to offer another account of logic. This obligation is discharged by the combination of two moves. The first is offering a criterion of demarcation for logical vocabulary that is couched in terms of the semantically expressive role played by such vocabulary in making implicitly content-conferring inferential commitments explicit in the form of judgments. This move depends on having a view about what it is for something to be explicit in the form of a judgeable, that is propositional, content. Such a view is precisely what the account of propositional contents in terms of material proprieties implicit in inferential practice, mentioned above, is intended to supply. The second element required to discharge the obligation to show how the notion of logically good inferences grows out of that of materially good inferences is the substitutional account of formal logical validity of inference—according to which an inference is valid or good in virtue of its logical form if it is primitively good and cannot be turned into one that is not primitively good by any (grammatical) substitution of nonlogical for nonlogical vocabulary.

An explanatory demand exhibiting the same structure as that just rehearsed for the anti-intellectualist about norms and the antiformalist about logic is incumbent on the inferentialist account of conceptual content in virtue of its commitment to invert the representationalist order of semantic explanation.⁹⁰ A viable working-out of the inferentialist order of explanation

must, to begin with, offer an account of correctness of inference that is not parasitic on correctness of representation. This demand is addressed in Chapters 3 and 4, which specify sufficient conditions for an attribution of implicitly normative social practices to a community to count as interpreting them as engaging in practices of giving and asking for reasons—as practically assessing inferences as correct or incorrect, and so as instituting material inferential proprieties that confer propositional conceptual content on their states and performances. It is not enough, however, for the inferentialist explanatory strategy to produce an account of the pragmatic basis of its own semantic primitives that does not rely on the prior intelligibility of representational concepts. It must also show how representational relations and the sorts of representational content they underwrite can be made intelligible in terms of those inferential primitives. That is, another critical criterion of adequacy of inferentialism is the extent to which, if this approach is granted its preferred starting point, it can develop it into an account of the sort of objective representational content other approaches begin with.

3. *Objective Representational Content*

Meeting this demand involves offering accounts of three important dimensions along which the notion of objective representational content is articulated. First is the referential dimension. The representationalist tradition has, beginning with Frege, developed rich accounts of *inference* in terms of *reference*. How is it possible conversely to make sense of reference in terms of inference? In the absence of such an account, the inferentialist's attempt to turn the explanatory tables on the representationalist tradition must be deemed desperate and unsuccessful.

The second dimension is categorial. An account must be offered not just of reference and representation but of reference to and representation of particular *objects* and general *properties*. That is, the peculiar kind of representational content expressed by *subsential* expressions, paradigmatically singular terms and predicates, must be explained. For reasons already indicated, inferential approaches to conceptual content apply directly only to what is expressed by declarative sentences, which can play the role of premises and conclusions of inferences. Somehow the inferential approach to conceptual content must be extended to apply to subsential parts of speech as well. The discussion of Dummett's model of circumstances and consequences of application provides some suggestive hints. But these must be developed far beyond the remarks already offered in order to put the inferentialist in a position to claim to have shown that the nominalist order of explanation standard prior to Kant, beginning with a doctrine of terms or concepts and moving from there to a doctrine of judgments, can successfully be stood on its head.

Finally, there is the *objective* dimension of representational content. It

must be shown how on inferentialist grounds it is possible to fund *objective* proprieties of inferring and judging—to make intelligible the way in which what it is correct to conclude or to say depends on how the objects referred to, talked about, or represented actually are. Even if, to begin with, attention is restricted to inferential proprieties, it is clear that not just any notion of correctness of inference will do as a rendering of the sort of content we take our claims and beliefs to have. A semantically adequate notion of correct inference must generate an acceptable notion of conceptual content. But such a notion must fund the idea of *objective* truth conditions and so of *objectively* correct inferences. Such proprieties of judgment and inference outrun actual attitudes of taking or treating judgments and inferences as correct. They are determined by how things actually are, independently of how they are taken to be. Our cognitive attitudes must ultimately answer to these attitude-transcendent facts.

This means that although the inferentialist order of explanation may start with inferences that are correct in the sense that they are accepted in the practice of a community, it cannot end there. It must somehow move beyond this sense of correctness if it is to reach a notion of propositional conceptual content recognizable as that expressed by our ordinary empirical claims and possessed by our ordinary empirical beliefs. Pursuing the inferentialist order of explanation as outlined above accordingly requires explaining how—if actual practical attitudes of taking or treating as correct institute the normative statuses of materially correct inferences, and these material proprieties of inference in turn confer conceptual content—that content nonetheless involves objective proprieties to which the practical attitudes underlying the meanings themselves answer.⁹¹ How is it possible for our use of an expression to confer on it a content that settles that we might all be wrong about how it is correctly used, at least in some cases? How can normative attitudes of taking or treating applications of concepts as correct or incorrect institute normative statuses that transcend those attitudes in the sense that the instituting attitudes can be assessed according to those instituted norms and found wanting? This issue of objectivity is perhaps the most serious conceptual challenge facing any attempt to ground the proprieties governing concept use in social practice—and the pragmatist version of inferentialism being pursued here is a view of this stripe.

In the terms set up in Section I of this chapter, the referential, categorial, and objective can be thought of as three interlocking dimensions of the project of explaining *object-representing* contentfulness in terms of *propositional* contentfulness, according to a semantic rendering of propositional contentfulness in terms of *material proprieties of inference* and a *pragmatic* rendering of those basic inferential proprieties.⁹² The relation between inference and reference is discussed in an introductory way in Chapter 5, which examines the use of the semantic vocabulary (paradigmatically 'refers' and 'true') by whose means the implicit referential dimension of conceptual

contents is made explicit. The anaphoric relations that are invoked in the explanation offered there are then explained in more primitive pragmatic terms in Chapter 7, which relates them to the sort of substitution-inferential commitments discussed in Chapter 6. The categorial issue is addressed by further development of Frege's substitutional methodology, in Chapter 6, which shows how the contents expressed by the use of singular terms and predicates can be understood in terms of substitution inferences.

The objectivity issue, which concerns the relation between what is properly *said* and what is talked *about*, intimately involves both the referential or representational dimension and the categorial. As with all substantive semantic and pragmatic concepts officially employed in this work in describing the contents and significances of discursive commitments, the concept of states, attitudes, and performances that are objectively contentful in purporting to represent how things are independently of anyone's states, attitudes, and performances is discussed at two levels. (This two-leveled account is part of the effort to secure for the use of such vocabulary in this project an analog of the kind of *expressive equilibrium* already adverted to as achieved by Frege's treatment of sentential logical vocabulary in the *Begriffsschrift*.) The first is a story about what it is for such purport and its uptake to be *implicit* in the practices of those whose states, attitudes, and performances are properly interpreted as having such content.

The second is a story about what it is for such purport and its uptake to be made *explicit* in the specification of contents of ascribed states, attitudes, and performances. Although technical philosophical vocabulary such as 'refers' or 'denotes' (discussed in Chapter 5) can play this explicating role, the fundamental locutions used in ordinary talk to express representational commitments are those used to form *de re* specifications of the contents of ascribed intentional states, attitudes, and performances—paradigmatically 'of' and 'about'. The use of *de re* ascriptions makes it possible to specify explicitly what is said in terms of what is talked about. What such ascriptions express and how those objective content-specifications are made explicit by their use is the topic of Chapter 8. The account of what we are doing when we interpret ourselves and each other as making claims with objective representational content that is offered in that chapter requires the expressive resources of all of the sorts of locutions whose use is introduced in prior chapters. It is accordingly only in the last substantive chapter of this work that this critical explanatory obligation of an inferential approach to semantics is finally discharged.

The foundation of that account is laid in the next chapter. It consists in the *social* structure of the inferential norms that confer propositional content. (Government by such norms is what such contentfulness is.) The development into a more full-blooded notion of conceptual content of the abstract notion of inferential role introduced in this chapter proceeds by taking account of the social dimension of inferential practice—which is

implicit in the idea that abstract talk about inferential relations must be rooted in consideration of what Sellars calls "the game of giving and asking for reasons." The pragmatic significance of making a claim or acquiring a commitment whose content could be expressed by the use of a declarative sentence cannot be determined by associating with that sentence a set of sentences that entail it and a set of sentences that it entails—not even if these are enriched by throwing in nonlinguistic circumstances and consequences of application as well. This is because of the interaction of two features of inferentially articulated commitments.

First, as Frege acknowledges in his original definition of *begriffliche Inhalt*, specification of the inferential role of a sentence requires looking at *multipremise* inferences.⁹³ Many of the important "consequences of application" of a sentence are not consequences it has all on its own; they consist rather in the differential contribution its inclusion makes to the consequences of a set of collateral premises or auxiliary hypotheses. Similarly, its purely inferential antecedents must be thought of not as individual sentences but as sets of them.

Second, the collateral concomitant commitments available as auxiliary hypotheses in multipremise inferences vary from individual to individual (and from occasion to occasion or context to context). If they did not, not only the notion of communication but even that of empirical information would find no application. The significance of acquiring a commitment or making a claim whose content could be expressed by the use of a particular sentence, when it would be appropriate to do so and what the appropriate consequences of doing so would be, depends on what other commitments are available as further premises in assessing grounds and consequences. What is an appropriate ground or consequence of that commitment from the point of view of one set of background beliefs may not be from the point of view of another. In view of the difference in their other attitudes, a single commitment typically has a different significance for the one undertaking it, a speaker or believer, from that which it would have for those attributing it, an audience or intentional interpreter. Of course this does not make communication or interpretation impossible—on the contrary. As was just mentioned, it is only the prevalence of situations in which background commitments do differ that give communication and interpretation their point.

The fact that the implicitly normative inferential significance of a commitment may be different from the point of view of one undertaking the commitment and one attributing it means that the inferential articulation of conceptual contents has a fundamental *social* dimension. It introduces a relativity to social perspective into the specification of such contents. The practical attitudes of taking or treating as committed, which ultimately institute the normative status of commitment, come in the two socially distinct flavors of undertaking or acknowledging a commitment (oneself) and

attributing a commitment (to another). Inferentially articulated contents are conferred on states, attitudes, and performances by the norms instituted by *social* practices: those that essentially involve the interaction of attitudes corresponding to both social perspectives. Investigation of the use of locutions that make explicit various aspects of the social perspectival character of conceptual contents will reveal what they express as the source of objective representational content. So, it will be claimed, what must be added to the *normative* approach to pragmatics and the *inferential* approach to semantics in order to make intelligible the *representational* dimension of conceptual contents is a *social* account of the interaction between them.