3

Linguistic Practice and Discursive Commitment

Language most shows a man: Speak, that I may see thee.
BEN JONSON, Timber or Discoveries

Language is called the Garment of Thought: however, it should rather be,
Language is the Flesh-Garment, the Body, of thought.
THOMAS CARLYLE, Sartor Resartus

Clearly human beings could dispense with all discourse, though only at
the expense of having nothing to say.
WILFRID SELLARS, "A Semantical Solution to
the Mind-Body Problem"

I. INTENTIONAL STATES AND LINGUISTIC PRACTICES

1. Discursive Practice as Deontic Scorekeeping

This chapter introduces a particular model of language use: the deontic scorekeeping model of discursive practice. The implicitly normative social practices it describes are inferentially articulated in such a way as to confer specifically *propositional* contents on expressions and performances that play suitable roles in those practices. The basic idea is the one motivated by the discussion in Chapter 2, namely that propositional contentfulness must be understood in terms of practices of giving and asking for *reasons*. A central contention is that such practices must be understood as *social* practices—indeed, as *linguistic* practices. The fundamental sort of move in the game of giving and asking for reasons is making a *claim*—producing a performance that is propositionally contentful in that it can be the offering of a reason, and reasons can be demanded for it. Other theoretically important concepts are defined in terms of this one: *linguistic* practice is distinguished by its according some performances the significance of claimings, and (declarative) *sentences* are distinguished as expressions whose utterances, inscriptions, or other tokenings have the default significance of claimings. The basic explanatory challenge faced by the model is to say what structure a set of social practices must exhibit in order properly to be understood as includ-

ing practical attitudes of taking or treating performances as having the significance of claims or assertions.

According to the model, to treat a performance as an assertion is to treat it as the undertaking or acknowledging of a certain kind of *commitment*—what will be called a 'doxastic', or 'assertional', commitment. To be doxastically committed is to have a certain social status. Doxastic commitments are normative, more specifically *deontic*, statuses. Such statuses are creatures of the practical attitudes of the members of a linguistic community—they are instituted by practices governing the taking and treating of individuals *as* committed. Doxastic commitments are essentially a kind of deontic status for which the question of *entitlement* can arise. Their inferential articulation, in virtue of which they deserve to be understood as propositionally contentful, consists in consequential relations among the particular doxastic commitments and entitlements—the ways in which one claim can commit or entitle one to others (for which it accordingly can serve as a reason).

Competent linguistic practitioners keep track of their own and each other's commitments and entitlements. They are (we are) *deontic scorekeepers.* Speech acts, paradigmatically assertions, alter the deontic score; they change what commitments and entitlements it is appropriate to attribute, not only to the one producing the speech act, but also to those to whom it is addressed. The job of pragmatic theory is to explain the *significance* of various sorts of speech acts in terms of practical proprieties governing the keeping of deontic score—what moves are appropriate given a certain score, and what difference those moves make to that score. The job of semantic theory is to develop a notion of the *contents* of discursive commitments (and the performances that express them) that combines with the account of the significance of different kinds of speech act to determine a scorekeeping kinematics.

The basic elements of this deontic scorekeeping model of discursive practice are presented in this chapter. The next chapter develops it further by attending in particular to the inferential articulation of *perception* and *action.* These are the entries to and exits from the realm of discursive commitments and entitlements—the source respectively of the empirical and practical dimensions of conceptual content, which are usually (and in one sense correctly) thought of as *non*inferential. Before plunging into a description of the details of the features of a system of social practices in virtue of which they should be understood as incorporating practical scorekeeping attitudes that institute deontic statuses and confer propositional contents, however, some methodological preliminaries are in order. The rest of this work presents not only an account of linguistic intentionality (thought of as one sophisticated species in a genus comprising other, more primitive sorts) but a linguistic account of intentionality generally. It is claimed that the propositional contentfulness even of the beliefs and other states intentional

interpreters attribute to nonlinguistic animals cannot properly be understood without reference to the specifically linguistic practice of the interpreters, from which it is derived. Original, independent, or nonderivative intentionality is an exclusively linguistic affair. The reasons for insisting on the conceptual primacy of linguistic intentionality cannot be presented until all the materials needed for the analysis of the representational dimension of propositional content (and of conceptual content generally) have been assembled, in Chapter 8. Nonetheless, the explanatory strategy being pursued will be easier to understand if the picture of the relations between language and belief that it incorporates has been sketched, even if the warrant for that picture cannot emerge until it is more fully developed (in Part 2 of this work).

2. Philosophical Semantics and Formal Semantics

One of the fundamental methodological commitments governing the account presented here is *pragmatism* about the relations between semantics and pragmatics. Pragmatism in this sense is the view that what attributions of semantic contentfulness are *for* is explaining the normative significance of intentional states such as beliefs and of speech acts such as assertions. Thus the criteria of adequacy to which semantic theory's concept of content must answer are to be set by the pragmatic theory, which deals with contentful intentional states and the sentences used to express them in speech acts. The idea that *philosophical* theories of meaning or content must be concerned with the larger pragmatic context within which attributions of contentfulness play an explanatory role may seem to be brought into question by the autonomy of *formal* semantics. But the independence of formal semantics from pragmatic concerns is only apparent.

The project of formal semantics entitles the theorist at the outset to stipulate an association of semantic interpretants with primitive interpreteds, typically linguistic expressions. Then this interpretation is extended to interpreteds that are derived from those primitives by syntactic operations— which for standard compositional syntactic structures include category-sensitive concatenation and various grammatical transformations of such concatenations. This is achieved by defining, for each syntactic operation on interpreteds, a corresponding operation on their associated interpretants that yields a new interpretant, which is thereby associated with the result of the syntactic operation. So the formal semantic theorist might begin by associating truth-values with sentence-letters, and then for each connective that produces compound sentences introduce a function taking sets of truth-values into truth-values that can then be assigned to the corresponding compound sentences. Or instead of truth-values, the semantic interpretants might be sets of possible worlds, and the operations corresponding to sentential connectives be set-theoretic operations on them (such as intersection for conjunction).[1]

So formal semantics is concerned generically with structure-preserving mappings. But not every mathematical representation theorem, which shows such a correspondence between structures of one kind and those of another, deserves to be called a semantics. What else ought to be required for a set of such mappings to count as presenting a specifically *semantic* interpretation of something? To ask this question is already to begin to move from the domain of purely formal semantics to that of philosophical semantics. When Tarski proved an algebraic representation theorem in which the interpretants assigned to quantificational expressions are topological closure operators, what qualifies that as a formal semantics for the first-order predicate calculus is not anything about the intrinsic properties of those interpretants but just that he is able in those terms to reproduce the relation of logical consequence appropriate to that idiom. From a purely formal or mathematical point of view, the task would be no different if the property to be reproduced were specified simply by randomly partitioning the elements of one grammatical category, placing an asterisk next to some of them and not others (and similarly for the relation in question). From the point of view of the philosophical motivation of calling what one is doing thereby 'semantics', however, it makes all the difference that the elements involved be interpretable as *sentences*, and that the property distinguished be interpretable as theorem-hood, a kind of *truth*, and that the relation distinguished be interpretable as derivation, a kind of *inference*. Indeed, to take the elements as subject to evaluations concerning propriety of *judging* and propriety of *reasoning*, truth and inference, is just what it is to interpret them as sentences.

What gives semantic theory its philosophical point is the contribution that its investigation of the nature of contentfulness can make to the understanding of proprieties of practice, paradigmatically of judging and inferring. That semantic theory is embedded in this way in a larger explanatory matrix is accordingly important for how it is appropriate to conceive the semantic interpretants associated with what is interpreted. It means that it is pointless to attribute semantic structure or content that does no pragmatic explanatory work. It is only insofar as it is appealed to in explaining the circumstances under which judgments and inferences are properly made and the proper consequences of doing so that something associated by the theorist with interpreted states or expressions qualifies as a *semantic* interpretant, or deserves to be called a theoretical concept of *content*. Dummett puts the point this way:

> The term 'semantics', at least as commonly applied to formalized languages, usually denotes a systematic account of the truth-conditions of sentences of the language: the purpose of thus assigning a value, true or false, to every well-formed sentence of the language is taken as already understood, and receives no explanation within the semantic theory itself . . . The classification of the sentences of a formalized

language into true ones and false ones relates to the purposes for which we want to use the language. But in the case of natural language, it is already in use: the only point of constructing a semantics for the language can be as an instrument for the systematic description of that use, that is, as part of a whole theory of meaning for the language, which as a whole constitutes an account of its working. If the semantic part of the theory is taken as issuing an assignment of conditions under which each sentence of the language, as uttered on a particular occasion, has this or that truth-value, the rest of the theory must connect the truth conditions of the sentences with the use to which they are put, that is, with the actual practice of speakers of the language . . . a semantic theory which determines the truth-conditions of sentences of a language gets its point from a systematic connection between the notions of truth and falsity and the practice of using those sentences.[2]

The essential point is that philosophical semantic theory incorporates an obligation to make the semantic notions it appeals to intelligible in terms of their pragmatic significance. Formal semantics qualifies as *semantics* only insofar as it is implicitly presupposed that this obligation can be satisfied by conjoining the semantics with some suitable pragmatics.[3]

Philosophical semantics is distinguished from formal semantics by its explicit concern with the relation between the use of semantic concepts, on the one hand, and pragmatic accounts of the proprieties of practice governing the employment of what those concepts apply to, on the other. Philosophical semantics is committed to explaining the content of concepts such as *content, truth, inference, reference,* and *representation,* while formal semantics is content to use such concepts, assuming them (and so the pragmatic significance of applying them) already to be implicitly intelligible. The difference between doing either sort of semantics for artificial languages and for natural languages is that in the former case there are no antecedent proprieties governing the use of the expressions, to which the semantic theorist is responsible. Since the language is not already in use, the theorist is free to *stipulate* an association of contents with expressions, in order to determine how they are to be understood to be correctly used. In the case of natural languages, however, the theorist's use of semantic concepts is not *synthetic* (to settle the proper employment of expressions that antecedently are subject to no such proprieties) but *analytic* (to codify and express antecedently existing proprieties of employment).

3. Associating Content Explicitly by Stipulation and Implicitly Conferring It by Practice

Philosophical semantic theories of expressions and states that already play normatively articulated roles in linguistic practice or in the

practical reasoning of rational agents accordingly cannot afford the luxury (enjoyed by formal semantics of all sorts and even by philosophical semantics of artificial languages) of employing a *stipulative* method. Such theories are obliged to explain what the association of content with expression or state *consists in:* what one is saying or doing in attributing content to them. At this point it has seemed to many that the cases of contentful *sentences* and of contentful *beliefs* diverge. It makes sense to think of the contents of linguistic expressions as *conferred* on them by the way they are used. Noises and marks mean nothing all on their own. No one thinks they are *intrinsically* contentful. The sign-design 'dog' could as well be used to express the concept expressed by the sign-design 'horse', or to express none at all, like 'gleeb'. It is only by being caught up in linguistic practice that they come to express propositions, make claims, have or express conceptual or intentional contents. Apart from their role in human activity, apart from the norms thereby imposed on their employment—which make it the case that some uses are correct and others incorrect—these linguistic vehicles are semantically mute, inert, dead.

The philosophical semantics of natural languages must begin, then, with the observation that it is the practice of those who use the language that confers content on the utterances and inscriptions that are the overt, explicitly expressive performances whose propriety is governed by that practice. Is something similar true of intentional states? There are some important asymmetries between the two cases. There is a familiar line of thought, already adverted to, according to which quite a different story must be told about the association of content with the states and attitudes, paradigmatically beliefs, that are expressed by such linguistic performances. The critical question is how to understand the *use* of language in which the pragmatic significance of speech acts consists and which accordingly confers semantic content on those speech acts and so indirectly on the expressions they involve. One way to think about such use is *instrumentally*. This line of thought may be traced back to Locke, who thought of speech as an instrument for communicating thoughts or ideas. It is successful when the noise emitted by the speaker arouses in the audience an idea with the same content as that prompting the speech act.

Contemporary elaborations of this approach see "nonnatural" meaning as rooted in the capacity of individuals deliberately to imbue signals with significance by producing them with the intention that they be understood in a certain way by their auditors. According to Grice's picture,[4] linguistic practitioners *make* their expressions have a certain content by producing them with the intention that others *take* them to have that content. In particular, assertion is understood as the expression of belief in the sense that a sentence is produced with the intention that those who hear it will acquire a certain belief in virtue of their recognition that the speaker uttered the sentence intending those who hear it to acquire that belief, in virtue of their

recognition of that very intention. The notion of a linguistic expression's meaning something is in this way derived from the capacity of language users deliberately to mean something by their utterances. In somewhat different ways, Lewis, Bennett, and Searle develop this instrumental model by showing how a layer of *conventions* can be built on such communicative *intentions*, in such a way that members of a linguistic community are for the most part relieved of the necessity for elaborate deliberation about each other's beliefs and intentions in choosing and interpreting each other's remarks.[5] The foundation on which the conventional meaningfulness of linguistic acts and expressions rests remains their intentional employment as means to an explicitly envisaged communicational end.[6]

Rosenberg calls this explanatory strategy "agent semantics," because linguistic meaning is explained in terms of a prior capacity to engage in practical reasoning.[7] If the pragmatic use of language that confers semantic content on utterances and expressions is understood in these terms, it is clear that the contentfulness of intentional states such as beliefs, desires, and intentions must be understood antecedently, and hence according to some other model. Agent semantics treats the contentfulness of utterances as derivative from that of intentional states. The content of an assertion derives from the content of the belief it is the expression of, and from the content of the intention that it be understood as expressing that belief. It follows that it must be possible to make sense of the contents of beliefs and intentions prior to and independently of telling this sort of story about the use of linguistic expressions. Their content cannot be taken to be conferred on them by the way they are used or employed, according to this model of use or employment.[8]

That the content of intentional states cannot be understood as conferred on them by proprieties governing their significance—when it is appropriate to acquire them and what the appropriate consequences of acquiring them are—follows only if the only candidate for content-conferring use is deliberate, instrumental employment in order to secure the *explicitly* envisaged purpose of being understood as having a certain content. It would not follow that semantic content could not be conferred on intentional states by proprieties *implicit* in the way those states are treated in practice. According to such a conception, the conferral of content might be a side effect of the way they are treated, not requiring that anyone explicitly intend to confer it by their behavior. Broadly functionalist approaches to content are of this sort. They understand intentional states to be contentful in virtue of the role they play in the proper functioning of some system of which they are a part. Going into a certain state is something that is done appropriately under some circumstances, according to the functional interpretation of the system, and it has certain appropriate consequences. Together these proprieties of input and output, antecedents and consequences, determine the functional role of the state in the system. According to the functionalist explanatory strategy,

it is in virtue of playing this role, being subject to these transitional proprieties, that intentional states have the content they do.

4. Intentionality: Linguistic Practice versus Rational Agency

The first question that needs to be addressed in working out such an approach is how the relevant functional system should be understood. Is it possible to understand propositional and other genuinely conceptual contents as conferred on states and performances by their role in a functional system that comprises only a single individual? Or is discursive practice essentially *social* practice, so that the functional system must be taken to comprise the activities of an entire community? The most popular and promising way of developing the first answer looks to the role belief plays in the practical instrumental reasoning of intelligent agents. The most popular and promising way of developing the second answer looks to the role assertion—the explicit expression of belief—plays in linguistic practice.

The considerations assembled in the first two chapters suggest the motivation that these two approaches have in common: states, attitudes, and performances are intentionally contentful in virtue of the role they play in *inferentially* articulated, implicitly *normative* practices. It is by looking at the practices in which the status of some states, attitudes, and performances as providing *reasons* for others is implicitly (and constitutively) acknowledged that the pragmatic significance of associating them with intentional contents is to be understood. There are two different sorts of context in which the specifically inferential significance of intentional states such as belief is to the fore: rational agency and linguistic practice. On the one hand, beliefs and other intentional states are expressed in *actions*, nonlinguistic performances that are intelligible in virtue of the beliefs and desires that are reasons for them. On the other hand, beliefs are expressed in *claims*. Overt assertions are the fundamental counters in the game of giving and asking for reasons—they can both be offered as reasons and themselves stand in need of such reasons.

Stalnaker points out that these two sorts of context in which intentional states are significant for practice generate two basically different ways of looking at those states: what he calls the *pragmatic* picture of intentionality and the *linguistic* picture of intentionality. The pragmatic picture is one according to which "rational creatures are essentially agents . . . According to this picture, our conceptions of belief and of attitudes pro and con are conceptions of states which explain why a rational agent does what he does . . . Linguistic action, according to this picture, has no special status. Speech is just one kind of action which is to be explained and evaluated according to the same pattern. Linguistic action may be a particularly rich source of evidence about the speaker's attitudes, but it has no special conceptual connection with them."[9] This picture amounts to a generalization of the ap-

proach of agent semantics. It shares the emphasis on rational agency as the conceptually and explanatorily fundamental context in which to understand the significance of the contentfulness of intentional states. It is more general in that it does not necessarily involve a commitment to understanding the contentfulness of speech acts as deriving from their deliberate instrumental employment to secure antecedently envisageable goals. Thus this picture leaves room for a picture in which the way speech acts inherit the contents of the intentional states they express might be less intellectualized, the conferral of content being implicit in the practice of expression, rather than explicit as its instrumentally conceived motive. The contrasting linguistic picture is one according to which "rational creatures are essentially speakers . . . Representational mental states represent the world because of their resemblance to or relation with, the most basic kind of representations: linguistic expressions."[10]

This way of dividing up the fundamental orientations of various approaches to intentionality, accordingly as rational agency or linguistic capacity is taken as primary, evidently cuts at important joints. It is a measure of the robustness of this botanization that it is serviceable even across large differences in collateral theoretical commitments. Here is how Stalnaker puts the setting in which he sees the pictures as competing: "The linguistic and pragmatic pictures each suggest strategies for giving a naturalistic explanation of representation—both mental and linguistic representation—but the two strategies differ in what kind of representation they take to be more fundamental. The pragmatic picture suggests that we explain the intentionality of language in terms of the intentionality of mental states, while the linguistic strategy suggests that we explain the intentionality of mental states in terms of, or by analogy with, the intentionality of linguistic expressions."[11] On the side of semantic content, Stalnaker follows the tradition in seeing the issue as fundamentally one of representation, with inferential relations presumably to be explained further along in ultimately representational terms. And on the side of the pragmatic significance of intentional contentfulness, Stalnaker begins with a commitment to a naturalistic approach, with the normative character of the practice in which intentional states are significant (whether that practice is conceived in the first instance as rational agency or as essentially linguistic) presumably to be explained further along in ultimately naturalistic terms.

The order of explanation that frames his discussion is the reverse of that pursued in this work. For the semantic explanatory strategy being developed here looks first to inference, on the semantic side, and aspires to making the representational dimension of intentional content intelligible ultimately in inferential terms. And on the pragmatic side, the strategy is to begin with an account of norms implicit in practice and work out toward an understanding of their relation to their naturalistic setting, which the normative practices in their most sophisticated form make it possible to describe objectively. It

is noteworthy that in spite of these major differences in approach, the large division of options for explaining intentionality into those focusing on rational agency and those focusing on language seems compelling from both points of view.

5. Analogical and Relational Versions of Linguistic Approaches to Intentionality

Broadly linguistic approaches comprise many importantly different variants, however, and these correspond to importantly different motivations. Stalnaker implicitly acknowledges one significant subdivision within the linguistic approach in his general characterization of the linguistic picture as seeking to explain the contentfulness of intentional states by appealing to their "resemblance to or relation with" contentful linguistic performances. The disjunction links two very different ways in which it might be thought that taking account of specifically linguistic practice is essential to making the contentfulness of intentional states explicit (in the sense of theoretically intelligible). The resemblance limb, according to which the contentfulness of states is modeled on that of expressions, involves a commitment to the claim that the theorist's or interpreter's capacity to attribute (and understand attributions of) contentful intentional states is parasitic on the theorist's or interpreter's capacity to attribute (and understand attributions of) contentful speech acts. It need not entail, for instance, that the intentional states attributed to nonlinguistic creatures are somehow second class. It requires only that what one is doing in attributing contentful states to nonlinguistic creatures cannot be understood apart from the capacity to attribute them to linguistic ones. By contrast, the relational limb—according to which the contentfulness of intentional states consists in or essentially involves the contentfulness of the speech acts that express them—involves a commitment to the claim that the theorist's or interpreter's capacity to attribute (and understand attributions of) contentful intentional states is in the first instance parasitic on the theorist's or interpreter's capacity to attribute (and understand attributions of) contentful speech acts to the *same* individuals who are taken to have the intentional states. It does entail that the intentional states attributed to nonlinguistic creatures are in important regards second-class statuses.[12]

One reason it is important to distinguish the claim that the intentionality of speech is conceptually prior to the intentionality of belief (as analogical theories have it) from the claim that as believers in the full sense, we are essentially rather than only accidentally speakers (as relational theories have it) is that only theories committed to the former thesis are obliged to offer accounts of linguistic practice that do not make reference to intentional states. It is open to one who subscribes to the second view to hold, as Davidson does, that attributions of contentful intentional states and content-

ful speech acts go hand in hand, that neither sort of attribution is intelligible apart from its relation to the other. So a relational account can understand the possibility of speech as essential to intentionality (in the paradigm cases from which we derive our grip on what intentional interpretation is) without thereby becoming obliged simply to invert the order of explanation characteristic of agent semantics—though such an account evidently cannot appeal to the sort of independently intelligible role of contentful states in rational agency presupposed by agent semantics.

In contrast, analogical linguistic theories of intentionality are committed to that converse order of explanation. Agent semantics employs an antecedent and independent notion of the contentfulness of intentional states to explain the derivative contentfulness of speech acts and linguistic expressions. A theory insisting that the contentfulness of intentional states is intelligible only by analogy to the contentfulness of speech acts and linguistic expressions would be obliged correspondingly to appeal to an antecedent and independent notion of the contentfulness of speech in order to explain the derivative contentfulness of intentional states. A relational linguistic theory of intentionality maintains rather that the understanding of intentionally contentful states that permits us to stretch the application of that notion and apply it in a second-class way to nonlinguistic animals (simple intentional systems) derives from and essentially depends on an understanding of the relation between the intentional states and the linguistic performances of language-using animals (communicating or interpreting intentional systems).

According to this sort of approach, understanding intentionality requires looking at practices that essentially involve *both* intentional states and linguistic performances. Neither sort of intentionality need be understood as conceptually prior to the other, and linguistic practice and rational agency can be presented as two aspects of one complex of jointly content-conferring practices. Davidson puts the characteristic contention of relational linguistic views of intentionality this way: "Neither language nor thinking can be fully explained in terms of the other, and neither has conceptual priority. The two are, indeed, linked, in the sense that each requires the other in order to be understood; but the linkage is not so complete that either suffices, even when reasonably reinforced, to explicate the other."[13] The account of intentionality introduced in the rest of this chapter is a linguistic theory in this relational sense. The key to motivating a theory of this sort is to show what it is about the contents of intentional states that can be explained only by appealing to the relation between such states and specifically *linguistic* performances.

Davidson suggests that an argument for a relational theory can be provided in two pieces. He claims first that "someone cannot have a belief unless he understands the possibility of being mistaken, and this requires grasping the contrast between truth and error—true belief and false belief."[14] What a

creature has does not function as a belief for that creature unless it has a certain kind of significance for that creature. It must be able to adopt a certain kind of practical attitude toward that state, to treat it in its practice or behavior *as* contentful in a special sense. In particular, Davidson is claiming, what it has is not recognizable as belief unless the creature whose state it is somehow in its practice acknowledges the applicability of a distinction between beliefs that are correct and those that are incorrect, in the sense of being true and false. We are not permitted to attribute the belief that *p* (a propositionally contentful intentional state) unless somehow the putative believer acknowledges in practice the *objective representational* dimension of its content—that its being held is one thing, but its being *correct* is another, something to be settled by how it is with what it is *about*. The second piece of Davidson's argument is the claim that a grasp of the contrast between correct and incorrect belief, true and false belief, "can emerge only in the context of interpretation, which alone forces us to the idea of an objective, public truth." The key claim is that "the concepts of objective truth and of error necessarily emerge in the context of interpretation."

The rest of this work focuses on the development of an account of linguistic social practices within which states, attitudes, and performances have, and are acknowledged by the practitioners to have, pragmatic significances sufficient to confer on them objective representational propositional contents. The view propounded is like Davidson's in seeing intentional states and speech acts as fundamentally of coeval conceptual status, neither being explicable except in an account that includes the other. It deserves nonetheless to be called a *linguistic* view of intentionality (of the relational rather than the analogical variety) because linguistic practice is nonetheless accorded a certain kind of explanatory priority over rational agency. The intentionality of nonlinguistic creatures is presented as dependent on, and in a specific sense derivative from, that of their linguistically qualified interpreters, who as a community exhibit a nonderivative, original intentionality. The sort of derivation in question is explicated in terms of the context that must be appealed to in making intelligible the sort of contents (*conceptual*—paradigmatically *propositional*—contents) that are associated with the intentional states attributed by interpreters. The contents of the intentional states attributed to nonlinguistic creatures can be understood only in a way that involves the activities of the language users who attribute them, and not entirely in terms of the activities of those who exhibit them. By contrast, the contents of the intentional states attributed to a community of language users can be understood as conferred on their states, attitudes, and performances entirely by the practices of those community members.

The argument that provides the ultimate justification for treating specifically *linguistic* practice as central in this way to intentionality has just the two-part form outlined by Davidson and rehearsed above. For what he

has really given us is not so much an argument as the form of one. Turning it into an actual argument requires filling in various notions of content, of objective representational correctness of content, of practical acknowledgment of the significance of assessments of correctness of content, and so on. That is the task of the rest of this work; the final justification for giving pride of place to language will thus not be complete until the end (in fact, in Chapter 8). At that point it will be possible to return again to the beginning, and know it for the first time—the warrant for this fundamental theoretical commitment will then be explicit.

6. Believing and Claiming

Claiming and believing are linked by the principle that assertions are one way of expressing beliefs. A fundamental question of explanatory strategy is then whether this principle can be exploited so as to account for one of these notions in terms of the other. Since there can be beliefs that are not avowed, the temptation is to start with belief and explain assertion as a speech act by which belief is expressed. But when the representational dimension of propositional and other conceptual contents is examined (in Chapter 8), it turns out to be intelligible only in the context of *linguistic* social practices of communicating by giving and asking for reasons in the form of claims. So if assertion were to be explained in terms of a prior notion of belief, the propositional contents of beliefs would have to be taken for granted. Their association with beliefs would have to be stipulated, rather than made intelligible as established by the functional role of beliefs in the behavioral economies of believers. The only sort of inferential practice that is socially articulated in the way that turns out to be required for the conferral of *propositional* content, in the form of objective truth conditions, is *assertional*, and therefore *linguistic* practice.

The idea pursued here is that the state or status of *believing* is essentially, and not merely accidentally, related to the linguistic performance of *claiming*. Beliefs are essentially the sort of thing that can be expressed by making an assertion. Dummett offers a crisp formulation of a view along these lines: "We have opposed throughout the view of assertion as the expression of an interior act of judgment; judgment, rather, is the interiorization of the external act of assertion."[15] Although the satisfyingly symmetric phrasing of this remark can obscure the point, it is important to realize that Dummett is not committed by it to the possibility of making sense of the speech act of asserting without mentioning anything but speech acts. For instance, this stance does not preclude an account of asserting that incorporates an account of the particular sort of *commitment* (a deontic status) one undertakes in making an assertion. What is precluded is only explaining assertion as the expression of a kind of intentional state or deontic status that is supposed to

be intelligible apart from the possibility of expressing it by asserting something.

The claim is, then, that speech acts having the pragmatic significance of assertions play an essential role in (social) functional systems within which states or statuses can be understood as propositionally contentful in the way beliefs are. A good way to see how explanatory progress might be made in thinking about beliefs by insisting on their linguistic expression in claims is provided by combining Dummett's thought with a suggestion of Hartry Field's concerning how an appeal to language might function as part of a divide-and-conquer strategy for explaining intentional states. His idea is that having a belief with a certain propositional content should be understood as standing in a certain relation to a sentence that expresses that content.[16] In particular, according to what he calls the "two-stage" explanatory strategy, A believes that p if and only if there is a sentence σ such that:

1. A believes* σ, and
2. σ means that p.

In the first stage of an account with this structure, the relevant relation between a believer and a sentence—what Field calls "belief*"—must be explained. In the second stage, what it is for a sentence to express a propositional content must be explained. Field's own way of pursuing this strategy takes the sentences involved to belong to a "language of thought," which is conceived by analogy to public languages. This is an additional theoretical commitment on his part; nothing about the two-stage decomposition dictates that the sentences that play the role of middle terms should not be sentences in a public language such as English. Filling in the appeal to sentences by invoking a language used in interpersonal communication makes the two-stage strategy available for duty in what were called above relational linguistic theories of intentionality.

In this form, Dummett's claim about the fundamental importance of the speech act of assertion can be pressed into service in addressing the first subproblem of the two-stage strategy. One way in which beliefs are manifested or expressed is by the utterance of sentences. Sentence-utterings can have many sorts of force or pragmatic significance, but when such performances have the significance of *assertions*, they express or purport to express beliefs. As Dummett suggests, this fact can be exploited by two different orders of explanation. If the theorist can get an independent grip on the notion of belief, typically from the consideration of its functional role in the sort of practical reasoning implicitly attributed by intentional interpretations that use the model of rational agency to make nonlinguistic performances instrumentally intelligible, then asserting might be explained in terms of it, as a further way in which beliefs can be manifested. Or, if the theorist can get an independent grip on the practices within which performances are accorded the significance characteristic of assertions, belief might be ex-

plained as what is thereby expressed. If this path is followed, then only parties to linguistic practices, those that institute assertional significances, will qualify as believers. The decision as to which direction of explanation to adopt is in part an issue concerning just how important language is to us.

Like all fundamental demarcational matters, however, it is only in part a factual issue. There clearly is a sense in which nonlinguistic animals can be said to have beliefs. But the sense of belief that Sellars, Dummett, and Davidson are interested in (and that is the subject of this work) is one in which beliefs can be attributed only to language users. The best reason for adopting the contrary order of explanation, for treating the sort of nonlinguistic belief that is manifested in behavior that can be construed instrumentally as fundamental—seeking thereby to explain the sort of belief that is essentially and not merely accidentally expressible in speech acts—is that it is clear that there were nonlinguistic animals before there were linguistic ones, and the latter did not arise by magic. If the instrumental sense of belief could be made sense of antecedently, and the linguistic sense explained in terms of it, the prospects for explaining how linguistic practice could come into the world in the first place would be bright. This is a laudable aspiration, and it may seem perverse to spurn it. Yet it is a consequence of the account of propositional contents to be offered here that they can be made sense of only in the context of *linguistic* social practices, which have as their core the interpersonal communication of information by assertions. Likewise, rational agency, on which instrumental behavior is modeled, depends essentially on specifically linguistic practices, including asserting. It follows that simple, nonlinguistic, instrumental intentionality can *not* be made fully intelligible apart from consideration of the linguistic practices that make available to the interpreter (but *not* to the interpreted animal) a grasp of the propositional contents attributed in such intentional interpretations.

A theory with such a consequence obviously involves a collateral commitment to show that the conceptual priority of the linguistic sense of belief need not make mysterious the advent of linguistic practices from the capacities of hitherto nonlinguistic creatures. The story to be told here assumes only that suitable social creatures can learn to distinguish in their practice between performances that are treated as correct by their fellows (itself a responsive discrimination) and those that are not. In accord with the pride of place being granted to the linguistic sense of belief, no appeal will be made to instrumental rationality on the part of fledgling linguistic practitioners. The primary explanatory target is what it is to grasp a propositional content and to have and to attribute to others states and performances with such contents—in the sense of explaining what doing the trick consists in, what would count as doing it, rather than how it is done by creatures wired up as we are. Thus no attempt will be made to show how the linguistic enterprise might have gotten off the ground in the first place. But it should be clear at each stage in the account that the abilities attributed to linguistic practition-

ers are not magical, mysterious, or extraordinary. They are compounded out of reliable dispositions to respond differentially to linguistic and nonlinguistic stimuli. Nothing more is required to get into the game of giving and asking for reasons—though to say this is not to say that an interpretation of a community as engaged in such practices can be paraphrased in a vocabulary that is limited to descriptions of such dispositions. Norms are not just regularities, though to be properly understood as subject to them, and even as instituting them by one's conduct (along with that of one's fellows), no more need be required than a capacity to conform to regularities.

If the strategy Field proposes is pursued by looking at the use of the sentences of a public language to perform communicative speech acts, the two subproblems into which he divides the problem of how to understand the attribution of intentionally contentful states correspond to a fundamental pragmatic question and a fundamental semantic question. The first concerns what it is for the utterance of a linguistic expression to have the pragmatic significance of an assertion. This can be rephrased as the question of what it is to use a sentence to make an assertion, provided it is remembered both (1) that it cannot be assumed that *sentences* can be distinguished from other linguistic expressions in advance of saying what it is to use an expression to make an assertion and (2) that *use* should not be assumed to involve a deliberate instrumental exercise of rational agency. Chapter 1 recommended a broadly *normative* approach to the pragmatic question; Chapter 2 recommended a broadly *inferential* approach to the semantic question. It is the task of the rest of this chapter to weave these approaches together into an account of discursive practice—comprising implicitly normative, inferentially articulated statuses, attitudes, and performances. It is the role they play in discursive practice that confers on them objectively representational content, in the most basic case objectively representational *propositional* content. The capacity of practice to confer such contents depends essentially, it will be argued, on its being not only *social* practice but *linguistic* social practice, in that at its core is *communication*, specifically by practitioners' interpretation of each others' *assertions*.

Addressing the pragmatic limb of the two-stage explanatory strategy by appealing to the speech act of assertion yields a further subdivision of issues. As just indicated, it requires an account of what it is for a performance to have the force or pragmatic significance of an assertion, for it to function as an assertion in the game of giving and asking for reasons. Such an account, however, would not by itself yield a suitable reading of Field's *belief** relation between a sentence and a potential subject of intentional interpretation or attribution of intentional states. For however tightly the two might be linked, there is a substantial difference between believing that p and claiming that p. (Commitment to a suitable resolution of the semantic subproblem entails a commitment to the eventual appropriateness of this sort of paraphrase in terms of propositional contents rather than sentences.) No sort of intentional

state (or normative status) that might be reconstructed in terms of assertion will provide a suitable analog for belief unless it preserves this contrast by leaving room for the possibility of being in the relevant intentional state without producing the corresponding assertional performance, and for producing the assertional performance without being in the corresponding intentional state. So besides an account of what the assertional significance of a speech act consists in, an account is required also of how attribution of such a speech act is to be understood to be related to attribution of the intentional state it expresses. Such an account might appeal to dispositions—for instance treating being in the state as being *disposed*, under appropriate conditions, to perform the speech act. Or it might appeal to norms—for instance treating the performance of the speech act as involving a *commitment* (which might or might not be fulfilled) involving the state. Or the account might involve both.

II. DEONTIC STATUS AND DEONTIC ATTITUDES

1. Doxastic Commitments

The leading idea of the account to be presented here is that belief can be modeled on the kind of inferentially articulated *commitment* that is undertaken or acknowledged by making an assertion. These may be called *doxastic* or *assertional* commitments. This is the basic kind of *discursive* commitment. The strategy is to describe a simplified system of social practices in which something can be taken or treated *as* (having the significance of) an assertion—the acknowledging of commitment to an assertible content. "Assertible content" is what Frege's "judgeable content"[17] becomes from the point of view of an explanatory commitment to understand judgments in the first instance as what is expressed by assertions. Specifically *propositional* contents (believables) are accordingly to be picked out by the pragmatic property of being assertible. Likewise, what is uttered or inscribed in producing an assertional performance is thereby recognizable as a declarative *sentence*. The role of *propositional* contents marks off *discursive* practice, and the role of *sentential* expressions of such contents is distinctive of *linguistic* social practice. In this way, everything comes down to being able to say what it is for what practitioners are doing to deserve to count as adopting a practical attitude of acknowledging the *assertional* significance of a performance: taking or treating it *as* an assertion. It is in terms of such attitudes that the pragmatic significance of assertional speech acts, the normative status of assertional commitments, and the possession or expression of propositional semantic contents are to be understood.

That the contents conferred by those practices are recognizable as *discursive* or *conceptual* contents (the genus of which propositional contents form the most basic species) depends on their *inferential* articulation and relation

to each other. The practices that institute the significance characteristic of assertional performances and the status characteristic of assertional commitments must be inferential practices. Asserting cannot be understood apart from inferring. So one fundamental question is, What makes something that is done according to a practice—for instance the production of a performance or the acquisition of a status—deserve to count as *inferring?* The answer developed here is that inferring is to be distinguished as a certain kind of move in the game of giving and asking for reasons. To say this is to say that inferring should be understood in an interpersonal context, as an aspect of an essentially *social* practice of communication.

The contentfulness of the states attributed as part of a simple intentional interpretation of an individual consists in a sort of inferential articulation that is not intelligible solely in terms of the role those states play in practical reasoning—if practical reasoning is conceived of as restricted to the sort of intrapersonal instrumental deliberation implicitly imputed by such interpretation. The explanation of behavior according to the model of rational agency depends on treating attributed intentional states as having propositional contents, which involve objective truth conditions. But there is, it will be argued, no pattern of moves a single individual might make that would qualify that individual's states as inferentially articulated in this sense. The inferential practice (including practical reasoning) that confers contents of this kind comprises not only first-person reasoning but also third-person attributions and assessments of it—and both aspects are essential to it. Deliberation is the internalization of the interpersonal, communicative practice of giving reasons to and asking reasons of others, just as judgment is the internalization of a public process of assertion. Inferring cannot be understood apart from asserting. To say this is to say that inferring should be understood as an aspect of an essentially *linguistic* practice. The practice of giving and asking for reasons must be conceived as including assertion because, although there are other kinds of performances besides assertings that can stand in need of reasons—indeed for a performance to be an action just is for it to be something it is appropriate to demand a reason for—to offer a reason is always to make an assertion.

The rest of this chapter is given over to developing a model of assertional and inferential practice. What is described is not our actual practice but an artificial idealization of it. Simplified and schematic though the model may be, it should nonetheless be recognizable as a version of what we do. The model is intended to serve as the core of a layered account of linguistic practice. Where our practice diverges from that specified in the model, those divergences should be explicable as late-coming additions to or modifications of the underlying practice. For instance, the model appeals only to semantic inferences, that is, inferences involving what is claim*ed.* Pragmatic inferences such as Gricean implicatures have to do rather with the antecedents and consequents of the performance of claim*ing* it. These pragmatic inferen-

tial practices form a shell around the more basic semantic ones, which they presuppose. The critical criterion of adequacy the model answers to is that the core linguistic practices it specifies be *sufficient* to confer propositional and other conceptual contents on the expressions, performances, and deontic statuses that play appropriate roles in those practices. It is also claimed, however, that the fundamental structural features of the model provide *necessary* conditions for the conferral of such contents. So there is a subsidiary commitment to the effect that sophisticated linguistic practices of the sort not addressed by the model are ultimately intelligible only by showing how they could develop out of the sort of practices the model does specify.

2. Commitment and Entitlement

At the core of discursive practice is the game of giving and asking for reasons. Chapter 1 sought to motivate the claim that discursive practice is implicitly *normative;* it essentially includes assessments of moves as correct or incorrect, appropriate or inappropriate. The institution of these proprieties by practical assessments on the part of the practitioners is the ultimate source of the meanings of the noises and marks they make, and of the other things they do.[18] As the term is used here, to talk of practices is to talk of *proprieties* of performance, rather than of regularities; it is to prescribe rather than describe. The general notion of proprieties of practice in terms of which the discussion of implicit norms has been conducted up to this point, however, does not cut fine enough to pick out what is distinctive of discursive norms. For that purpose the pragmatics Dummett suggests—which specifies the significance of linguistic expressions (and implicitly of speech acts and alterations of intentional states) in terms of circumstances of *appropriate* application and *appropriate* consequences of such application—must be further refined. Different sorts of propriety must be acknowledged.

The fundamental normative concept required is the notion of *commitment.* Being committed is a normative status—more specifically a *deontic* status. The project of the central sections of this chapter is to introduce a notion of *discursive commitment* as a species of deontic status that can do much of the explanatory theoretical work that is normally assigned to the notion of intentional state. But deontic statuses come in two flavors. Coordinate with the notion of *commitment* is that of *entitlement.* Doing what one is committed to do is appropriate in one sense, while doing what one is entitled to do is appropriate in another. The model of linguistic practice described here elaborates on the Dummettian bipartite pragmatics by distinguishing on the side of consequences, for instance, what a particular speech act commits one to from what it entitles one to. This permits a finer-grained specification of functional roles in linguistic practice than does using a single-sorted notion of propriety of circumstances and consequences of performance.

Commitment and entitlement correspond to the traditional deontic primitives of obligation and permission. Those traditional terms are avoided here because of the stigmata they contain betraying their origin in a picture of norms as resulting exclusively from the commands or edicts of a superior, who lays an obligation on or offers permission to a subordinate. Framed this way, the question of what one is obliged or permitted to do can slip insensibly into the question of who has a right to impose those statuses (as it does explicitly for Pufendorf, for instance). The picture presented here does not depend on a hierarchy of authority. The concepts of obligation and permission, as of duties and rights, can be reconstructed in terms of commitment and entitlement as they will be construed here.

Another way in which the treatment here of the deontic primitives of commitment and entitlement differs from that usually accorded to obligation and permission concerns the relation between them. It has been traditional to acknowledge the relations between these deontic modalities by defining one in terms of another: being permitted to do something is to be rendered as not being obliged not to do it, or being obliged as not being permitted not to. It does make sense to think of being committed to do something as not being entitled not to do it, but within the order of explanation pursued here it would be a fundamental mistake to try to exploit this relation to define one deontic status in terms of the other. Doing so requires taking a formal notion of negation for granted. The strategy employed here is rather to use the relation between commitment and entitlement (which are not defined in terms of this relation) to get a grip on a *material* notion of negation, or better, incompatibility. Two claims are incompatible with each other if commitment to one precludes entitlement to the other. One of the prime advantages normative-functional analyses of the notion of intentional states have over causal-functional analyses is that rendering the phenomenon of 'inconsistent' beliefs[19] as incompatible commitments makes it intelligible in a way not available to causal accounts.

It may also be remarked, in a preliminary fashion, that supposing that sense can be made of the underlying deontic statuses of commitment and entitlement, the notion of material incompatibility of commitments they give rise to leads in a straightforward way to a notion of the *contents* of such commitments. For the content of a commitment can for many purposes be represented by the theorist as the set of commitments that are incompatible with it. For instance, a kind of entailment relation is induced on commitments by inclusion relations among such sets of incompatibles. The commitment p incompatibility-entails the commitment q just in case everything incompatible with q is incompatible with p. Thus "Wulf is a dog" incompatibility-entails "Wulf is a mammal," since everything incompatible with his being a mammal is incompatible with his being a dog. The notion of material incompatibility that is made available by *not* defining commitment and entitlement in terms of one another accordingly provides a route from

the pragmatics that deals with these deontic statuses to the semantics that deals with their contents. How this hint might be exploited will become clearer presently.

3. Attributing and Undertaking Commitments

Deontic statuses of the sort to be considered here are creatures of practical attitudes. There were no commitments before people started treating each other as committed; they are not part of the natural furniture of the world. Rather they are social statuses, instituted by individuals attributing such statuses to each other, recognizing or acknowledging those statuses. Considered purely as a natural occurrence, the signing of a contract is just the motion of a hand and the deposition of ink on paper. It is the undertaking of a commitment only because of the significance that performance is taken to have by those who attribute or acknowledge such a commitment, by those who take or treat that performance as committing the signatories to further performances of various kinds. Similarly for entitlements. A license, such as a ticket, entitles one to do something. Apart from practices of treating people as entitled or not, though, there is just what is actually done. The natural world we consider when bracketing the influence of such social practices contains no distinction between performances one was entitled to and those one was not.

It will be useful to see how this basic vocabulary can be used to discuss the authority and responsibility involved in familiar sorts of deontic statuses. The way in which such statuses can be instituted by practical attitudes can be illustrated by artificially simplified versions of some fundamental normative practices. Authority may be considered first, apart from responsibility. A license, invitation, or entrance ticket entitles or authorizes one to do something one was otherwise not entitled to do. It is always a license in the eyes of someone, for example a ticket-taker or doorman. These "consumers" of licenses (along with the others whose attitudes make the practice into a going concern) constitute them by attributing the authority they thereby come to possess. They do so by treating the authorized one as entitled to a performance. It is not appropriate to enter unless one is authorized by a ticket acquired in the appropriate way. Being given a ticket by the ticket-giver is being authorized or entitled to enter, because and insofar as the ticket-taker will not treat entry as appropriate unless so authorized. In the simplest case, the ticket-taker is the attributor of authority, the one who recognizes or acknowledges it and who by taking the ticket as authorizing, makes it authorizing, so instituting the entitlement.

Practices of this sort can be described in purely responsive terms for prelinguistic communities. The entitlement given and recognized in these practices has a content for an attributor insofar as that attributor practically partitions the space of possible performances into those that have been

authorized and those that have not, by being disposed to respond differently in the two cases. These sanctioning responses (for instance admitting versus ejecting) and the performances they discriminate (enterings of the theater) can be characterized apart from and antecedent to specification of the practice of conferring and recognizing entitlement defined by their means. For this reason the entitling authority will be said to be *externally* defined. The sanctions applied in taking or treating someone as entitled can be specified in nonnormative terms.

The basic structure just considered involves entitlement without commitment, authority without responsibility. A corresponding way into the basic structure of commitments or responsibilities is provided by describing a similarly simplified and artificial version of an actual eighteenth-century British practice. According to this practice, taking "the queen's shilling" from a recruiting officer counts as committing the recipient to military service. A performance of this kind has the same significance that signing a contract would have—in either case one has joined the army and undertaken all the commitments entailed by that change of status. (The official rationale was that some such overt irrevocable nonlinguistic performance was required to do duty for signing a contract, given that those enlisting were largely illiterate. The actual function of the practice was to enable "recruiting" by disguised officers, who frequented taverns and offered what was, unbeknownst to their victims, the queen's shilling, as a gesture of goodwill to those who had drunk up all of their own money. Those who accepted found out the significance of what they had done—the commitment they had undertaken, and so the alteration of their status—only upon awakening from the resulting stupor.) The significance of a commitment is to be understood in terms of the practical attitude of those attributing it, that is in terms of what taking or treating someone *as* committed consists in. In this case, it consists in being liable for punishment by a court-martial if one fails to discharge or fulfill the commitment. The content of the commitment attributed corresponds to the subsequent behavior that would or would not elicit a sanction. Or more precisely, that content can be thought of according to Dummett's model of circumstances and consequences; the particular sanctions (court-martial) are the consequences, and the various failures to perform as a soldier that elicit them are the circumstances.

Two features of this simple commitment structure are worthy of note. First, for someone to *undertake* a commitment, according to this story, is to do something that makes it appropriate to *attribute* the commitment to that individual. That the performance of accepting the coin has the significance of altering the *status* of the one whose performance it is consists in the change it brings about in what *attitudes* are in order. It is by reference to the attitudes of others toward the deontic status (attributing a commitment) that the attitude of the one whose status is in question (acknowledging or undertaking a commitment) is to be understood. So all that is required to make

sense of the normative significance of the performance as an undertaking of commitment is an account of what it is to take or treat someone as committed to do something. The possibility of sanctioning failure to perform appropriately—that is, as one is (thereby) taken to be committed to do—offers a way of construing this fundamental practical deontic attitude.

Second, the basic notion of responsibility or commitment that is introduced by consideration of this simple practice can be understood in terms of the notion of authority or entitlement already discussed. For undertaking a commitment can be understood as authorizing, licensing, or entitling those who attribute that commitment to sanction nonperformance. Such sanctioning would be inappropriate (and so itself subject to sanction) unless it had been authorized by the undertaking of a commitment. One may not court-martial someone who has never joined the service. Thought of this way, the effect of undertaking a commitment is not a matter of in fact eliciting punishment if one does not fulfill the commitment but rather of making such punishment *appropriate.* It is not a matter of the actual conditional dispositions to sanction of those who attribute the commitment but a matter of the conditional normative status of such sanctions. What is being considered is a slightly more sophisticated practice, in which the significance of taking the queen's shilling (the consequences of undertaking a commitment and thereby acquiring a new deontic status) is itself defined in terms of deontic statuses— in particular, of *entitlement* to punish.

The significance of that deontic status (entitlement to punish) might itself be defined in normatively *external* terms; those who attempt to court-martial someone who has not committed himself to the service (and so entitled their superiors to hold such a court) are taken out and shot. Or the difference that entitlement makes might itself be cashed out only in normative terms; those who attempt to court-martial someone who has not committed himself to the service (and so entitled their superiors to hold such a court) thereby *license* or *entitle* their superiors to have them taken out and shot, make it *appropriate* or *correct* for them to be taken out and shot. Normatively *internal* definitions of the significance of changes of deontic status, which specify the consequences of such changes in terms of further changes of deontic status, link various statuses and attitudes into *systems* of interdefined practices.

4. The Example of Promising

These points can be illustrated by considering an idealized version of a more familiar sort of practice of undertaking commitments, namely *promising.* Promising is another way of undertaking a task responsibility: committing oneself to perform in a certain way. A special defining feature of promising performances is that they involve specifying what one is committed to do by explicitly *saying* it—*describing* the performances that would

count as fulfilling the commitment. This feature of promising will not be officially intelligible in terms of the theory presented here until the practice of assertion has been introduced. For the promise offers a linguistic characterization of a performance, and that characterization (or a grammatical transform of it) must turn out to be assertible, on pain of the promisor's having failed to fulfill the commitment undertaken. Only one who claims that MacArthur returned will take it that MacArthur's promise to return was fulfilled. Furthermore, promises are typically made *to* someone. The promisees are the ones who are entitled to hold the promiser responsible.

A social-practical description of promising displays how the deontic attitudes of undertaking and attributing the commitment that is a promise are two sides of one coin, aspects of a single social practice. Promising (like taking the queen's shilling) can be understood as consisting in a dimension of responsibility and a coordinate dimension of authority. The authority of the licensing, in the eyes of the attributors of the promise-commitment, is an entitlement on the part of others to rely (as perhaps they were not previously entitled to do) upon the promised performance. The responsibility consists in conditionally entitling others to sanction the committed one. A commitment to perform includes a license to those who attribute the commitment to hold the committed one responsible for nonperformance. To be entitled to hold responsible is to be conditionally entitled to sanction, in case of nonperformance. In the simplest sort of promising practice, both the promised performances and the conditionally authorized sanctions can be specified antecedently (that is, defined externally to the definition of the practice that then ties them together). The promised performance might be washing the dinner dishes, and the conventionally authorized sanction for nonperformance might be being beaten with sticks.

Commitments to perform that are externally defined by practical attitudes manifested in fulfilling performances, and sanctions that are antecedently specifiable in nonnormative terms, share some features with assertional commitments. As already indicated, however, this category must be developed in the direction of a system of practices governing interlocking internally specified significances, defined in terms of fulfilling performances and sanctions. The toy promising practice just mentioned can be developed further by introducing an internally defined sanction for failure to perform as promised. Instead of responding to the failed promiser by beatings or refusals of entry to the theater, attributors of commitments to perform might rather withhold their recognition of that individual's entitlement to undertake such commitments. If the commitment is not fulfilled, the cost is that those who attribute both the commitment and the failure will not (or will become less disposed to) recognize the authorization (making entitled) of reliance on performances, that is, will not recognize the promises of the failed commitment undertaker as counting for anything. Undertaking the commitment is still doing something that conventionally has the significance of entitling others to attribute the commitment, that is both to attribute the commit-

ment and to hold the undertaking performer to it. One who succeeds in making a promise still authorizes others (makes them entitled) to rely on one's future performance, to hold one responsible for a failure to perform according to one's commitment.

These are both authority and responsibility, adding up to commitment. The commitment is undertaken by the one who authorizes or licenses, and undertaking is always licensing the attribution of—that is the holding of the undertaker responsible for—the commitment. Responsibility is of author-izer-undertaker to attributor. The difference is that in the sort of promising being considered, uttering certain words is not always sufficient to undertake the commitment that defines making a promise. For as a sanction consequen-tial on previous failures to keep commitments undertaken by promises, attributors may withhold their recognition of the undertaker's entitlement to issue licenses of that sort, refuse to recognize as successful attempts to undertake such commitments. Unlike the queen's shilling case, no perfor-mance compels attribution of authority and responsibility—that is, commit-ment. Such a sanction is defined internally to the practice being considered, since apart from the practice of promising, one cannot specify what the sanction is.[20]

Part of the definition of what it is to undertake a commitment to perform by promising is the significance of this status for the deontic attitudes of others—the practical interpretation that attributors are entitled to place on the undertaking—namely their right to rely on future performance. The second sort of responsibility undertaken is not a task responsibility (as the commitment to perform could be called). It is a becoming liable to be held responsible (taken to be responsible) for failure to perform as one promised. What the promiser entitles others to do, in this simple practice, is to with-draw their recognition (taking) of one's entitlement to issue a license, of one's authority. When responsibility of this sort is added to authority of the invi-tation or ticket-giving variety, besides the conditional task responsibility, there must also be the special sort of authority involved in undertaking responsibility. This is authorizing others to hold the undertaker responsible. It is a license to do something, conditional upon the undertaker's failure to fulfill his or her task responsibility once its conditions are satisfied. A dis-tinct and important species of such authority to hold responsible arises when the licensed consequences of failure to perform consist in withholding rec-ognition of entitlement to undertake further responsibilities of the same form. Both promising, as reconstructed here, and asserting, as discussed below, exhibit this special commitment-entitlement structure.

5. Social Practice: Deontic Statuses and Deontic Attitudes

The discussion of these simplified examples introduced a number of points. The sort of implicitly normative practice of which language use is a paradigm is to be discussed in terms of two sorts of *deontic status*, namely

commitment and entitlement. The notion of normative status, and of the significance of performances that alter normative status, is in turn to be understood in terms of the practical deontic *attitude* of taking or treating someone as committed or entitled. This is in the first instance *attributing* a commitment or entitlement. Adopting this practical attitude can be explained, to begin with, as consisting in the disposition or willingness to impose *sanctions.* (Later, in more sophisticated practices, *entitlement* to such a response, or its *propriety,* is at issue.) Attributors of these statuses may punish those who act in ways they are not (taken to be) entitled to act, and those who do not act in ways they are (taken to be) committed to act. What counts as punishment may (according to the one who interprets a community as exhibiting practices of this sort) be specifiable in nonnormative terms, such as causing pain or otherwise negatively reinforcing the punished behavior. Or what counts as punishment with respect to a particular practice may be specifiable only in normative terms, by appeal to alterations in deontic status or attitude. A performance expresses the practical attitude or has the significance of an *undertaking* of a commitment in case it *entitles* others to attribute that commitment. So there are two sorts of practical deontic attitudes that can be adopted toward commitments: attributing them (to others) and acknowledging or undertaking them (oneself). Of these, attributing is fundamental.

Here, then, is a way of thinking about implicitly normative social practices. Social practices are games in which each participant exhibits various deontic statuses—that is, commitments and entitlements—and each practically significant performance alters those statuses in some way. The significance of the performance is how it alters the deontic statuses of the practitioners. Looking at the practices a little more closely involves cashing out the talk of deontic *statuses* by translating it into talk of deontic *attitudes.* Practitioners take or treat themselves and others as having various commitments and entitlements. They *keep score* on deontic statuses by attributing those statuses to others and undertaking them themselves. The significance of a performance is the difference it makes in the deontic score—that is, the way in which it changes what commitments and entitlements the practitioners, including the performer, attribute to each other and acquire, acknowledge, or undertake themselves. The significance of taking the queen's shilling lies in its being an undertaking of a commitment on the part of the recipient, altering the attributions of commitment by those who appreciate the significance of the performance. It entitles other authorities—those who according to the antecedent score already had undertaken various commitments or duties and entitlements or sorts of authority, those who therefore play a certain role or hold a certain office in the system of practices in question—to punish the performer in particular ways under particular circumstances. The normative significances of performances and the deontic states of performers are instituted by the practice that consists in keeping score by adopting attitudes of attributing and acknowledging them.[21]

III. ASSERTING AND INFERRING

1. Linguistic Practice: Assertion and Inference

The discussion of the significance of performances as altering the deontic attitudes that keep track of normative statuses has so far addressed implicitly normative social practices—whether or not they are specifically *linguistic* (and so more generally *discursive*) practices. What distinguishes the latter sort is the *inferential* articulation of the normative significances they involve—and so their conferral of specifically *conceptual* content on the states, attitudes, performances, and expressions they govern. The challenge is to show how these two approaches (normative pragmatics modeled on deontic scorekeeping, and inferential semantics) can be combined into a single story about social practices of treating speech acts as having the significance of assertions.

Describing practices sufficient to institute such a significance is the way to fill in the notion of assertional commitment. Such an account provides an answer to the question, What is it that we are *doing* when we assert, claim, or declare something? The general answer is that we are undertaking a certain kind of *commitment.* Saying specifically *what* kind is explaining what structure must be exhibited by the practices a community is interpreted as engaging in for that interpretation to be recognizable as taking the practitioners to be keeping score for themselves and each other in virtue of the alterations of their practical deontic attitudes of attributing and undertaking assertional commitments and their corresponding entitlements.

The key to seeing how the scorekeeping model of deontic social practices can be used to make sense of asserting is Sellars's notion of a "game of giving and asking for reasons." The idea is that assertings (performances that are overt undertakings of assertional commitments) are in the fundamental case what reasons are asked for, and what giving a reason always consists in. The kind of commitment that a claim of the assertional sort is an expression of is something that can stand in need of (and so be liable to the demand for) a reason; and it is something that can be offered as a reason. This is the principle motivating the present strategy for discriminating assertional commitments from other species of commitment. Other things besides assertional commitments involve liability to demands for justification or other demonstration of entitlement—for instance, the practical commitments involved in *actions.* Other things besides assertional commitments can entitle interlocutors to assertional commitments—for instance *reliability* in the responsive acquisition of assertional commitments of a certain kind. For being a reliable reporter of currently visible red things who responsively acquires a disposition to claim that there is something red in the vicinity may entitle someone to that commitment.[22] But only assertional commitments stand in *both* these relations.

That claims play the dual role of justifier and subject of demand for justification is a necessary condition of their kind properly being called

assertional commitments. It is here employed as well as part of a sufficient condition, in an idealized artificial practice constructed to model this central aspect of the use of natural language. Specifically *linguistic* practices are distinguished as just the social practices according to which some performances have the significance of undertakings of assertional commitment (in virtue of their role in giving and asking for reasons); *declarative sentences* are picked out as the expressions uttered or inscribed in such assertional performances. What is expressed by such performances and determine the particular features of their significance within the assertional genus count thereby as *propositional* contents.

The idea exploited here, then, is that *assertions* are fundamentally fodder for *inferences.* Uttering a sentence with assertional force or significance is putting it forward *as* a potential reason. Asserting is giving reasons—not necessarily reasons addressed to some particular question or issue, or to a particular individual, but making claims whose availability as reasons for others is essential to their assertional force. Assertions are essentially *fit* to be reasons. The function of assertion is making sentences available for use as premises in inferences. For performances to play this role or have this significance requires that assertional endorsement of or commitment to something entitles or obliges one to other endorsements. The pragmatic significance of assertional commitments and entitlements to such commitments consists in the ways in which they are *heritable;* their heritability is the form taken by the *inferential* articulation in virtue of which they count as semantically contentful.

2. *Three Dimensions of Inferential Articulation*

The basic model of the inferential practices that institute assertional significance—and thereby confer propositional contents on states, attitudes, and performances playing suitable roles in those practices—is defined by a structure that must be understood in terms of the interaction of three different dimensions. First, there are two different sorts of deontic status involved: *commitments,* and *entitlements* to commitments. Inheritance of commitment (being committed to one claim as a consequence of commitment to another) is what will be called a *committive,* or commitment-preserving, inferential relation. Deductive, logically good inferences exploit relations of this genus. But so do materially good inferences, such as inferences of the form: *A* is to the West of *B,* so *B* is to the East of *A;* This monochromatic patch is green, so it is not red; Thunder now, so lightning earlier. Anyone committed to the premises of such inferences is committed thereby to the conclusions.

Inheritance of entitlement (being entitled to one claim as a consequence of entitlement to another) is what will be called a *permissive,* or entitlement-preserving inferential relation. Inductive empirical inferences exploit rela-

tions of this genus. The premises of these inferences entitle one to commitment to their conclusions (in the absence of countervailing evidence) but do not compel such commitment. For the possibility of entitlement to commitments incompatible with the conclusion is left open. In this way the claim that this is a dry, well-made match can serve as a justification entitling someone to the claim that it will light if struck. But the premise does not commit one to the conclusion, for it is compatible with that premise that the match is at such a low temperature that friction will not succeed in igniting it. The interplay between the two sorts of deontic status is at the center of the model of assertional and inferential practices presented here.

The broadly inferential roles that are identified with propositional contents involve not only commitment- and entitlement-preserving inferential connections among such contents but also relations of *incompatibility*. To say that two claims have (materially) incompatible contents is to say that commitment to one precludes entitlement to the other.[23] No candidate notion can count as a construal of the sort of propositional content we take assertions, judgments, and beliefs to have unless it underwrites incompatibility relations among them. (That possible worlds, for instance, must be understood as corresponding to maximal sets of *compatible* propositions is acknowledged both by those who want to exploit that principle to define propositions and their material compatibility in terms of possible worlds and by those who would reverse that order of explanation.) The explanatory strategy adopted here is to begin with practices that institute deontic statuses of commitment and entitlement and then to show how those practices thereby confer specifically propositional conceptual contents on what is assertible—contents recognizable as such in virtue of the deontic inferential and incompatibility relations they stand in.

The second dimension of broadly inferential articulation that is crucial to understanding assertional practice turns on the distinction between the *concomitant* and the *communicative* inheritance of deontic statuses. This is the *social* difference between *intra*personal and *inter*personal uses of a claim as a premise. Undertaking a commitment or acquiring an entitlement has consequences for the one whose statuses those are. One commitment carries with it other concomitant commitments as consequences. Its consequences are those that it entails according to the commitment-preserving inferential relations that its content stands in to other contents of possible commitments. Similarly, being entitled to a commitment can entitle one to others, which stand to it in suitable permissive or entitlement-preserving inferential relations. Again, the definition of incompatibility of contents in terms of commitment and entitlement means that acquiring a commitment may have as a consequence the loss of entitlement to concomitant commitments one was heretofore entitled to.

But these intrapersonal inferential consequences of changes in deontic status do not exhaust the significance of assertional performances. Such

performances also have a significance for interpersonal *communication*. Putting a sentence forward in the public arena *as* true is something *one* interlocutor can do to make that sentence available for *others* to use in making further assertions. Acknowledging the undertaking of an assertional commitment has the *social* consequence of licensing or entitling others to *attribute* that commitment. The adoption of that deontic attitude on the part of the audience in turn has consequences for the commitments the audience is entitled to undertake. Putting a claim forward *as* true is putting it forward as one that it is appropriate for others to *take* true, that is to endorse themselves. Assertion that is communicatively successful in the sense that what is put forward as true by a speaker is taken as true by the audience consists in the interpersonal inheritance of commitment.

The third dimension of broadly inferential articulation that is crucial to understanding assertional practice is that in which discursive *authority* is linked to and dependent upon a corresponding *responsibility.*[24] In uttering a sentence assertively, the claim one makes involves an *endorsement.* One aspect of this sort of endorsement was indicated above in a preliminary fashion in terms of the function of an asserting as licensing or *authorizing* further assertions (and eventually, actions—but consideration of practical rather than doxastic commitments is postponed until the next section). But unless some independent grasp is offered of the status or significance that must be bestowed on a performance for it to count as an asserting, invoking the inferential warranting of further assertions merely traces out a rather small circle. It is the second aspect of endorsement, of the sort of *responsibility* involved in assertional commitment, that permits a larger horizon. Understanding that aspect requires putting together the distinction between the deontic statuses of commitment and entitlement, on the one hand, and of intra- and interpersonal inference, on the other. Putting forward a sentence as true or as information—that is, asserting it—has been glossed as putting it forward as fit to be a reason for other assertions, making it available as a premise from which others can be inferred. This means that others can inherit *entitlement* to an assertional commitment from the one who makes an assertion and thereby licenses or warrants its reassertion and the assertion of what follows from it. To understand this warranting function, the heritability of entitlement, requires understanding the social significance of the distinction between warranted and unwarranted assertional commitments.

Ordinarily the relation of an authorizing event to the performances it licenses requires at least that in the context of that event, performances become socially appropriate that otherwise would not be. For example purchasing a ticket entitles one to take a seat in the theater, which it would be inappropriate to do without the ticket. This observation presents a dilemma. If asserting a sentence is not a performance requiring prior authorization, then it seems one cannot understand the function of assertion as inferentially licensing further assertions. If, however, asserting is a performance requiring

authorization, how does one become entitled to the original licensing assertion? Talk of *inheritance* of entitlement makes sense only in an explanatory context that includes a story about the significance of *possession* of entitlement. It is this question that is addressed by an account of the dimension of *responsibility* characteristic of asserting. In asserting a claim, one not only authorizes further assertions (for oneself and for others) but undertakes a responsibility, for one commits oneself to being able to vindicate the original claim by showing that one is entitled to make it. Others cannot inherit an entitlement that the asserter does not possess. Overtly acknowledging or undertaking a doxastic commitment by issuing an assertional performance can warrant further commitments, whether by the asserter or by the audience, only if that warranting commitment itself is one the asserter is entitled to. Only assertions one is entitled to make can serve to entitle anyone to their inferential consequences.

The function of asserting as the giving of reasons is intelligible only as part of a practice in which reasons can be asked for or required. That some performances admit or stand in need of reasons is presupposed by the practice of offering them. Many kinds of performances are subject to demands for or explanation according to reasons. The two fundamental sorts discussed here are intentional actions that are not speech acts and assertions themselves. Both actions and assertions—overt performances corresponding to practical and to doxastic commitments respectively—are essential and ineliminable aspects of discursive practice as here construed. Nonetheless, the significance of assertional performances can be filled in to a considerable extent before it is necessary to look at the role of assertions as reasons for anything other than more assertions.

The converse is not the case. Actions just are performances for which it is appropriate to offer reasons, and offering a reason is making an assertion. So actions are not intelligible as such except in a context that includes assertional giving of reasons. Where intentional explanations are offered of the behavior of nonlinguistic creatures (those that are not understood as interpreters of others), the reasons are offered, the assertions are made, by the interpreter of the simple intentional system, who seeks to make its behavior intelligible by treating it *as if* it could act according to reasons it offers itself. That is why what is attributed in such interpretations is *derivative* intentionality. Assertions play both roles; reasons can be offered for them, and they can be offered as reasons. Actions play only the first role in the game of giving and asking for reasons. Thus a description of that game that has a certain autonomy can be offered if, to begin with, attention is restricted to assertions alone, and as a result asserting has a certain explanatory priority over action. Although for this reason the discussion of actions and the practical commitments they express is postponed until the next chapter, it remains that asserting is a *doing,* and the responsibility it involves should also be understood as a responsibility to *do* something.

Besides its specifically linguistic use, 'assert' has a broader normative use according to which one can assert one's authority or one's rights. An important component of this sense of asserting is defending, championing, or justifying. Milton uses the word this way in the famous expression of his intent in *Paradise Lost:* "That to the highth of this great Argument I may assert Eternal Providence, and Justifie the wayes of God to men." This use points to the fact that the sort of commitment involved in linguistic asserting involves the undertaking of a specifically *justificatory* responsibility for what is claimed. In asserting a sentence, one not only licenses further assertions (for others and for oneself) but commits oneself to *justifying* the original claim. The responsibility in question is of the sort Baier calls "task-responsibility"; it requires the performance of a task of some kind for its fulfillment.[25] Specifically, in making a claim, one undertakes the conditional task responsibility to demonstrate one's entitlement to the claim, if that entitlement is brought into question. Justifying the claim when it is queried, giving reasons for it when reasons are asked for, is one way to discharge this obligation. If the commitment can be defended, entitlement to it demonstrated by justifying the claim, then endorsement of it can have genuine authority, an entitlement that can be inherited.

3. Assertion as a Doing: Authority and Responsibility

The position maintained here is that *discursive* (in the Kantian sense of concept-mongering) practice can only be *linguistic* practice, and that what distinguishes a practice as specifically linguistic is that within it some performances are accorded the significance of *assertions.* It is only because some performances function as assertions that others deserve to be distinguished as *speech* acts. The class of questions, for instance, is recognizable in virtue of its relation to possible answers, and offering an answer is making an assertion—not in every individual case, but the exceptions (for example, questions answered by orders or by other questions) are themselves intelligible only in terms of assertions. Orders or commands are not just performances that alter the boundaries of what is permissible or obligatory. They are performances that do so specifically by *saying* or *describing* what is and is not appropriate, and this sort of making explicit is parasitic on claiming. Saying "Shut the door!" counts as an order only in the context of a practice that includes judgments, and therefore assertions, that the door is shut or that it is not shut. (The "slab" *Sprachspiel* Wittgenstein describes in the opening paragraphs of the *Investigations* is not in this sense a *language* game—it is a set of practices that include only vocal, but not yet verbal, signals.) In the same way, promises are not just undertakings of responsibility to perform in a certain way. They are performances that undertake such responsibility by *saying* or *describing* explicitly what one undertakes to do. One promises in effect to make a proposition true, and the propositional

contents appealed to can be understood only in connection with practices of saying or describing, of taking-true—in short, of asserting what are, in virtue of the role they play in such assertions, declarative sentences. As it is with these examples, so it is with other speech acts. Asserting is the fundamental speech act, defining the specific difference between linguistic practice and social practices more generally.

A crucial measure according to which a theory of speech acts ought to be assessed, then, is its treatment of what one is *doing* in producing an assertion. This challenge is not always accepted. One prominent theorist defines the assertion of the declarative sentence p as "an undertaking to the effect that p."[26] One does not have to subscribe to the pragmatist project of explaining the propositional contents that are assert*ed* in terms of the practices of assert*ing* them in order to find this disappointing. What sort of an undertaking is this? What, exactly, is the effect? The theory being presented here aims to answer just these questions.

In producing assertions, performers are doing two sorts of things. They are first *authorizing* further assertions (and the commitments they express), both concomitant commitments on their part (inferential consequences) and claims on the part of their audience (communicational consequences). In doing so, they become responsible in the sense of answerable for their claims. That is, they are also *undertaking* a specific task *responsibility*, namely the responsibility to show that they are *entitled* to the commitment expressed by their assertions, should that entitlement be brought into question. This is the responsibility to *do* something, and it may be fulfilled for instance by issuing other assertions that *justify* the original claim. The semantic content of the commitment expressed by the performance—that the authority it claims and the justificatory responsibility it undertakes are specifically "to the effect that p" (rather than some other q)—consists in its specific inferential articulation: what else it commits the asserter to, what commitments it entitles its audience to, what would count as a justification of it, and so on. On this account, then, the pragmatic force or significance characteristic of asserting (and therefore ultimately also the concepts of declarative *sentence, propositional* content, and specifically *linguistic* practice) is to be understood in terms of performances with the dual function in the game of giving and asking for reasons of being givings of reasons, and themselves also performances for which reasons can be asked. The conceptual contents expressed by assertional performances are to be explained by appeal to the inferential roles they play in that reason-mongering practice. What is *done* in asserting— the pragmatic *significance* or *effect* of producing an assertional performance—consists in the way in which, by authorizing particular further inferentially related performances and undertaking responsibility to produce yet other inferentially related performances, asserters alter the score interlocutors keep of the deontic statuses (commitments and entitlements) of their fellow practitioners.

The analysis being suggested divides this significance into a component having to do with *authority* and one having to do with *responsibility*. The particular way these components are intertwined defines the sort of pragmatic significance that is being identified with assertional force. The constellation of authority and responsibility characteristic of the assertional significance of speech acts is *socially* articulated. In producing an assertion, one undertakes a responsibility oneself. The authority of that performance (which is conditioned on the responsibility) in turn consists in opening up a new avenue along which those in the audience can fulfill the responsibilities associated with *their* assertions. At the core of assertional practice lie three fundamental ways in which one can demonstrate one's entitlement to a claim and thereby fulfill the responsibility associated with making that claim. Two of these—justifying the content of an assertion and deferring to the authority of an asserter—can be discussed here. The third—invoking one's own authority as a reliable noninferential reporter—is best left for later.

First, as already mentioned, one can demonstrate one's entitlement to a claim by *justifying* it, that is, by giving reasons for it. Giving reasons for a claim always consists in making more claims: asserting premises from which the original claim follows as a conclusion.[27] Interlocutors who accept such a vindication of the commitment—who accept the reasons offered as a justification demonstrating entitlement to the conclusion—thereby implicitly endorse a certain inference. These practical attitudes of taking or treating inferences as correct (distinguishing them from the incorrect ones by responses to attempted justifications) institute inferential proprieties relating the performances of asserters (and the commitments adopted thereby) and so confer contents on them. For it is the practical inferential proprieties acknowledged by such attitudes that make noises and marks mean what they mean. Assertions play a dual role in justification: as justifiers and as justifieds, premises and conclusions. That it plays this dual role, that it is caught up in justificatory inferences both as premise and as conclusion, is what makes it a specifically *propositional* (= assertible, therefore believable) content at all. That it exhibits the particular inferential grounds and consequences that it does is what makes it the particular determinate content that it is—settling, for instance, what information it conveys, the significance that undertaking a commitment with that content would have for what else one is committed and entitled to. Thus the inferential articulation of speech acts depends on this practice of demonstrating entitlement to the commitment acknowledged by the performance of a speech act.

The second way of vindicating a commitment by demonstrating entitlement to it is to appeal to the authority of another asserter. The *communicational* function of assertions is to license others who hear the claim to reassert it. The significance of this license is that it makes available to those who rely on it and reassert the original claim a special way of discharging their responsibility to demonstrate their entitlement to it. They can invoke

the license or authority of the asserter, thus deferring to the interlocutor who communicated the claim and passing along to that other individual any demands for demonstration of entitlement. The authority of an assertion includes an offer to pick up the justificatory check for the reassertions of others. That A's assertion has the social significance of authorizing B's reassertion consists in the appropriateness of B's deferring to A the responsibility to demonstrate entitlement to the claim. B's responsibility can be discharged by the invocation of A's authority, upon which B exercised the right to rely. The buck is passed to A.

So communication does not involve only the sharing of *commitments*—their spreading from one individual to another as the speaker who produces an assertion communicates to, and possibly infects, an audience. It involves also the way that *entitlements* to claims can be inherited by the consumers of an assertion from its producer. In this way the *authority* of an assertional performance consists in part in making available a new way in which those to whom it is communicated can discharge their *responsibility* for demonstrating entitlement to commitments they undertake. Furthermore, assertions can be seen to play a dual role on the side of communication, just as they do on the side of justification. For assertions are on the one hand what is communicated (made available to others), and on the other hand they are what communication is *for:* one interlocutor's claim is fodder for inferences by others to further claims. The audience not only attributes to the one producing an assertional performance commitment to claims entailed (according to commitment-preserving inferences) by the assertion, but it also may undertake commitments and acquire entitlements that are its consequences.

In such inheritance of entitlement by communication, the content of the commitment is preserved intact and merely transferred from one scorekeeper to another. The communicational mechanism for fulfilling the responsibility to demonstrate entitlement appeals to *inter*personal, *intra*content inheritance of entitlement to a propositional commitment. By contrast, the justificatory mechanism appeals to *intra*personal, *inter*content[28] inheritance of entitlement to a propositional commitment—since the contents of premises and conclusions will differ in any inference that is nontrivial in the sense of being available to do justificatory work. This combination of the *person-based* authority (invoked by *deferring* to the claim of another) and *content-based* authority (invoked by *justifying* the claim through assertion of other sentences from which the claim to be vindicated can appropriately be inferred) is characteristic of asserting as a doing. This constellation—of commitment and entitlement, of authority and responsibility, and of an inheritance of entitlement to assertional commitments that is *inter*personal and *intra*content as well as *intra*personal and *inter*content—constitutes a fundamental substructure of the model of assertional practices presented here.

4. The Default and Challenge Structure of Entitlement

More clearly needs to be said about the practices that govern the attribution of entitlement to assertional commitments. The two mechanisms considered so far for demonstrating such entitlement are ways in which entitlement to commitments can be *inherited*. Entitlement to commitment to one claim can be extended to entitlement to another either according to the *inferential* pattern appealed to by justification, in which case that entitlement is inherited by another commitment (to a different claim) undertaken by the same interlocutor, or according to the *communicational* pattern appealed to by deferral, in which case that entitlement is inherited by another commitment (to the same claim) by a different interlocutor. Tracing back an entitlement secured by either of these sorts of inheritance potentially sets off a regress.

The justificatory style of vindication, in which one interlocutor offers premises with different contents as reasons for a claim, threatens a regress on claim contents. At each stage vindication of one commitment may involve appeal to commitments that have not previously been invoked, for which the issue of demonstrating entitlement can arise anew, so the issue is merely put off. Or at some point a circle is closed by appeal to a set of premises whose entitlement has already been brought into question (and put off). Then the argument offered for a claim amounts to something that could be made explicit (eliminating the intervening steps) by a stuttering inference of the form "*p*, therefore *p*," which cannot create entitlement.

The communicational style of vindication, in which one interlocutor appeals to another interlocutor's assertional avowal of a commitment with the same content, threatens a corresponding regress on interlocutors. At each stage vindication of one interlocutor's commitment may involve appeal to the commitment of some interlocutor who has not yet been appealed to, for whom the issue of demonstrating entitlement can arise anew—so the issue is merely put off. Or at some point a circle is closed by appeal to the assertion of some interlocutor whose entitlement has already been brought into question (and put off). Then the deferral that seeks to vindicate the claim amounts to something that could be made explicit (eliminating the intervening steps) by a self-citation of the form "I am relying for my entitlement to *p* on the authority of my own claims that *p*," which cannot create entitlement.

The situation is not fundamentally altered by the fact that tracing back a single entitlement might involve both inferential and communicational appeals—that the chain of inheritance might comprise both justifications and deferrals. These are mechanisms for spreading entitlements, not for originating them; combining the two merely results in more complicated regresses and circles. What gets the process off the ground? What gives these multiplicative mechanisms something to work with in the first place, so that chains of vindication can come to an end? This question arose above in connection

with the authorizing function of assertions and was pursued through the notion of fulfilling the responsibility to vindicate the authorizing commitment, by demonstrating one's entitlement. But looking at those mechanisms raises the same issue all over again.

The worry about a regress of entitlements is recognizably foundationalist. It can be responded to by appealing to the fundamental pragmatic commitment to seeing normative statuses (in this case entitlement) as implicit in the social practices that govern the giving and asking for reasons. Those practices need not be—and the ones that actually confer content on our utterances are not—such that the default entitlement status of a claim or assertional commitment is to be guilty until proven innocent. Even if all of the methods of demonstrating entitlement to a commitment are regressive (that is, depend on the inheritance of entitlement), a grounding problem arises in general only if entitlement is never attributed until and unless it has been demonstrated. If many claims are treated as innocent until proven guilty—taken to be entitled commitments until and unless someone is in a position to raise a legitimate question about them—the global threat of regress dissolves.

One of the lessons we have learned from thinking about hyperbolic Cartesian doubt is that doubts too sometimes need to be justified in order to have the standing to impugn entitlement to doxastic commitments. Which commitments stand in need of vindication (count as defective in the absence of a demonstration of entitlement to them) is itself a matter of social practice—a matter of the practical attitudes adopted toward them by the practitioners. The different circumstances under which various claims are taken or treated as requiring justification (or vindication by deferral) is part of what confers on the sentences that express them the different meanings that they have. It is part of the inferential role they play, in the broad practical sense of that expression, that includes the conditions under which inferential performances of various sorts are appropriate or obligatory. Claims such as "There have been black dogs" and "I have ten fingers" are ones to which interlocutors are treated as prima facie entitled. They are not immune to doubt in the form of questions about entitlement, but such questions themselves stand in need of some sort of warrant or justification. Entitlement is, to begin with, a social status that a performance or commitment has within a community.[29] Practices in which that status is attributed only upon actual vindication by appeal to inheritance from other commitments are simply unworkable; nothing recognizable as a game of giving and asking for reasons results if justifications are not permitted to come to an end.

The model presented here has what might be called a *default and challenge structure* of entitlement. Often when a commitment is attributed to an interlocutor, entitlement to it is attributed as well, by default. The prima facie status of the commitment as one the interlocutor is entitled to is not permanent or unshakeable; entitlement to an assertional commitment can

be challenged. When it is *appropriately* challenged (when the challenger is *entitled* to the challenge), the effect is to void the inferential and communicative authority of the corresponding assertions (their capacity to transmit entitlement) unless the asserter can vindicate the commitment by demonstrating entitlement to it.

This is what was meant by saying that the broadly justificatory responsibility to vindicate an assertional commitment by demonstrating entitlement to it is a *conditional* task-responsibility. It is conditional on the commitment's being subject to a challenge that itself has, either by default or by demonstration, the status of an entitled performance. Indeed, the simplest way to implement such a feature of the model of asserting is to require that the performances that have the significance of challenging entitlements to assertional commitments themselves be assertions. One then can challenge an assertion only by making an assertion incompatible with it. (Recall that two claims are incompatible just in case commitment to one precludes entitlement to the other.) Then challenges have no privileged status: their entitlement is on the table along with that of what they challenge. Tracing the provenance of the entitlement of a claim through chains of justification and communication is appropriate only where an actual conflict has arisen, where two prima facie entitlements conflict. There is no point fixed in advance where demands for justification or demonstration of entitlement come to an end, but there are enough places where such demands *can* end that there need be no *global* threat of debilitating regress.

This is the sort of picture of the practices of giving and asking for reasons that Wittgenstein suggests, but it is recognizable already in Socratic elenchus. In the present context, the proper question is not whether practices that incorporate such a default-and-challenge structure of entitlements are somehow in principle defective in view of some a priori rationalistic criterion of what it is to be *really* entitled to a claim. The proper question is rather, What sort of propositional contents can reason-constituting practices of this sort confer on the scorekeeping attitudes they govern, the deontic statuses they institute, and the performances they acknowledge as having the significance of assertions? The claim eventually to be made is that such practices suffice to confer objectively representational propositional contents on claims, objective truth conditions according to which the correctness of an assertion can depend on how things are with the objects represented by it, to the extent that the entire linguistic community could be wrong in its assessment regarding it.

5. Internal Sanctions: Doxastic Commitments without Entitlements Lack Authority

The picture, then, is one in which giving reasons is obligatory only if they have been appropriately asked for. What has the significance of a

challenge (a demand for reasons) is just more assertions, whose entitlements are subject to the same sort of assessments as any others. All are weighed in the same balance. The fundamental concept in terms of which the default-and-challenge structure is adumbrated is the deontic attitude of *attributing* entitlement, of one interlocutor's taking or treating another *as* entitled to a commitment or performance. Now that the background presupposed by an interlocutor's conditional task-responsibility to demonstrate entitlement to the commitment undertaken by an assertional performance has been filled in, it is possible to be a bit clearer about this deontic attitude. The question is, What practical difference does it make whether the asserter is entitled to an assertional commitment? That is, What is the pragmatic significance of the distinction between warranted and unwarranted assertional commitments? What is it about a scorekeeper's treatment of an attributed commitment that makes it appropriate to describe that practical attitude as one of taking the commitment to be one the asserter is not *entitled* to?

Answering this question requires considering all three dimensions of assertion mentioned above: relations between commitments and entitlements, relations between intra- and interpersonal inferential significance, and relations between authority and responsibility. Since the *authority* of an assertion consists in its inferentially licensing or warranting further commitments, and this is a matter of *inheritance* of entitlement, an assertion expressing a commitment that is not taken to be one the performer is entitled to will not be taken to have inferential authority. Although making a claim by asserting a sentence is putting it forward *as* a fit premise for oneself and others to draw conclusions from, it will be accorded that status (its authority recognized) only by those scorekeepers who attribute not only the *commitment* the performance expresses but also an *entitlement* to that commitment. Absent such entitlement, assertion is an attempt to lend what one does not possess. Failure to shoulder the justificatory *responsibility* associated with entitlement to an assertional commitment (supposing it to have been appropriately challenged) renders void its *authority* as an inferential warrant for further commitments. Inferential authority and justificatory responsibility are coordinate and commensurate.

In the ideal Sprachspiel being described, making a claim one is not entitled to (even as a challenge) is a kind of impropriety, the violation of a norm. For a performance to have this sort of status or significance within or according to a set of practices—for this sort of norm to be implicit in or instituted by those practices—requires that the practices include attitudes of taking, treating, or acknowledging performances *as* incorrect in that particular way. Some sort of sanction must be available, with respect to which it can be specified what a practitioner *does* in adopting those practical attitudes. The practical significance of lack of entitlement consists in liability to *punishment* of some kind.

As has already been pointed out, however, such punishment need not

consist in *external* sanctions—responses such as beating with sticks, which are interpretable as punishments (perhaps in virtue of functioning as negative reinforcement in a behavioral-statistical sense) apart from the normative significance they in turn have within the practices in question. One can coherently interpret a community as engaging in practices in which performances are treated as having the significance of promises (or of the more primitive sort of nonlinguistic undertaking of task-responsibility, of which taking the queen's shilling is an example) even if the *only* sanction for failing to perform as one has committed oneself to do is to disqualify oneself from counting in the future as eligible to undertake such commitments. Something like this is what happened to the boy who cried "Wolf." Having several times committed himself to the claim that a wolf was present (thereby licensing and indeed obliging others to draw various conclusions, both practical and theoretical) under circumstances in which he was not entitled by the evident presence of a wolf to undertake such a commitment and to exercise such authority, the boy was punished—his conduct practically acknowledged as inappropriate—by withdrawal of his franchise to have his performances treated as normatively significant.

Unlike the case of the liar who ceases to be believed or of the irresponsible promiser who ceases to be relied upon, however, the internal sanctions constituting the practical recognition of an assertional performance as one the performer is not entitled to do not, on the present model, deal with the significance accorded to *other* performances of the same sort by that individual. Those sanctions have rather to do with the significance assigned to that very performance. Treating the commitment expressed by an assertional performance as one the asserter is not entitle*d* to is treating it as not entitl*ing* that interlocutor and the audience of the performance to commitments whose contents follow inferentially from the asserted content. The practical sanction constitutive of the implicit norm governing entitlement to assertional commitments is internal to the system of scorekeeping attitudes the practice comprises. Taking someone to be (or not to be) entitled to a claim has consequences for what deferring performances one acknowledges as in order—and this in turn affects what deferrings one is oneself disposed to produce under various circumstances. But there need be no social pattern of performances and dispositions describable in nonnormative terms that is either necessary or sufficient for the constitution of such deontic attitudes.

IV. SCOREKEEPING: PRAGMATIC SIGNIFICANCE AND SEMANTIC CONTENT

1. Lewis's Version of Scorekeeping in Language Games

The particular way in which the pragmatic significance of speech acts and deontic statuses is related to their semantic contents can be clarified

by looking more closely at the metaphor of *scorekeeping* by linguistic prac-
titioners. The use made of the notion here is an adaptation of an idea intro-
duced by David Lewis, in his paper "Scorekeeping in a Language Game." He
suggests thinking about the rule-governedness of conversation by using some
of the concepts appropriate to games that evolve according to a score func-
tion. This notion is explained in terms of baseball:

> At any stage in a well-run baseball game there is a septuple of numbers
> $<r_v, r_h, h, i, s, b, o>$ which I shall call the *score* of the game at that stage.
> We can recite the score as follows: the visiting team has r_v runs, the
> home team has r_h runs, it is the h^{th} half (h being 1 or 2) of the i^{th} inning;
> there are s strikes, b balls, and o outs.[30]

The *constitutive rules* of the game are then of two sorts:

> *Specifications of the kinematics of score.* Initially the score is $<0, 0, 1,$
> $1, 0, 0, 0>$. Thereafter, if at time t the score is s, and if between time t
> and t' the players behave in manner m, then at time t' the score is s',
> where s' is determined in a certain way by s and m.
> *Specifications of correct play.* If at time t the score is s, and between
> time t and time t' the players behave in manner m, then the players
> have behaved incorrectly. (Correctness depends on score: what is cor-
> rect play after two strikes differs from what is correct play after three.)
> What is not incorrect play according to these rules is correct.[31]

He then points out that it is possible to use specifications of these sorts to
define 'score' and 'correct play', by using the notion of *score function*, which
is

> the function from game-stages to septuples of numbers that gives the
> score at every stage. The specifications of the kinematics of score, taken
> together, tell us that the score function evolves in such-and-such way.
> We may then simply define the score function as the function which
> evolves in such-and-such way . . . Once we have defined the score
> function, we have thereby defined the score and all its components at
> any stage. There are two outs at a certain stage of a game, for instance,
> if and only if the score function assigns to that game-stage a septuple
> whose seventh component is the number 2.[32]

Correct play is specified in terms of current score and current behavior. Since
the required relation between these is codified in the score function, it also
defines correct play.

The idea is, then, that the evolution of a linguistic interchange or conver-
sation can be thought of as governed by implicit norms that can be made
explicit (by the theorist) in the form of a score function. Here are some of the
analogies Lewis points to:

—Like the components of a baseball score, the components of a conversational score at any given stage are abstract entities. They may not be numbers, but they are other set-theoretic constructs . . .

—What play is correct depends on the score . . .

—Score evolves in a more-or-less rule-governed way. There are rules that specify the kinematics of score:

> If at time t the conversational score is s, and if between time t and time t' the course of conversation is c, then at time t' the score is s', where s' is determined in a certain way by s and c.

Or at least:

> . . . then at time t' the score is some member of the class S of possible scores, where S is determined in a certain way by s and c . . .

—To the extent that conversational score is determined, given the history of the conversation and the rules that specify its kinematics, these rules can be regarded as constitutive rules akin to definitions. Again, constitutive rules could be traded in for definitions: the conversational score function could be defined as that function from conversation-stages to n-tuples of suitable entities that evolves in the specified way.[33]

As Lewis applies this idea, the elements of the conversational score are things such as sets of presupposed propositions and boundaries between permissible and impermissible actions. Then the acceptability of uttering a particular sentence at a given stage can depend on what is being presupposed. Similarly, the saliencies established by the current score can determine the extension or even the intension of terms such as 'the pig'.

2. Deontic Scores and the Pragmatic Significance of Speech Acts

This idiom can be adapted to the model of linguistic practice introduced in this chapter by specifying scores in terms of *deontic statuses*. Linguistic practice as here described can be explained in terms of a score function that determines how the *deontic score* at each stage in a conversation constrains both what performances are appropriate and what the consequences of various performances are—that is, the way they alter the score. The concept of the *pragmatic significance* of a speech act is central to the theoretical metalanguage being employed here. It is a generalization of Dummett's idea of specifying the use of an expression in terms of the pair of its *circumstances* of application and *consequences* of application. (Recall that Dummett's idea was adopted in Chapter 2 as a way of connecting a normative

pragmatics with an inferential semantics.) In scorekeeping terms, the significance of a speech act consists in the way it interacts with the deontic score: how the current score affects the propriety of performing the speech act in question, and how performing that speech act in turn affects the score. Deontic scores consist in constellations of commitments and entitlements on the part of various interlocutors. So understanding or grasping the significance of a speech act requires being able to tell in terms of such scores when it would be appropriate (circumstances of application) and how it would transform the score characterizing the stage at which it is performed into the score obtaining at the next stage of the conversation of which it is a part (consequences of application). For at any stage, what one is permitted or obliged to do depends on the score, as do the consequences that doing has for the score. Being rational—understanding, knowing how in the sense of being able to play the game of giving and asking for reasons—is mastering in practice the evolution of the score. Talking and thinking is keeping score in this sort of game.

In baseball the components of the score (for instance the status a performance can have as a *ball* or a *strike*, or an *out*) are defined in formal terms by the role they play in the process of keeping score—that is, their function in determining what counts as correct play according to the kinematics of score, as codified in the score function. So it is with the components of deontic score in terms of which linguistic practice is to be understood. The deontic statuses of commitment and entitlement are defined in formal terms by the way they can be used to keep track of the moves made in the game of giving and asking for reasons—that is, their function in determining what counts as correct play according to the kinematics of score, as codified in the score function. In this way the notion of *commitment* in linguistic practice plays a role like that of *strike* in baseball: each is an artificial, scorekeeping device.

Besides this formal characterization, however, there is also a material aspect to each of the components in the score, in virtue of which a particular performance qualifies as a ball, a strike, or an out. This material aspect is represented in Lewis's formulation of the kinematics of score quoted above by "the manner *m*," which characterizes the behavior that changes the score from one stage to the next. In games such as baseball, which are not purely formal games (by contrast to chess or tic-tac-toe), the manner in which the score is changed cannot itself be specified entirely in terms of the concepts by means of which the score itself is specified. The complex manner in which a concrete performance qualifies as having the status of a strike or an out invokes such further concepts as the swinging of a bat, the passage of the baseball through a certain region of space specified relative to the position of the batter's body, catching the baseball on the fly, and so on. These further concepts give a material content to the scorekeeping concepts, beyond the formal content they have in virtue of their role in scorekeeping. So it is as well with deontic scorekeeping in linguistic practice. In that case the mate-

rial element concerns such issues as which utterances count as undertaking which commitments, or as deferring to the authority of another asserter, or as invoking noninferential responsive authority.

In baseball the application of scorekeeping vocabulary to particular performances is governed by rules, which are expressed largely in nonscorekeeping vocabulary. The use of this nonscorekeeping vocabulary accordingly answers to norms implicit in the practices of using that vocabulary in contexts other than baseball, where terms like 'inch', 'touch', 'between', and so on already have well-established circumstances of application. Final authority over the application of these rules is vested in the practice of officials. The exact character of this authority is a somewhat complicated matter. There is sometimes an inclination to think of it as *constitutive* of the correct application of the scorekeeping vocabulary, as it is taken to be for instance in the escalating claims of the competing umpires in the familiar tale:

> *First Umpire:* I calls 'em as I sees 'em.
> *Second Umpire:* I calls 'em as they is.
> *Third Umpire:* Until I calls 'em, they ain't.

On such a view, the rules function as something like guides or advisory maxims for the judgment of the umpire, who *makes* a throw into a strike when he *takes* it *as* a strike. But though the attitude of the umpire does determine the status of a throw as a strike for official scorekeeping purposes (that is, does determine what the score is), the use of nonscorekeeping vocabulary in stating the rules that determine how the scorekeeping vocabulary *ought* to be applied to particular cases establishes a perspective from which the judgment of an umpire can nonetheless be understood to be *mistaken.* Metarules explicitly envisage the possibility of such mistakes and, without obliterating their status as such, set up a default-and-challenge system that leaves the umpire entitled to scorekeeping judgments even in the case where they are in fact mistaken, so long as they survive any appropriate challenges that are actually offered.

Linguistic scorekeeping on assertional commitments and entitlements has analogs to both of these dimensions of authority concerning the score. On the one hand, the actual attitudes of scorekeepers are essential in determining the score. On the other hand, the formation of those attitudes is itself subject to norms; scorekeeping is something that can be done correctly or incorrectly. This is not, of course, because it is in general governed by explicit rules; the regress that Wittgenstein and Sellars point to shows that. It is of the utmost importance to the present project to offer an account of what one is doing in taking a scorekeeper to have gotten things wrong, to have attributed commitments different from what the one to whom they are attributed is *really* committed to. For it is in terms of this practical attitude that the possibility of understanding the application of concepts as subject to objective representational constraint—as subject to assessment for being correct

or incorrect in a sense that involves answering to how the things the concepts are applied to actually are, rather than to anyone's attitudes toward them—is eventually to be explained.[34]

One fundamental difference between a game such as baseball and the game of giving and asking for reasons is the *perspectival* nature of the scorekeeping involved in the latter case. As Lewis sets things up, each stage of a baseball game has a single score. One might instead associate a different score with each of the two teams—though some elements, such as the specification of the inning, would be common to both teams at each stage. Linguistic scorekeeping as here construed is more like that: each interlocutor is assigned a different score. For to each, at each stage, different commitments and different entitlements are assigned. There may be large areas of overlap, since almost everyone is committed and entitled to such claims as that 2 + 2 = 4, that red is a color, and that there have been black dogs. But there will also be large areas of difference, if for no other reason than that everyone has noninferentially acquired commitments and entitlements corresponding to different observable situations. These differences ramify both because of the inferential consequences of such observations and because their public availability for inheritance of attitudes and attribution based on testimony varies with conversational exposure. As a result, no two individuals have exactly the same beliefs or acknowledge exactly the same commitments. As with baseball, instead of thinking of these as different scores associated with different interlocutors (corresponding to different teams), they can be aggregated into one grand score for each stage of the conversation of a linguistic community, so long as it is kept clear (in a way corresponding to Lewis's use of subscripts) which deontic statuses are being attributed to which interlocutors.

But linguistic scorekeeping is also perspectival in a way that has no analog at all in baseball. Not only are scores kept *for* each interlocutor, scores are also kept *by* each interlocutor. In baseball there is just one official score, whether it is thought of as the score of the whole game or as the set of scores of each of the teams (and this is, as was pointed out, compatible with there nonetheless being a sense in which the umpire who determines the official score may make a mistake in calling a certain performance a strike). But part of playing the game of giving and asking for reasons is keeping track of the commitments and entitlements of the other players, by *attributing* commitments and entitlements. Just as each interlocutor is typically at each stage attributed a different set of deontic statuses, so each interlocutor typically has at each stage a different set of attitudes or attributions. What C is committed to according to A may be quite different, not only from what D is committed to according to A, but also from what C is committed to according to B. Linguistic scorekeeping practice is *doubly* perspectival.

The idea is that the deontic attitudes of each interlocutor A constitute one perspective on the deontic statuses of the whole community. There are, to

begin with, the commitments that A acknowledges and the entitlements that A claims. Then for each other interlocutor, there are the commitments and entitlements that A attributes to that individual. The different sorts of speech acts are to be understood in terms of the different consequences they have for the score that each interlocutor keeps, that is, in terms of how they affect the deontic attitudes of various interlocutors. If B asserts that p, B thereby acknowledges (and so undertakes) a commitment to p. So such a commitment ought to be attributed to B by anyone in a position to overhear or otherwise find out about that remark.

The pragmatic significance of an assertion goes far beyond this simple shift in deontic attitude on the part of other scorekeepers, however. For the speech act B performs has an inferentially articulated *content*, which relates it to other contents. Undertaking commitment to p is undertaking commitment as well to its inferential consequences—to those claims q that are related to it as conclusions of commitment-preserving inferences having p as premise. So if, as a result of B's assertion, A's deontic attitudes change in that A comes to attribute to B a commitment to p, then A is obliged also to attribute to B commitment to q. Or rather, A's treating this as a good inference consists in A's being disposed to keep score in this way, linking the attribution of commitment to q consequentially to attribution of commitment to p. Again, for r to be incompatible with p is for commitment to p to preclude entitlement to r. For A to treat these contents as incompatible is for A to be disposed to withhold attribution of entitlement to r whenever A attributes commitment to p.

Besides these intercontent, intrapersonal scorekeeping consequences of B's speech act, the assertion may have intracontent interpersonal consequences regarding A's attitudes. For if A takes B to be entitled to the claim that p (either noninferentially or as the conclusion of an inference), then this may result in A's taking C (who also overheard the remark) to be entitled to that claim—but on the basis of testimony, to be defended by deference, rather than either noninferentially or inferentially. The effects of a speech act on the practical attitudes by means of which A keeps score on the deontic statuses of various interlocutors depends both on the antecedent score—what they were already taken to be committed and entitled to—and on the content expressed.

3. Inferentially Articulated Significance: Force and Content

Specifying the pragmatic significance of a speech act kind such as assertion requires showing how the transformation of the score from one conversational stage to the next effected by such a speech act systematically depends on the semantic content of the commitment undertaken thereby. Starting with a notion of the pragmatic significance of speech acts—understood in terms of transformations of the deontic attitudes by which inter-

locutors keep track of each other's commitments and entitlements—it is possible to understand both what it is for two commitments to have the same content and what it is for two commitments to be undertaken by or attributed to the same interlocutor. Not only can these scorekeeping attitudes and shifts of attitude be used to define both contents and interlocutors, the justificatory and communicational links between them can be used to define the notion of *representation*. This is the burden of the discussion of the hybrid deontic attitudes that are made assertionally explicit in the form of *de re* ascriptions of propositional attitudes, in Chapter 8. So the notion of linguistic scorekeeping is intended to play a more fundamental explanatory role here than Lewis has in mind for it. For he is happy to think of conversational scores as kept track of in "mental scoreboards," consisting of attitudes he calls "mental representations" of the score (representations, presumably, whose content is *that* some component of the score is currently such and such).[35] Clearly he does not envisage a project such as the present one, in which both the nature of mental states such as belief and their representational contents are themselves to be understood in terms of their role in scorekeeping practices, rather than the other way around.

Consider first the notion of the *content* of a speech act or an intentional state. It is motivated first by the idea that speech acts, attitudes, and states of different kinds might share a content—in Fregean terminology, that different sorts of *force* can attach to the same sense. It requires further that the significance of a speech act depends in a systematic way on the content and the sort of force that is attached to it. Dummett's way of putting the point is this: "The implicit assumption underlying the idea that there is some one key concept in terms of which we can give a general characterization of the meaning of a sentence is that there must be some uniform pattern of derivation of all the other features of the use of an arbitrary sentence, given its meaning as characterized in terms of the key concept. It is precisely to subserve such a schema of derivation that the distinction between sense and force was introduced: corresponding to each different kind of force will be a different uniform pattern of derivation of the use of a sentence from its sense."[36] It does not simply go without saying that such a notion of content is to be had. Use of the theoretical concept of content involves a commitment to displaying the "uniform pattern of derivation of the use from the content," which Dummett talks about. As he goes on to indicate, one way of reading some of Wittgenstein's remarks is as "rejecting the whole idea that there is any one key idea in the theory of meaning: the meaning of each sentence is to be explained by a direct characterization of all the different features of its use; there is no uniform means of deriving all the other features from any one of them."[37] This is the point at which semantic theory and pragmatic theory must mesh. It is possible to associate many sorts of things with sentences and other linguistic expressions. What makes the association a *semantic* one is precisely the possibility of appealing to it to explain the

proprieties that govern the *use* of those expressions. Calling what one associates with expressions 'contents', 'propositions', 'sets of possible worlds', 'truth conditions', 'extensions', or 'referents' is at best issuing a promissory note that hints at how what are put forward as their semantic correlates ought to be taken to be relevant to determining how those expressions are correctly used. In the absence of a pragmatics offering an account of what it is to express a content or proposition, to take the actual world to be contained in a set of possible worlds, to try to utter truths, or to employ expressions with various extensions and referents, the theorist's entitlement to the commitment undertaken in treating these associations as semantic is liable to challenge. Semantics answers to pragmatics, attributions of content to explanations of use.

In the present case the pragmatics (comprising the practical proprieties governing linguistic expressions and intentional states alike) is couched in terms of social deontic scorekeeping. The *force* of an utterance, the *significance* of a speech act, is to be understood in terms of the difference it makes to what commitments and entitlements are attributed and undertaken by various interlocutors—that is, in terms of the alteration of deontic scorekeeping attitudes it underwrites. Indeed deontic statuses are to be understood just as ways of keeping such scores. The paradigmatic speech act kind of asserting is specified as having the significance of an undertaking of a commitment (and so the licensing of attributions of that commitment), the licensing or authorizing of further undertakings of such commitments, and the undertaking of a conditional task-responsibility to demonstrate entitlement to the commitment undertaken, if appropriately challenged. This is the sort of significance that must be determined "according to a uniform pattern" (in Dummett's phrase) by the sort of semantic content that is associated with the expressions that qualify as sentences in virtue of their freestanding utterances having this kind of assertional significance.

To be entitled to an *inferential* conception of the contents that qualify as propositional in virtue of their being assertible, then, requires showing how particular assertional significances result from the general account of the speech act of asserting when particular inferential roles are associated with what is asserted. The model of asserting has been constructed with just this criterion of adequacy in mind. To specify the inferential content associated with a sentence, one must, to begin with,[38] indicate the role it plays (in relation to the contents expressed by other sentences) in three different sorts of broadly inferential structure: committive inferences, permissive inferences, and incompatibilities. Doing so is saying what it follows from, what follows from it, and what it precludes or rules out. These are characterized as "broadly" inferential because all of them involve alterations of deontic status that have other alterations of deontic status as their consequences.

In the same sense, the sort of authority that observation reports exhibit counts as broadly inferential because of the reliability inference it involves on the part of the attributor of such authority (discussed in the next chapter).

Although it sounds paradoxical, for this reason the role of a sentence in noninferential reporting should also be understood as falling under the rubric "(broadly) inferential role." Two features of such specifications are worth focusing on in this context. First, the informal explanations of these inferential relations are in terms of precisely the deontic statuses of commitment and entitlement that are kept track of by the scorekeeping attitudes of interlocutors—that is, just the terms in which the pragmatics is couched. Second, the account of assertional significance in general requires nothing more than inferential roles articulated along these three dimensions in order to determine the significance for social deontic scorekeeping of an assertional utterance.

On the first point, specifying the committive-inferential role of a sentence is specifying the commitment-preserving inferences in which that sentence serves as a conclusion, and those in which (along with other auxiliary hypotheses) it plays an essential role as a premise—essential in that if it is omitted, one could be committed to the remaining premises without therefore counting as committed to the conclusion. This sort of inference is the material-inferential genus of which deductively valid logical inferences are a formal species. Similarly, specifying the permissive-inferential role of a sentence is specifying the entitlement-preserving inferences in which (along with other auxiliary hypotheses) it plays an essential role—essential in that if it is omitted, one could be committed and entitled to the remaining premises (and to no incompatible defeasors) without therefore counting as being entitled to the conclusion. This sort of inference is the material-inferential genus of which inductively good inferences are a species. Incompatibility relations are thought of as broadly inferential because their construal in terms of the deontic statuses of commitment and entitlement is analogous to the first two. The difference is that where the deontic statuses relevant to the two species of strictly inferential relation are homogeneous—both premises and conclusions being assessed in terms of commitments, or both in terms of entitlements—in the case of incompatibility, they are heterogeneous. For to say that two claims are incompatible is to say that *if* one is committed to the first, *then* one is not entitled to the second.[39]

To anticipate the discussion of the next chapter: the sort of default entitlement characteristic of observation reports (perhaps the most important species of this genus) is thought of as broadly inferential because the one who attributes such authority implicitly endorses the reliability of the reporter (under these circumstances and with regard to such contents). Treating someone as a reliable reporter is taking the reporter's *commitment* (to this content under these circumstances) to be sufficient for the reporter's *entitlement* to that commitment. This is endorsing an inference, in the broad sense that corresponds to the consequential relationship between attributing commitment and attributing entitlement to it.[40] Like that involved in incompatibility relations among contents, and in contrast to the two strictly inferential components of content, this inference is heterogeneous with respect to the

deontic statuses involved. Unlike incompatibility, the content of the commitment and the content of the entitlement involved in the reliability inference are the same. The essential role played by endorsement by the attributor of this reliability inference qualifies the role of utterances as elicited observationally as the exercise of reliable differential responsive dispositions for inclusion as contributing to the broadly inferential role of sentences.

The connection between these four kinds of broadly inferential proprieties and the deontic statuses of commitment and entitlement has a double significance. On the one hand, it means that it is possible to understand how social practices of keeping score on commitments and entitlements could confer inferential roles articulated along these four dimensions on the expressions that are caught up in them. That is, social deontic scorekeeping provides an explanation of how expressions must be *used* in order to have contents of this sort associated with them—associated not by the theorist's stipulation but by the practical attitudes of the practitioners whose linguistic conduct is being interpreted. On the other hand, the connection between broadly inferential proprieties of these four kinds and the two sorts of deontic status they involve (homogeneously or heterogeneously) means that once contents articulated in terms of these proprieties have been associated with expressions, it is possible to derive "according to a uniform pattern" the significance that uttering them assertively has for scorekeeping in terms of commitments and entitlements. *Inferring* is accordingly the key concept linking semantic content and pragmatic significance. For not only can propositional semantic contents be understood as inferential roles, but proprieties of inference can be made sense of pragmatically, and specifically assertional significance can be understood in terms of them.

4. How Inferential and Incompatibility Relations among Contents Affect the Score

The significance of an assertion of p can be thought of as a mapping that associates with one social deontic score—characterizing the stage before that speech act is performed, according to some scorekeeper—the set of scores for the conversational stage that results from the assertion, according to the same scorekeeper.[41] Suppose A is such a scorekeeper, and B is such an asserter. The way A's score ought to be transformed is settled by the content of B's claim, thought of as its tripartite inferential role in commitment- and entitlement-preserving consequence relations and in incompatibility relations connecting commitments and entitlements. To begin with, A must add p to the list of commitments attributed to B (supposing the more interesting case in which A does not already attribute to B a commitment with that content). A should then add also commitment to any claims q that are committive-inferential consequences of p, in the context of the other claims attributed to B. These will vary, depending on the auxiliary hypotheses available, according to what other commitments A already attributes to

B. This is closing *A*'s attributions to *B* under commitment-preserving inferences. This closure is determined, in the context of the prior score, by whatever committive-inferential role *A* associates with *p* as part of its content.[42]

Next, the incompatibility relations that *p* (and so its commitment-inferential consequences) stand in must be consulted to determine which, if any, of the entitlements *A* previously attributed to *B* are precluded by the newly attributed commitment. Assertions add new commitments, but they can not only add but also subtract entitlements. Then, in the light of the incompatibility relations associated with all of the commitments attributed to *B*, *A* can attribute *entitlements* to any claims that are *committive*-inferential consequences of commitments to which *B* is already taken to be entitled, closing the attributed score under commitment-preserving inferences, where the resulting attributions of entitlement are not defeated by incompatibilities.[43]

Next, constrained by the entitlement-precluding incompatibility relations associated with all of the other commitments attributed to *B*, *A* can attribute entitlements to any claims that are permissive-inferential consequences of commitments to which *B* is already taken to be entitled, closing the attributed score under entitlement-preserving inferences not defeated by those incompatibilities.[44] Then *A* needs to assess *B*'s entitlement to the claim that *p*, by looking at good inferences having it as a conclusion and premises to which *B* is committed and entitled. This is determined by the first two elements of the broadly inferential role *A* associates with *p*. Similarly, *A* must assess the possibility of *B*'s noninferential default entitlement to *p*. Whether *B*'s undertaking of commitment to *p* falls within the scope of any good reliability inferences, according to *A*, depends on what else *A* is committed to—concerning the conditions under which the deontic status was acquired, implicitly, whether they are among those (if any) in which *B* is a reliable reporter with respect to contents such as *p*. Again, *A* must assess *B*'s entitlement to *p* as testimony, by inheritance of the entitlement *A* attributes to some other interlocutor (possibly even *A*) who has asserted it at an earlier stage. If *A* takes *B* to be entitled to *p* by any of these mechanisms of inheritance and default, then *A* will take *B* to have successfully entitled others (including *A*) to that claim (in the absence of incompatible defeasors). In this way the broadly inferential content that *A* associates with *B*'s claim determines the significance *B*'s assertional speech act has from the point of view of *A*'s scorekeeping, the difference it makes to the deontic attitudes of attributing and acknowledging commitments and entitlements by means of which *A* keeps track of everyone's deontic statuses.

5. Deferrals, Disavowals, Queries, and Challenges

The model of assertion defined by a scorekeeping function that appeals in this way to broadly inferential assertible contents to determine the significance of assertions of those contents can be enriched by allowing

various auxiliary sorts of speech acts. *Deferrals* have already been mentioned. No new sorts of content need be considered in order to specify the significance of deferring for social deontic scorekeeping. The content associated with a deferral is just the assertible content of the commitment that the deferrer is seeking to vindicate by indicating a testimonial path whereby entitlement to it can be inherited. It is the force associated with that content that is different from the assertional case.

A assesses *C*'s deferral to *B* concerning *p* by assessing first *B*'s entitlement to *p* (as considered above), and then *C*'s entitlement to inherit it. This latter is a matter not only of its being the case that *A* does not attribute to *C* commitment to anything incompatible with *p* (an issue in general independent of whether *A* attributes to *B* commitment to anything incompatible with *p*). For even if *A* does not attribute to *C* commitments incompatible with the *claim B* made (the commitment undertaken), it is still possible that *A* attributes to *C* commitments incompatible with inheriting *B*'s *entitlement* to it. This would happen if *B*'s entitlement, according to *A*, depends on justifying the claim *p* by appeal to the claim *q*, where *A* takes it that *C*, but not *B*, is committed to some claim incompatible with *q*. Similarly, it might be that *C* is committed, though *A* is not, to some claim incompatible with one of the conditions (according to *A*) for *B*'s observational authority with respect to *p*. Thus if *C* takes it that *B* is looking through a tinted window, *A* may take this to preclude *C*'s inheritance of entitlement to *B*'s noninferential report of the color of a piece of cloth, even though *A* takes it that *C* is wrong about the conditions of observation.

Disavowals, queries, and *challenges* are three other speech acts auxiliary to assertion that it is useful—from a scorekeeping point of view—to include in a model of the game of giving and asking for reasons. Disavowals permit one to repudiate or disclaim a commitment one has previously undertaken or to make it clear that one does not acknowledge such a commitment. Again no new sorts of content need to be considered; the force or significance of speech acts of this sort (the difference they make to the score that interlocutors keep on each other's deontic statuses) is determined by associating with them broadly inferential, assertible contents of the sort already discussed. It is the force of the speech act to which such contents attach that is different. For *A* to take *B*'s disavowal of commitment to *p* to be successful is for *A* to cease attributing commitment to *p* to *B*, and to reinstate any attributed entitlements that were withheld because they were defeated by their incompatibility with *p*.

Why might a disavowal not be successful? Because of the two fundamental ways *B* can undertake a commitment (and thereby license *A* to attribute it)—about which more below. For *B* might acquire commitment to *p* directly, by avowing it—that is, by overtly asserting it. Or *B* might acquire that commitment indirectly, as a consequence of a commitment (perhaps itself avowed) to *q*, from which it follows by a good commitment-preserving infer-

ence (according to *A*). In such cases, *B*'s disavowal of *p* can be successful (according to *A*) only if *B* is also prepared to disavow *q*. Indeed, disavowing *p* is indirectly disavowing *q*. But if *B* persists in asserting *q*, that commitment is incompatible with the disavowal, and the disavowal of *p* cannot accordingly reinstate entitlement to claims *A* withholds attribution of entitlement to only because of the flaw in *B*'s title represented by commitment to the incompatible *p*. Again, the sort of significance for *A*'s scorekeeping that disavowals of *p* have requires no further elements of content beyond those involved in its assertional use, even though the significances of the two sorts of speech act are quite different—indeed, in some ways complementary.

It would also be useful to those keeping score if there were some way of eliciting the avowal or disavowal of a particular claim—a way for *A* to find out whether *B* acknowledges commitment to *p*. Such a speech act is a basic query: *p?* By itself, such a speech act would have no effect on the deontic score; only responses to speech acts of this kind would alter the score. In the basic model, there is no reason not to allow anyone to be entitled to such a query at any point in a conversation.

Another sort of speech act that might be distinguished is *challenging* the testimonial authority of an assertion. As was indicated above, this might consist in no more than making an incompatible assertion. But it might be useful from a scorekeeping point of view to have a way of addressing an assertion as a challenge to another assertion. The significance of such a challenge is to bring attributions of entitlement by default into question wherever the challenging assertion is one the challenger is at least prima facie entitled to. For *A* to treat *C*'s challenge of *B*'s assertion of *p* as successful is for *A* to respond to it by withholding attribution of entitlement to *B* for that claim, pending *B*'s vindication of it, whether inferentially or deferentially. This has the effect of making that assertion unavailable (according to *A*'s score) to other interlocutors who might otherwise inherit entitlement to commitments to the same content testimonially from *B*. There is no reason in principle that conflicts of this sort need to be resolvable. The public status of competing claims may remain equivocal in that neither the challenged nor the challenging claim can be vindicated successfully, or in that both can be—though of course *A* will not take it that any one interlocutor could inherit entitlements to commitments to both of the incompatible contents. "Let a thousand flowers blossom. Let a hundred schools of thought contend."

6. Acknowledged and Consequential Commitments

In the next chapter the model of assertion is enriched by adding another variety of discursive commitment. Besides the cognitive commitments undertaken by assertions, *practical* discursive commitments—that is, commitments to *act*—are considered. The speech acts that undertake such commitments, namely expressions of intention, have quite a different sort of

significance from assertions. Attributions of them (the attitudes in terms of which score is kept on these deontic statuses) also behave differently. Yet the significances of these further sorts of performances, statuses, and attitudes can be understood straightforwardly by analogy to the sorts of scorekeeping that have been introduced for the pure assertional case. That extension of the model of the game of giving and asking for reasons provides a way of understanding intentional action and intentional interpretation of agents. It has been suggested that the doxastic commitments undertaken by speech acts having the significance of assertions can serve as analogs of belief—that such deontic statuses can do much of the explanatory theoretical work usually done by the paradigmatic sort of intentional state. That claim clearly cannot be assessed until the model is extended so as to include the possibility of manifesting such commitments not only in what is *said* but also in what is *done*—in action as well as assertion. Before turning to that wider context, however, it is worth pausing briefly to consider what can be said about the relation between belief as an intentional state and the deontic statuses with assertible contents as considered so far, already in the more restricted context of purely assertional practice.

It was indicated above that the assertional practices described so far generate two different senses in which one could be taken to be committed to a claim. Interlocutors undertake some commitments directly, by avowing them overtly: performing speech acts that have the significance of assertions. The commitments one is disposed to avow[45] are *acknowledged* commitments. But in virtue of their inferentially articulated conceptual contents, assertional commitments have consequences. Undertaking a commitment to a claim with one content involves undertaking commitments to claims whose contents are (in the context of one's other commitments) its commitive-inferential consequences. Undertaking commitment to the claim that Pittsburgh is to the West of Philadelphia is one way of undertaking commitment to the claim that Philadelphia is to the East of Pittsburgh. These *consequential* commitments may not be acknowledged; we do not always acknowledge commitment to all the consequences of the commitments we do acknowledge. They are commitments nonetheless. For the only way that deontic *statuses* enter into the scorekeeping specification of assertional practices is as the objects of deontic *attitudes*. Indeed, *all* one can do with a commitment (or entitlement), in the model presented here, is take up a deontic attitude toward it—attribute it or undertake it, either directly by acknowledging it, or indirectly and consequentially. The scorekeeping model trades in talk about the status of *being* committed for talk about proprieties of practical attitudes of *taking* to be committed. Deontic statuses are just something to keep score with, as balls and strikes are just statuses that performances can be treated as having for scorekeeping purposes. To understand them, one must look at actual practices of keeping score, that is, at deontic attitudes and changes of attitude.

These ways in which one can come to be committed to a claim—by acknowledgment and consequentially—correspond to two ways in which we talk about *belief*. In one sense, one believes just what one takes oneself to believe, what one is prepared to avow or assert. In another sense, one believes, willy-nilly, the consequences of one's beliefs. Believing that Pittsburgh is to the West of Philadelphia *is* believing that Philadelphia is to the East of Pittsburgh, whether one knows it or not. This second sense looms particularly large for those who take their role in intentional explanations of behavior as the touchstone for identifying beliefs. For such explanations work only to the extent the individual in question is (taken to be) *rational*—not to have contradictory or incompatible beliefs and to believe the consequences of one's beliefs. From the point of view of the present project, this is because the conclusion of intentional explanations, strictly construed, is always a *normative* one—given intentional states whose contents are thus and so, one *ought* rationally or is rationally *obliged* or *committed* to act in such and such a way. Drawing conclusions about what actually will be or was done requires an additional premise, to the effect that the individual in question has mastered the practices of giving and asking for reasons sufficiently to be disposed to respond to acknowledgment of a (practical) commitment by producing a performance that satisfies it.

The sense of belief in which one is taken actually to believe what one ideally *ought* to believe (at least given what else one believes), call it *ideal* or *rational* belief, can conflict with the sense of belief for which avowal is authoritative. Dennett distinguishes these but thinks of them as competing norms to which a univocal sense of 'belief' must answer: "These two interdependent norms of belief, one favoring the truth and rationality of belief, the other favoring accuracy of avowal, normally complement each other, but on occasion can give rise to conflict."[46] The conflict arises precisely because one can avow incompatible beliefs, and fail to avow even obvious consequences of one's avowals:

> What better source could here be of a system's beliefs than its avowals? Conflict arises, however, whenever a person falls short of perfect rationality and avows beliefs that are either strongly disconfirmed by the available empirical evidence or are self-contradictory or contradict other avowals he has made. If we lean on the myth that man is perfectly rational, we must find his avowals less than authoritative: "You *can't* mean—understand—what you are saying!"; if we lean on his 'right' as a speaking intentional system to have his word accepted, we grant him an irrational set of beliefs. Neither position provides a stable resting place; for, as we saw earlier, intentional explanation and prediction cannot be accommodated either to breakdown or to less than optimal design, so there is no coherent intentional description of such an impasse.[47]

The notion of incompatible beliefs offers no difficulties for a normative construal of intentional states as deontic statuses. There is nothing incoherent or unintelligible about the idea of undertaking incompatible commitments—incompatible or inconsistent beliefs just go into a box with incompatible or inconsistent promises. This is one of the benefits of this sort of approach over causal-functional accounts of intentional states. Yet the tension that Dennett identifies is a real one. The decision to treat belief just as what one is prepared to avow "amounts to the decision to lean·on the accuracy-of-avowal norm at the expense of the rationality norm . . . If we demand perfect rationality, we have simply flown to the other norm at the expense of the norm of accuracy of avowal."[48]

Dennett does not offer any way to reconcile the competing demands that the norm of rationality and the authority of avowals place on attributions of belief. The terminology employed here is animated in part by the thought that 'belief' may simply be *ambiguous* between a sense in which one believes just what one is prepared to avow and a sense in which one also believes what one *ought* rationally to believe, as a consequence of what one is prepared to avow (as already indicated, failures of rationality due to incompatibility cause no particular trouble once intentional states are construed as deontic statuses). An unambiguous, univocal technical term 'doxastic *commitment*' is introduced, which comprises both commitments one is prepared to avow and commitments that follow from those one acknowledges. But attention to the attitudes in terms of which those deontic statuses are explained makes it possible also to distinguish clearly between these two *kinds* of commitment, as 'belief'-talk does not. The proposal is accordingly not to *analyze* belief in terms of commitment but to discard that concept as insufficiently precise and replace it with clearer talk about different sorts of commitment.

The fundamental concept of the metalanguage employed in specifying the model of assertional practice is that of the deontic attitude *attributing* a commitment. For the deontic attitude of *undertaking* a commitment is definable in terms of attribution: undertaking a commitment is doing something that licenses or entitles others to attribute it. Assertional performances or avowals are performances that express the deontic attitude of acknowledging doxastic commitments. They license attribution of (and insofar as they are successful, deferral with respect to) both the commitments they express and those whose contents are appropriate inferential consequences of the contents, commitment to which is overtly acknowledged. The attitude of acknowledging a commitment is in effect that of attributing it to oneself.[49] The fact that one thereby undertakes consequential commitments that may reach beyond what one acknowledges just shows that the generic attitude of *undertaking* a commitment is not to be identified with its species *attributing* a commitment to oneself, which is acknowledging it. The *social* dimension (invoking the perspective of *other* attributors) is essential to understanding

undertaking in terms of proprieties of attributing. The way in which the collaboration of attitudes adopted from two socially distinct perspectives—attributions of commitment to oneself and by others—is required to institute discursive commitments is the central theme of this work. It is in terms of the social-perspectival character of discursive deontic statuses that the notion of *objectivity* is to be made intelligible—both the general normative distinction between what one is *really* committed to do (or ought to do) and what one is merely *taken* by someone to be committed to do, and the more specific version that underwrites the notion of objective representational content, of a claim's correctness answering to how things are with what it represents, rather than to what anyone *takes* to be correct.

The roots of this social-perspectival account can already be discerned in the distinction that a scorekeeper can make between the commitments an interlocutor has undertaken and those that interlocutor acknowledges, and so is prepared to assert. For the attributions of the scorekeeper distinguish between the actual deontic *status* of the one for whom score is being kept, what that interlocutor is *really* (consequentially) committed to, and the deontic *attitudes* of that subject, what that interlocutor acknowledges commitment to by self-attribution. In other words, the notion of consequentially undertaking commitments provides the basis for distinguishing (in terms of the attitudes of someone keeping score) between deontic statuses and deontic attitudes. Indeed (as will appear in Chapter 8, where this issue is explored) the very notion of one propositional content being an inferential *consequence* of another essentially involves a crucial relativity to social perspective: are the auxiliary hypotheses (the premises to be conjoined with the claim in question in assessing its consequences) to be those the scorekeeper assessing the propriety of the inference *undertakes* commitment to, or those the scorekeeper *attributes* to the one whose statuses are being assessed?

Neither answer is correct. The fact that proprieties of inference a claim is involved in can be assessed from either of two social perspectives—that of the one *attributing* commitment to the claim or that of the one *undertaking* that commitment—is fundamental to the very notion of a *propriety* of inference. And since propositional and so conceptual contents of all sorts are constituted by the broadly inferential proprieties of practice in which they are caught up, such contents are essentially social and perspectival in nature. The propositional content of a claim or commitment can be specified only from some point of view; that it would be differently specified in definite ways from other particular possible social perspectives (that is, scorekeepers occupying such perspectives) is an essential part of its being the content it is.

At this point the phrase "social-perspectival character of the contents of discursive commitments" can be little more than a label attached to a promissory note—though the discussion of scorekeeping in this section is intended to give it enough resonances to make it at least a suggestive label. Before that promissory note is redeemed, in Chapter 8, it is necessary to look

much more closely at the sorts of contents that can be conferred on expressions playing various roles in assertional practices of the sort described here. In the intervening chapters the notion of *inferential* articulation is deepened and extended by adding *substitutional* machinery. The notion of substitution inferences permits the extension of the notion of conceptual content introduced here to essentially *sub*sentential expressions such as singular terms and predicates, which can play only indirectly inferential roles—not serving themselves as premises and conclusions of inferences but only occurring in the sentences that can serve in those capacities. The conceptual content expressed by the use of singular terms and predicates is articulated by substitution-inferential commitments (in Chapter 6). Extending this account to the sort of content expressed by the token-reflexive or indexical use of unrepeatable expressions, paradigmatically the sort of *deictic* tokenings that play such an important role in empirical knowledge claims, requires looking still further at *anaphoric* connections among tokenings. (The importance of anaphoric relations for understanding what is expressed by traditional semantic vocabulary, paradigmatically 'true' and 'refers', is argued in Chapter 5.) Anaphora is explained as a certain structure of *inheritance* of substitution-inferential commitments (in Chapter 7). The result is a layered account of the semantic contents that can be conferred by assertional practice as here described—an account whose key concepts are those of *inference, substitution*, and *anaphora*, each permitting a finer-grained analysis of the structures that precede it and which it presupposes.