

Conclusion

The meaning of words is to be determined by their use.

ISAAC NEWTON, *Principia*¹

I. TWO CONCEPTS OF CONCEPTS

1. *Three Kantian Dualisms*

The semantic core of the account of discursive practice presented here is the theory of conceptual content it incorporates. The distinctive features of that theory emerge most clearly when it is contrasted with more traditional ways of thinking about concepts. The most familiar conception, one that is pervasive in contemporary philosophical thought, traces its ancestry back to Kant. Its debt to Kant is most evident in its essentially *dualistic* character: the ways in which the conceptual is contrasted with the nonconceptual. It is in this regard that traditional views of concepts differ most strikingly from the nondualist alternative endorsed here.

Kant's account begins by elaborating two of his epoch-making insights: first, that *judgments* are the fundamental form of awareness, so that concept use must be understood in terms of the contribution it makes to judging; second, that cognition and action are distinguished from their analogs in nonrational beings by their liability to certain sorts of *normative* assessment (see Chapter 1). Kant combined these insights with a *classificatory* theory of concepts, in terms of which he synthesized the teachings of his rationalist and empiricist predecessors. It is this aspect of his account that has been most influential in subsequent thought—becoming so much a matter of

course as to be almost invisible as a presupposition. Although it is based on important dimensions of ordinary concept use, the classificatory conception generalizes inappropriately as a result of running together substantially different phenomena.

For Kant, concepts provide only one of the two elements required for judgment. Concepts without intuitions are empty, and intuitions without concepts are blind.² Kant's theory is essentially dualistic in that the notion of the conceptual element in judgments is that of one of a *pair* of contrasting aspects. That neither is intelligible apart from its collaboration with the other is one reflection of his healthy emphasis on the primacy of judgment. It remains unclear, however, how much remains of the picture of judgment as the joint product of two distinct faculties if those faculties can be understood only by abstraction, that is, in terms of their contribution to the activity of judging. Insofar as Kant's embrace of both intellectual and sensible faculties is construed as his saying "You're both right" to Leibniz and Locke, his insistence on their mutual presupposition is bound to look like the bit where he takes it back. On the other hand, insofar as sense can be made of the notion of distinct contributions to judgment made by concepts and the unconceptualized given, the nature of their collaboration seems bound to remain mysterious. What sort of 'fit' is envisaged between concepts and intuitions, in virtue of which it is *correct* (or just possible?) to apply some, but not other, concepts to the manifold of (preconscious) representations with which intuition in some sense presents the understanding? How does intuition *constrain* the application of concepts? Kant's appeal to the schematization of the concepts just moves the bump in the rug around. The capacity for judgment, for applying rules to particular instances, subsuming intuitions under concepts, is something that in the end³ we must just accept *that* we have, without understanding just *what* we have. A distinction becomes a dualism when its components are distinguished in terms that makes their characteristic relations to one another ultimately unintelligible. (Descartes's dualism is, as always, the paradigm.)

Essential elements of Kant's dualistic conception of concepts are still with us today. They are the basis for the suspicion evinced by some (for instance Davidson) that talk of concepts inevitably commits us to a picture in which they play the role of *epistemological intermediaries*, which stand between us and the world we conceptualize and forever bring into question the very possibility of genuine cognitive access to what lies beyond them.⁴ To see why such suspicions are justified, and to bring out the contrast between dualistic and nondualistic conceptions of the conceptual, it is helpful to disentangle three different sorts of contrast between the conceptual element in thought and some nonconceptual element in thought, all of which are in play in Kant's usage. Each of these contrasts represents a genuine distinction, but as these distinctions are elaborated and run together in Kant's classificatory model of concepts, each becomes a dimension of an unworkable dualism. For

Kant, concepts contrast with intuitions first as *form* to *matter*, which they structure or organize. Second, they contrast with intuitions as *general* to *particular*. Finally, they contrast with intuitions as products of *spontaneity* or intellectual activity, as opposed to products of *receptivity*.⁵

In the first, the conceptual is distinguished from the *material*, that which provides *content*, as opposed to the form (more specifically the *normative* form or *rulishness*), which is the contribution of concepts. In the second, the conceptual is distinguished from the particular, as what *classifies* to what is *classified*. In the third, the conceptual is distinguished from what is imposed on us from without, as what *we do* as opposed to what is *done to* or imposed on us. It is the beginning of wisdom in reading the first Critique to distinguish the roles played in various arguments by these different distinctions. It is central to Kant's account that the three contrasts (though different) all line up together. Once this is questioned, a host of alternatives to his arguments present themselves. The lines of thought developed in this work support the conclusions that:

1. there are genuine distinctions underlying the contrasts Kant points to, but
2. far from coinciding, they are each independent of and orthogonal to the others, and
3. none of them is properly understood as distinguishing the conceptual from some nonconceptual element in judgment.

What a judgment expresses or makes explicit, its content, is conceptual all the way down.

The first idea is that of concepts as organizing something else. This can take many forms. It can be claimed that what is organized is *experience*,⁶ which is carved up by concepts, or alternatively lumped together by them. The material on which concepts work can be conceived of as perceptions or observations, sense data, or patterns of sensory stimulation. The concepts are supposed to be the source of *structure*, while something else provides the *content* or matter. Davidson has this picture in mind when he objects to the "scheme/content" dualism that he takes to be implicit in talk of alternative conceptual schemes.⁷ Concepts contrast with the unconceptualized matter that they conceptualize, which thereby provides content to the judgments that result. The worry inevitably raised by this picture is that unless its activity is entirely unnecessary, in conceptualizing the unconceptualized the understanding that is deploying the concepts must somehow *alter* what it works on and is therefore liable to the possibility of systematically *falsifying* that matter in rendering it digestible to the intellect.⁸ It should be admitted that it always remains pretty obscure what can be meant by either the form or the matter side of this opposition. (C. I. Lewis's heroic expository effort in *Mind and the World Order* is probably as clear a setting-out of this way of conceptualizing intuitions about concepts and intuitions as can be had.)

The second idea is that concepts are something *general*, something best expressed by the use of *predicates*. Along this dimension they contrast with nonconceptual *particularity*, as expressed by the use of (at least some kinds of) singular terms. The idea here is that predicates *classify* things, *say* something about them, as opposed to simply picking them out. This thought is the heir of Kant's treatment of intuitions as representations of particularity. The association of concepts with general terms rather than proper names is pervasive in the tradition. One important example is Frege,⁹ for whom concepts are functions from (sequences of) singular terms to truth-values, and so essentially things that can be *true* of the objects picked out by singular terms, by contrast to those *objects*, which concepts can be true of.

Finally, the third idea is that the *conceptual* order contrasts with something like the *causal* order, which constrains it. This distinction is the heir of Kant's distinction between judgments as the joint products of the activity ('spontaneity') of the intellect and the receptivity of the senses. According to this line of thought, whatever is conceptually articulated shows the effects of the mind working on it, whereas the nonmental world that thought is largely about is not in itself conceptually articulated. Because of special features of Kant's view, he could not put this contrast in terms of concepts versus causes (since talk of causation is for him already talk that betrays traces of the activity of the concept-mongering intellect).¹⁰ Nonetheless the tradition he inspired contrasts conceptually articulated expressions such as definite descriptions with those that are taken merely to register causal impingements—above all the uses of *demonstratives* that are so important in expressing the noninferential reports in virtue of which our concepts have empirical content. Kaplan's work is a prime example of contemporary versions of this distinction, as he worries about how to characterize the relation between the conceptual element in propositional contents, expressed by the use of predicates and descriptions, and the *nonconceptual, contextual, or causal* element, expressed by the use of indexicals.

In this contemporary form, Kant's distinction survives as the contrast between the *unrepeatable* character of indexical tokenings, reflecting their token-reflexive embeddedness in a causal context, and the *repeatable* concepts that are epitomized by definite descriptions.¹¹ Kant ran the second and third dualistic thoughts together—that is, the distinction between predicates and singular terms, on the one hand, and between repeatable and unrepeatable elements of thought, on the other, by systematically failing to distinguish between *representations of particularity* and *particular representations* (though elsewhere he is clear-headed about the distinction between representations of relations and relations among representations). In fact, however, singular terms, which represent particulars, are typically themselves as repeatable as predicates, while unrepeatable or token-reflexive indexical expressions can be of either grammatical category.

According to this broadly Kantian, dualistic, classificatory conception of

concepts, they function as epistemological intermediaries. They stand between the understanding mind and a world that is the source of their content or matter—a world composed of particulars that are grasped by means of general concepts and that imposes itself causally on a mind obliged somehow to conform itself to those causal impingements. As long as the conceptual is conceived in this way, Davidson is quite right to object to talk of conceptual schemes by means of which we render the world intelligible to and digestible by thought. But one of the lessons that ought to be drawn from the stories told here that this is not the only way to think about concepts. In particular, this broadly Kantian approach can be laid alongside another, inspired by Sellars, which avoids the dualistic understanding of each of the three distinctions that is characteristic of the Kantian one.

2. The Inferential Conception of Concepts Is Not Dualistic in Any of the These Ways

The approach developed here thinks of concepts to begin with as inferential roles (see Chapter 2). It treats a reliably differentially elicited response as *conceptually* classifying the stimulus to which it is keyed just in case that response occupies a position in an *inferentially* articulated space of claims that can be offered as, and stand in need of, reasons. In order for it to count as a conceptually contentful performance, that response must be able to serve as a premise for inferences to the applicability of further concepts. The particular content of a given concept is accordingly thought of as the content of an inferential commitment: roughly the commitment to the propriety of the inference from any of the appropriate *circumstances of application* of that concept to any of the appropriate *consequences of application* of the concept.¹² In this way even the *empirical* content some concepts have in virtue of their connection with *noninferential* circumstances of application in *perception*, and the *practical* content some concepts have in virtue of their connection with noninferential consequences of application in *action*, can be seen to be inferentially articulated.

It is essential to this inferential approach to concepts that the inferences in question are what Sellars calls *material* inferences. This is to say that their correctness involves the particular *contents* of the concepts invoked by their premises and conclusions; it is not underwritten purely by the *form* of those premises and conclusions. A paradigm is the inference from “*A* is to the East of *B*” to “*B* is to the West of *A*,” whose correctness expresses part of the *content* of the concepts *East* and *West*.

The first point to notice, then, is that thinking of conceptual contents as articulated by the material inferences that determine their role in giving and asking for reasons involves no contrast between concepts as *form* and something else as *matter* or content. The inferential role, which is the conceptual role, *is* the content. If one likes, one can say that on this conception the form

of that content is inferential. But the concept itself is identified with the particular constellation of material-inferential transitions the concept is involved in. This is not a structuring of something *else* that contrasts with the concept. The inferences materially relate one concept to other concepts, not to something of another kind. Thus the first of the Kantian dualisms, contrasting the conceptual and the material, is simply not involved in the inferential conception of nonlogical concepts.

It is possible, however, to go on to erect a superstructure of formal proprieties of inference on this base of material proprieties of inference (see 2.4.2 above). This sort of inferential articulation is an essential part of the expressive role of specifically *logical* vocabulary, by means of which we make explicit to ourselves the contents of our nonlogical concepts. So a distinction between matter and form is discerned and exploited by the inferential approach, though not in a form recognizable as distinguishing a conceptual from a nonconceptual element in judgment. Indeed, the same Fregean procedure of noting invariants under substitution that gives rise to the notion of formal proprieties of inference—an inference being valid in virtue of its form with respect to some distinguished vocabulary-kind *K* (paradigmatically logical vocabulary) just in case it is a materially good inference and cannot be turned into one that is not good by substituting non-*K* for non-*K* vocabulary—is what makes it possible to distinguish various formal categories of subsentential expressions, such as singular terms and predicates.

Only claims can literally function as premises and conclusions in inference; so only what is expressed by *sentences* can directly have an inferential role and so be in the most basic sense inferentially articulated. This is the version of Kant's insight concerning the primacy of judgments in cognition that survives into the inferential conception of concepts. But subsentential expressions can nonetheless be conceptually articulated according to that conception—their occurrence in a sentence can have an *indirect* inferential significance. For substitution of one subsentential expression for another in a sentence can either result in preserving the goodness of inferences in which the sentence is involved or fail to preserve it. In this way subsentential expressions can be sorted into indirect inferential equivalence classes, by noting direct inferential invariances of the sentences that result from their substitution one for another. Thus the *inferential* approach to the conceptual articulation of sentences can be extended *substitutionally* to include the conceptual articulation of *subsentential* expressions.

When this is done, the subsentential categories of singular terms and predicates can be distinguished by the different patterns of *substitution inferences* in which they are involved. In particular, singular terms are distinguished by the de jure *symmetric* significance that their occurrence in a sentence has for substitution inferences involving it. For example, if the inference from "Benjamin Franklin spoke French" to "The popularizer of lightning rods spoke French" is a good one, then so is the converse inference.

By contrast, all predicates are involved in some *asymmetric* substitution inferences. For instance, the inference from "Benjamin Franklin could dance" to "Benjamin Franklin could move," is a good one, but the converse inference need not be. On the basis of such differences in substitution-inferential behavior, the difference between the sort of conceptual role played by singular terms and that played by predicates can be characterized (see Chapter 6).

This means that the second of the Kantian dualisms, though based on a genuine categorial distinction, also does not set off concepts as conceived by the material-inferential model. There is no restriction of the conceptual to the general, as expressed by predicates, in contrast to the particularity invoked by singular terms. Singular terms have an inferential role, represented by the set of terms intersubstitutable for them, just as predicates do. The difference between them is a formal difference of symmetric versus asymmetric substitution inference. It is not a difference that involves contrasting the conceptual as embodied in predicates, which express generalities, with something else, embodied in singular terms, which express particularity. Singular terms and predicates, the particular and the general aspects of claims, are equally (though not identically) inferentially articulated, and so equally conceptually contentful. *Particularity is as much a conceptual matter as generality*, on this conception. Thus the second dualism gets no grip on the inferential rendering of conceptual contentfulness, once that account has been extended to the subsentential level by invoking the notion of substitution inferences.

The third of the Kantian dualisms contrasts the conceptual, as the product of cognitive activity, with the nonconceptual impingement on cognitive receptivity in virtue of which that cognitive activity is constrained. Outside the strictures of Kant's own system, we can think of this as the conceptual/causal contrast, in which the application of concepts is constrained by the causal order, thought of as not itself conceptually articulated. The point of contact between the conceptual order and the causal order, according to this conception, takes place in *deixis*, where something is indicated without being characterized. In grasping this conception it is helpful to focus on the use of deictic expressions in noninferential reports, such as "This is red." For it is in such reports that the world most directly imposes itself on suitably trained concept-mongers, who find themselves passively acknowledging empirically contentful commitments.

Once again, it ought not to be denied that this sort of receptivity is essential to our empirical knowledge and that it ought to be distinguished from other, more spontaneous applications of concepts, for instance in purely inferential theorizing. Yet according to the inferentialist conception, unrepeatably deictic tokenings—for instance particular uses of 'this'—are fully conceptually articulated. Indeed, were they not, they could serve no cognitive purpose. To see how occurrences of deictic tokenings are assigned an inferential significance, and so taken to be conceptually contentful, is accordingly

to see that the third of the broadly Kantian dualisms—contrasting conceptual constraint with causal constraint on the application of concepts—fails to get a grip on the inferential conception of the conceptual. Just as the idea of *inference* needed to be supplemented by that of *substitution* in order to be brought to bear on subsentential expressions, so the idea of substitution-inferential significance needs to be supplemented by that of *anaphora* in order to be brought to bear on *unrepeatable* tokenings of subsentential expressions, rather than just on their repeatable types. To take one expression to be anaphorically dependent on another is to take it as *inheriting* its substitution-inferential role from the tokening that is its anaphoric antecedent.¹³ If you say, “*That* is a porcupine” and I pick up that premise and conclude, “[so] *it* is a vertebrate,” the truth of the conclusion I have drawn is to be settled (according to an interpreter) by what substitutions are appropriate (according to the interpreter) for the demonstrative tokening that serves as the antecedent for my anaphorically dependent tokening. If (according to the interpreter) what the first speaker referred to by ‘*that*’¹⁴ is the most cunning wooden replica of a porcupine in the room, then since this identity claim is to be understood as an intersubstitution license, I have unwittingly claimed of a cunning wooden replica of a porcupine that it is a vertebrate, and what I said is false.

Anaphora permits the formation of *chains* of tokenings, anchored by antecedents that can be deictic and therefore strictly unrepeatable. *These chains of unrepeatables are themselves repeatables and play the same role in substitution inferences that sets of cotypical tokenings play for repeatable expressions such as proper names and definite descriptions.* It is by means of anaphora, then, that substitution-inferential potential can be inherited by one expression from an unrepeatable tokening. In virtue of this mechanism, unrepeatable tokenings such as uses of demonstratives become available for service as premises in inference. In this way they acquire an *inferential* significance and so can be understood as expressing a *conceptual* content. This function of anaphora is essential to the existence of deictic expressions. For without the capacity to be picked up anaphorically, and so to have some inferential significance, deictic expressions would just be noises wrought from us by exposure to things—rather than genuinely linguistic expressions that can be used to *say* something. Thus anaphora is more basic than deixis, for there can be languages that have anaphoric mechanisms but no deictic ones, while there cannot in principle be languages with deictic mechanisms but no anaphoric ones (see Chapter 7).

In any case, with anaphora available to bring deictic expressions into substitution inferences, such expressions have indirect inferential roles, and so conceptual contents. There is no contrast between expressions like definite descriptions and those like demonstratives over the issue of whether or not they are inferentially articulated and so conceptually contentful. The structure of their contents is specifically different, for the latter are involved

in substitution inferences via anaphoric chains of unrepeatable tokenings potentially of a variety of types, while the former¹⁵ are involved in substitution inferences via sets of repeatable, because cotypical, tokenings. But this difference plays a role analogous to that between symmetric and asymmetric substitution-inferential significances in distinguishing singular terms from predicates. In neither case is a contrast underwritten between the conceptual and something else, whether particular or causally responsive. Deictic tokenings play a role in the causal order, but they are not for that reason not also conceptually articulated.

Thus *none* of the Kantian dualisms—which contrast the conceptual as formal with the material, the conceptual as general with the particular, and the conceptual as spontaneous activity with the constraint of causes—applies to the inferential conception of concepts. That conception does not involve any commitment to a dualism of conceptual scheme and something else that it structures, classifies, or is about. So Davidson's proper fastidiousness about scheme/content dualisms and epistemological intermediaries ought not to motivate a rejection of appeals to concepts as here conceived. Concepts conceived as inferential roles of expressions do not serve as epistemological intermediaries, standing between us and what is conceptualized by them. This is not because there is no causal order consisting of particulars, interaction with which supplies the material for thought. It is rather because all of these elements are themselves conceived as thoroughly conceptual, not as contrasting with the conceptual.

The conception of concepts as inferentially articulated permits a picture of thought and of the world that thought is about as *equally*, and in the favored cases *identically*, conceptually articulated. *Facts* are just true claims.¹⁶ Facts, like other claims, are conceptually articulated by their inferential and incompatibility relations to other claims. It is a feature of the conceptual articulation of claims, and hence of facts, that they are about particular objects.¹⁷ (Indeed, the fact that we are accustomed to saying that facts, like claims, are *about* objects, rather than that they somehow *consist* of objects, is evidence for the correctness of identifying facts with true claims.) It is these facts and the properties and related objects they involve that are cited as stimuli by interpreters who are specifying the reliable differential responsive dispositions in which the contents of empirical contents originate. These noninferential dispositions (the locus of our empirical receptivity) accordingly do not constitute the interface between what is conceptually articulated and what is not, but merely one of the necessary conditions for a conceptually articulated grasp of a conceptually articulated world—the world consisting of everything that is the case, all the facts, and the objects they are about.

In this way a story has been told about how the three nonconceptual poles of Kant's tripartite division of the conceptual and the nonconceptual contributions to the contents of judgments ought to be incorporated within the

conceptual realm. An approach to concepts that moves beyond exclusive focus on classification to include inferential connections among concepts as essential to their identity and individuation:

1. incorporates *content* by employing a notion of *material* proprieties of *inference*,
2. incorporates *particularity* by distinguishing between the symmetric role of singular terms in *substitution* inferences and the asymmetric role of predicates in substitution inferences, and
3. incorporates the deictic *unrepeatability* by which causal context affects conceptual content by explaining how *anaphoric* chains initiated by unrepeatable tokenings function as type-repeatables in substitution inferences.

The key theoretical concepts used to characterize the articulation of conceptual roles are *material inference*, *substitution*, and *anaphora*, so this can be called the ISA approach to semantics.

II. NORMS AND PRACTICES

1. *The Normative and the Factual*

This inferential semantics is embedded in a normative pragmatics. Material proprieties of inference are understood as norms implicit in social practices that qualify as *discursive* inasmuch as they involve treating some performances as having the significance of assertions. Such inferentially articulated practices confer propositional contents on statuses, attitudes, and performances that are suitably caught up in them (since for an expression to have a certain conceptual content just is for its use to be governed by a corresponding set of norms). In this way the semantically primitive notion of material proprieties of inference is explained in the pragmatics—in the account of linguistic practice.

Such a pragmatic theory of the relations between meaning and use raises issues about the status of implicit practical norms. Does not talk of deontic statuses as instituted by social practices involve a residual dualism? When the orienting commitment to the normative character of discursive practice was first introduced and motivated, in Chapter 1, this insight of Kant's was presented in the context of a shift from a broadly Cartesian dualism of the mental and the physical to a broadly Kantian dualism of the normative and the factual. In these crude initial terms, Descartes's opposition of two kinds of descriptive properties (corresponding to ontological kinds of substances) was contrasted with a deeper opposition between descriptive and prescriptive attitudes—between attributing properties and attributing proprieties. Thus even if Kant's *semantic* dualisms have been overcome by the ISA approach, it would appear that the pragmatics in which that semantics is embedded

incorporates an overarching dualism that distinguishes the normative and the nonnormative. How should the relations between these categories be understood?

The deontic scorekeeping idiom acknowledges a distinction between normative and nonnormative claims, explained in terms of their different roles in practical reasoning, but that *distinction* does not underwrite a *dualism* of norm and fact. Indeed, looked at more carefully, neither does Kant's. (He is large; he contains multitudes.) The initial discussion of replacing one dualism with another (in Chapter 1, Section II) was only a temporary expository device, discarded in favor of a more nuanced treatment (in Section IV) once its purpose was served. For Kant, rules are the form of the normative as such. To call something 'necessary' is to say that it happens according to a rule, and everything that happens in nature, no less than everything done by humans, is subject to necessity in this sense. Concepts are rules, and concepts express natural necessity as well as moral necessity. So according to him there is strictly no nonnormative realm—no realm where concepts do not apply. Kant's fundamental innovation is best understood to consist in his employment of a normative metalanguage in specifying *both* what merely happens and what is done.

Of course he does distinguish between the realm of regularity and the realm of responsibility. This is the distinction between that to which concepts apply and those who apply concepts—between that which can acknowledge rules only implicitly by obedience (by having concepts be applicable to it) and those who can acknowledge them explicitly by the use of concepts (by applying concepts). It is only rules as explicitly acknowledged that can be both binding and *disobeyed*,¹⁸ and it is the capacity for such acknowledgment—acting not just according to rules but according to *conceptions* of rules—that institutes distinctively normative statuses such as duty and responsibility. The applicable distinction is accordingly not between the normative and the nonnormative but between what can adopt explicitly normative *attitudes* and what cannot. Only we discursive (that is concept-mongering) creatures can *take* ourselves and others to be bound by the norms that are our concepts.¹⁹

This is the idea that is followed out in the deontic scorekeeping pragmatics presented here. The idiom in which the account of discursive commitment is expressed is normative throughout. Propositional contents are understood in terms of their explanatory role in specifying proprieties of claiming, judging, and inferring—in general, in terms of the role they play in the game of giving and asking for reasons. What it is for something to state or express a fact is explained in normative terms, and what it is for something to be stated or expressed is explained in turn by appeal to that practice. So what it is to be a fact—that is, true claim—is explained in normative terms. It is explained phenomenally, by appeal to the practice of fact-stating, which comprises the practical attitudes of taking a performance to be the

stating of a fact and purporting to state a fact by producing a performance. In this order of explanation, normative notions such as commitment and entitlement—which articulate implicit proprieties of practice—are more fundamental than the nonnormative properties they enable discursive practitioners to express explicitly.

However, only some of the vocabulary on which conceptual content is conferred by implicitly normative discursive practice plays the expressive role of making explicit specifically normative attitudes—for instance the attribution or acknowledgment of commitments. As explained in Chapter 4, the distinctive function of normative vocabulary is to express endorsement of patterns of practical reasoning—that is, in the first-person case, reasoning that leads from doxastic to practical commitments (presystematically: from beliefs to intentions). Social practices are implicitly normative in a way that mere behavioral regularities are not. Put phenomenally, that is to say that what a scorekeeper or interpreter has attributed counts as a *practice* in this sense (rather than a behavioral regularity or disposition) only if it is specified in explicitly normative terms—in terms of what, according to the practice, it is *correct*, or *proper* to do, what one *ought* to do, what one becomes *committed* or *entitled* to by a certain sort of performance, and so on. The account of practical reasoning explains in deontic scorekeeping terms how words have to be used in order to mean what such terms as ‘correct’, ‘ought’, and ‘committed’ do. By doing that, it makes sense of the distinction between normative statuses and attitudes, on the one hand, and nonnormative states and dispositions, on the other.

Explicitly normative vocabulary can be used to make claims (for example “Bank employees are obliged to wear neckties,” “One ought not to torture helpless strangers”). Those claims can be taken-true, can be put forward as, or purport to be, true. Since facts are just true claims (in the sense of what is *claimed*, not the *claiming* of it), this means that norm-explicating vocabulary is in the fact-stating line of business. That is, corresponding to the distinction between normative and nonnormative vocabulary is a distinction between normative and nonnormative facts. (Indeed, this ontologically relaxed approach to facts finds nothing mysterious about negative, conditional, or modal facts, facts about the self-identity of objects, or in general facts expressed by any sort of declarative sentence at all.) In this way the normative is picked out as a subregion of the factual.

To revert to the previous point, however, this is a distinction made *within* the encompassing normative metalanguage in which the deontic scorekeeping roles characteristic of normative and nonnormative vocabulary are specified. The distinction between normative and nonnormative vocabulary, claims, and facts is itself drawn in normative terms. In this sense, the story is one in which it is norms all the way down—a Kantian story (on the pragmatic, rather than the semantic side).²⁰ Far from opposing one another, the realms of fact and norm mutually include one another: fact-stating talk

is explained in normative terms, and normative facts emerge as one kind of fact among others. The common deontic scorekeeping vocabulary in which both are specified and explained ensures that the distinction between normative and nonnormative facts neither evanesces nor threatens to assume the proportions of an ultimately unintelligible dualism.

2. *Where Do Norms Come From?*

The story told here is Kantian not only in that it is told in normative terms but also in the pride of place it gives to normative *attitudes* in explaining how we are both distinguished from and related to the non-us that surrounds us. On the one hand, such practical attitudes—taking or treating a performance as correct, attributing or acknowledging a commitment—have been appealed to in explaining our relations in perception and action to the causal order of nonnormative facts that we inhabit cognitively and practically. On the other hand, they have been appealed to in explaining where discursive norms come from—how sapience could have arisen out of the primordial nondiscursive ooze of mere sentience. For it has been claimed not just that we discursive beings are creatures of norms but also that norms are in some sense creatures of ours—specifically, that discursive deontic statuses are *instituted* by the practices that govern scorekeeping with deontic attitudes.

Norms (in the sense of normative statuses) are not objects in the causal order. Natural science, eschewing categories of social practice, will never run across *commitments* in its cataloging of the furniture of the world; they are not by themselves causally efficacious—any more than strikes or outs are in baseball. Nonetheless, according to the account presented here, there are norms, and their existence is neither supernatural nor mysterious.²¹ Normative statuses are domesticated by being understood in terms of normative attitudes, which *are* in the causal order. What is causally efficacious is our practically taking or treating ourselves and each other as having commitments (acknowledging and attributing commitments)—just as what is causally efficacious is umpires and players dealing with each other in a way that can be described as taking the score to include so many strikes and outs.²²

It must then be asked how such an apparently reductive story about norms as instituted by social practices can be understood to be compatible with an insistence on the *irreducibly* normative character of the metalanguage in which norm-instituting social practices are specified. Here is the short answer: The work done by talk of deontic statuses cannot be done by talk of deontic attitudes *actually* adopted or relinquished, nor of *regularities* exhibited by such adopting and relinquishing, nor of *dispositions* to adopt and relinquish such attitudes. Talk of deontic statuses can in general be traded in only for talk of *proprieties* governing the adoption and alteration of deontic attitudes—proprieties implicit in social scorekeeping practices.

The crucial inferential articulation of discursive commitments consists in part in the fact that unacknowledged commitments can be (taken by other scorekeepers to be) undertaken consequentially, by acknowledging commitments to claims that (according to those scorekeepers) entail them. So according to the attributions (normative attitudes) of another, my commitments (normative statuses) outrun those I acknowledge (normative attitudes). In this way the social articulation of deontic scorekeeping attitudes is essential to the inferential (and hence discursive) articulation of the contents of the commitments they address. But this social articulation of scorekeeping practice is essentially normative in force. That I acknowledge commitment to *p* does not (according to the scorekeeper) mean that I *do* or *will* acknowledge commitment to its consequence *q*, only that I *ought* to—that I am, whether I realize it or not, *committed* to *q*.

It was shown at the end of the last chapter that the contents of ordinary empirical claims—objective proprieties governing the application of concepts—are not equivalent to the contents of any claims about who is committed to what. The implicit scorekeeping attitudes expressed by this difference in explicit contents accordingly distinguish what follows from *p* from what I or anyone *takes* to follow from *p*. What follows from *p* cannot be identified with how I or anyone *actually* keeps score; it is rather to be identified with a feature of *correct* scorekeeping (for it depends on what else is true, not on what anyone *takes* to be true). Conceptual contents on this inferential conception—and so what interlocutors are really committed to by using particular expressions (performing particular speech acts)—codify *proprieties* of scorekeeping. Any scorekeeper who attributes a conceptually contentful commitment may get these wrong, just as anyone who acknowledges or otherwise acquires such a commitment may get them wrong. Talk of inferentially articulated contents is a way of talking about implicit norms governing deontic scorekeeping practice;²³ this is the cash value of the claim that conceptual contents are *conferred* by such practice. But since commitments must be individuated at least as finely as their contents, if those contents are determined only by how it is *correct* to acquire and alter deontic attitudes, the commitments themselves must be understood as instituted also by proprieties of scorekeeping, rather than by actual scorekeeping. The scorekeeping account incorporates a phenomenalist approach to norms, but it is a *normative* phenomenism, explaining having a certain normative status in effect as being *properly* taken to have it.

At this point it can easily look as though the account of normative statuses as instituted by social practices is marching around in an unproductive circle (at best, unilluminating; at worst viciously circular and incoherent). For clearly the prior question arises once more: What is the relation between normative specifications of practices and nonnormative specifications of behavior? Actual scorekeeping, the adoption and alteration of practical normative attitudes (acknowledgments and attributions of deontic statuses),

consists of causally efficacious events and dispositions. If normative statuses could be understood as instituted by actual attitudes of acknowledging and attributing them, then the use of normative vocabulary specifying proprieties, commitments, and entitlements would straightforwardly supervene on the use of nonnormative vocabulary specifying performances and performative dispositions and regularities. If, however, as has been claimed, the institution of discursive deontic statuses should be understood rather in terms of the implicit practical *proprieties* governing such scorekeeping—not how the score is actually kept, but how according to the implicitly normative scorekeeping practices it *ought* to be kept, how scorekeepers are *obliged* or *committed* to adopt and alter their deontic attitudes rather than how they *actually* do—then the source and status of these norm-instituting proprieties of scorekeeping practice must be inquired into.

3. Interpretation

Proprieties are normative statuses—the status a performance has as *correct* or *incorrect* according to a rule or practice. This is so even when the practice whose proprieties are in question is itself a deontic scorekeeping practice. In that case what is being evaluated as proper or improper is the acquisition and alteration of deontic attitudes—that is the acknowledgment and attribution of further deontic statuses (commitments and entitlements). The (normative) phenomenalist strategy that has been pursued throughout is to understand normative statuses in terms of normative attitudes—in terms of (proprieties of) *taking* to be correct or incorrect. This strategy dictates two questions concerning proprieties of scorekeeping practice. First (apropos of phenomenism about norms), what must one be doing in order to count as taking a community to be engaging in implicitly normative social practices—in particular in deontic-status-instituting, conceptual-content-conferring discursive scorekeeping practices? Second (apropos of its being a normative phenomenism), what is it about the actual performances, dispositions, and regularities exhibited by an interacting group of sentient creatures that makes it correct or appropriate to adopt that attitude—to interpret their behavior by attributing those implicitly normative discursive practices?

The first question can be addressed by considering the different sorts of intentional stance that interpreters can adopt, according to the story told here. The central task of the pragmatic part of this project (the account of discursive practice) has been to introduce the model of *deontic scorekeeping*. Keeping discursive deontic score by attributing inferentially articulated deontic statuses—propositionally contentful commitments and entitlements to those commitments—is treating the one so interpreted as being in the game of giving and asking for reasons. Social practices are *linguistic* practices when interlocutors take up the *discursive scorekeeping stance* toward one another. Adopting this stance is (implicitly, or in practice) taking or treating others as

producers and consumers of propositionally contentful speech acts. Performances count as propositionally contentful in virtue of their relation to a core class of speech acts that have the pragmatic significance of *claims* or *assertions*.

Assigning this sort of significance to performances is treating them as making *explicit* the adoption of a normative status—that is, acknowledging (undertaking) a doxastic commitment by *saying* what one is committed to. Keeping discursive score on others is adopting deontic attitudes—that is, attributing discursive commitments by implicitly or in practice taking or treating another *as* committed. Such scorekeeping (and so linguistic practice generally) does *not* require that one be able *explicitly* to *attribute* deontic statuses—to *say* (assert) *that* someone is committed to the claim that *p*. The logical locutions whose expressive role is to make the adoption of such pragmatic attitudes explicit in the form of claimable contents—propositional-attitude-ascribing vocabulary such as the regimented “. . . is committed to the claim that . . .” or its vernacular correlate “. . . believes that . . .”—form an optional superstratum whose expressive role can be understood in terms of what is implicit in ground-level linguistic practice, but which is not required for, or presupposed by, such practice.

The production and consumption of speech acts of which participants in these fundamental discursive practices are capable accordingly differ as to whether the adoption of deontic attitudes (toward normative statuses) they involve is explicit or implicit. They can *explicitly acknowledge* (and so undertake) discursive commitments, in their assertional performances, but only *implicitly attribute* them, in their scorekeeping practice. Since acknowledging a commitment (the basic sort of undertaking or acquisition of that deontic status) is producing (or being disposed to produce) performances whose pragmatic significance is to make it appropriate for scorekeepers to attribute that commitment, to take someone to be a producer of speech acts is implicitly to take that practitioner to be also a consumer of them—a scorekeeper. Givers of reasons must be able to understand what it is to give a reason. As Davidson says: “One cannot be a thinker unless one is an interpreter of the speech of others.”²⁴

Although performances cannot be accorded the significance of speech acts without implicitly treating the performer as a discursive scorekeeper, it is possible for those who are discursive scorekeepers to attribute a derivative sort of propositionally contentful discursive status and attitude to nonlinguistic creatures. This is adopting the *simple* intentional stance of interpreting something as a simple or practical intentional system. When this stance is adopted, the interpreter keeps a simplified sort of deontic score, by attributing propositionally contentful commitments, both doxastic and practical, which the subject is taken to acknowledge implicitly in its behavior. Its performances, dispositions, and behavioral regularities can be made intelligible by attributing sample pieces of practical reasoning, in the way Dennett

has described so well. The scorekeeping involved is simplified in that adopting the simple intentional stance does not involve attributing speech acts; it does not involve even implicitly treating the system in question as itself able to keep score (attribute, not just acknowledge deontic statuses); hence it does not involve treating it as a participant in the essentially social and linguistic game of giving and asking for reasons.

Discursive scorekeepers, participants in full-blooded linguistic practices, do two sorts of things that such simple, nonlinguistic intentional systems cannot: institute deontic statuses and confer conceptual contents. On the pragmatic side, both social flavors of deontic attitude—acknowledging and attributing—are needed to institute deontic statuses; reference to practical grasp of the possibility of *attributing* them is required to make sense of what is *acknowledged* as being inferentially articulated *commitments*. In the case of simple intentional systems, that essential pragmatic ingredient is supplied only by the interpreter, rather than attributed to the one being interpreted. On the semantic side, the social-perspectival dimension of inferential articulation is required to make sense of what states, attitudes, and performances exhibit as genuinely *propositional*, which includes having objectively *representational* conceptual content (see Chapter 8). In the case of simple intentional systems, that essential semantic ingredient is supplied only by the interpreter, rather than being attributed to the one interpreted. So the intentionality attributed by adopting this sort of stance is doubly derivative. On the side of pragmatics, the socially and inferentially articulated norms are derivative from the scorekeeping practices of the interpreter. As a result, on the side of semantics, the propositional and other conceptual contents employed to measure and systematize its behavior cannot be funded out of that behavior itself.²⁵

By contrast, if one attributes genuinely linguistic practices to a community—takes its members to adopt the discursive scorekeeping stance to one another, and so to accord some performances the significance of speech acts, in particular assertional ones—one thereby takes them to exhibit *original* intentionality. The social practices one interprets them as engaging in are sufficient by themselves to institute inferentially articulated deontic statuses and so to confer genuinely conceptual contents. Describing the model of inferentially articulated deontic scorekeeping social practices is specifying in detail what one must *take* the members of a community to be doing in order for it to be *talking*—giving and asking for *reasons*, *making* their words and performances mean something by *their taking* them to mean something²⁶—that one is thereby taking them to be doing. In short, the model specifies what structure an interpretation of the activities of a community must have in order for it to count as attributing original intentionality to that community—taking it as instituting socially and inferentially articulated deontic statuses and so conferring genuinely propositional conceptual content on them. This is adopting a further sort of stance.

So the difference between derivative simple and original discursive intentionality is presented in terms of the difference between two stances or forms of interpretation—in terms of the difference between the attitudes adopted in *attributing* them. The difference between these sorts of intentionality is not that one is construed in methodologically phenomenalist terms and the other is not. In keeping with the stance stance, this account is phenomenalist about both. The difference is that what one attributes in the case of genuinely discursive intentionality is (taken to be) autonomous in a way that what one attributes in the case of simple or practical intentionality is not.²⁷

4. *Semantic Externalism and the Attribution of Original Intentionality*

Interpreting a community as exhibiting original intentionality is taking its members to adopt the discursive scorekeeping stance toward each other. The content-conferring norms and proprieties that an interpreter who attributes discursive scorekeeping practices takes to be implicit in them have a number of important structural features. Central among them is the fact that the conceptual norms implicit in the practices attributed to a community outrun the nonnormatively specifiable behavioral discriminations members of that community are disposed to make. For this reason, conceptual norms can be understood as *objective*, and so as binding alike on all members of a discursive community, regardless of their particular attitudes. This feature of attributions of linguistic practices secures the sense in which concepts and the commitments they involve concerning appropriate circumstances and consequences of application can be understood to be *shared*, in spite of the many differences of attitude that correspond to the different scorekeeping perspectives of the discursive practitioners who keep track of each other's statuses. This normative surplus of practice (as attributed by an interpreter) over behavior (nonnormatively specified) is also what is appealed to in responding to the issue raised by the possibility of *gerrymandering* (introduced above in 1.3.5)—the problem of what privileges one of the many ways of projecting from actual applications of concepts (and regularities and dispositions regarding such performances) commitments regarding cases that have not arisen for practical adjudication.

The reason the conceptual contents conferred by the discursive scorekeeping practices a community is interpreted as engaging in can outrun the community's capacity to apply them correctly and to appreciate the correct consequences of their application is the empirical and practical *solidity* or concreteness of those practices. The assertible contents a discursive interpretation takes to be conferred by communal deontic scorekeeping practices are inferentially articulated, but they are not merely placeholders in abstract, purely formal, relational structures—hollow shells waiting to be filled up by supplying actual facts and objects that somehow 'fit' them. For the content-

conferring practices do not relate the deontic statuses that bear those contents only to other deontic statuses. Discursive practice comprises noninferential entries and exits as well, and these (according to the interpreter attributing those practices) relate contentful doxastic and practical commitments to the worldly states of affairs that properly elicit acknowledgments of those commitments and are properly elicited by such acknowledgments, respectively. Standard discursive practices—those that encompass both empirical and practical dimensions—are solid (even lumpy), in that they involve actual objects and states of affairs, as well as the deontic statuses in terms of which score is kept.²⁸

In such practices, the actual causal provenance or consequences of a deontic attitude—and not just the proprieties that connect its adoption to the adoption of other deontic attitudes—can matter (according to the external intentional interpreter attributing the content-conferring practices) for the content of the status it is an attitude toward. So an interpretation of this sort takes it that what an interlocutor who performs a certain speech act is committed to thereby, according to the practices of the relevant community, can depend on how things are in the nonlinguistic world. The interpreter takes it that the solid, corporeal communal practices determine what is being talked about (whether or not any scorekeepers in the community realize it), for those practices incorporate it. And the interpreter also takes it that what is being talked about determines what it is correct to say and infer, including practically (whether or not any scorekeepers in the community attribute the right claims and consequences). Interpretations that attribute original intentionality are accordingly *semantically externalist* in Davidson's sense.²⁹ This is part of what was called above (8.5.6) 'tactile Fregeanism': our practice puts us in touch with facts and the concepts that articulate them—we grasp them. But what we grasp by our practice extends beyond the part we have immediate contact with (its handles, as it were); that is why what we grasp is not transparent to us, why we can be wrong even about its individuation. How the world really is determines what we have gotten a hold of; but even though for that reason we do not know all the details about it, we still genuinely grasp *it*.

In this way the proprieties governing the application of a community's concepts are in part determined (according to the interpreter) by the actual properties of and facts concerning the things the linguistic practitioners are perceiving, acting on, and so talking about—which are just features of *their* practice (according to the interpreter). How the things and properties they are talking about actually are determines the correctness of the commitments of all community members alike. They are all bound by the *same* conceptual norms, regardless of the differences in collateral commitments that make particular claims have different inferential significances for different scorekeepers. According to the practices the interpreter takes them to be engaging in, they *share* a common set of concepts, which determines how the attitudes of those who keep score on each other are answerable to the facts.

When concrete discursive practices (including perceptual reporting and intentional agency) are ascribed to a community, the states of affairs that properly noninferentially elicit the acknowledgment of doxastic commitments and those that are properly noninferentially elicited by the acknowledgment of practical commitments are specified in the interpreter's own language. For instance, in assessing the extent to which the claims made by various community members do express facts, and so are correct uses of their concepts, the interpreter compares the commitments he or she attributes to them to those the interpreter undertakes—and similarly for assessments of their reliability as perceivers and agents. Semantic externalism is perspectival externalism.

To treat those interpreted as linguistic practitioners who use particular concepts is to treat them as bound by proprieties that project beyond their actual behavior and dispositions. The interpreter uses the norms implicit in his or her own concepts in specifying how the conceptual norms that bind the community being interpreted extend beyond the practitioners' actual capacity to apply them correctly. All the resources of the interpreter's home language are available in distinguishing one such set of proprieties from another; taking the interpreted interlocutors to have bound themselves by even a slightly different set of proprieties would be offering a different interpretation, attributing a different set of practices. The general point is that while normative interpretation of a community as engaged in one set of practices rather than another is underdetermined by nonnormatively specified actual behavior, regularities of behavior, and behavioral dispositions, *relative to such an interpretation*, concepts nevertheless are objective, shared, and unambiguously projectable.

5. *Sharing Inferentially Individuated Concepts*

It has been acknowledged throughout this exposition that an inferential conception of concepts raises *prima facie* difficulties for understanding what is involved in communication between individuals with different repertoires of commitments. The inferential significances of utterances of the same sentence produced by different performers are different—even where anaphoric and indexical phenomena are not in play. For their different collateral commitments make available different auxiliary hypotheses; hence what consequential commitments the performer undertakes by producing those performances and what would entitle their utterer to them (according to the scorekeeper who attributes the collateral commitments) are different. So something special needs to be said about the sense in which interlocutors with different collateral commitments can nonetheless be said to be able to make the same claims and express the same inferentially articulated concepts. It is worth rehearsing briefly the features of the discursive scorekeeping model that are appealed to in providing such an account.

What is from many points of view the most natural way out of this difficulty is not the path taken here. The most straightforward approach would be to adopt an inegalitarian attitude toward the different inferences a concept is involved in. A privileged class of inferences would be distinguished, which are taken to be constitutive of the concept, while the rest are accorded a secondary status as turning out to be correct ways of using the concept so constituted. There is an undeniable intuitive basis for such a distinction: The inferences from "This tractor is completely green" to "This tractor is not completely red" and from "This cloth is scarlet" to "This cloth is red," for instance, have a different status from the inferences from "This tractor is completely green" to "This tractor is made by John Deere" or from "The apple in the box is a ripe Winesap" to "The apple in the box is red." The correctness of the first inference plausibly is taken to be part of the concepts *green* and *red*, while the correctness of the second sort is equally plausibly taken to be just a matter of empirical facts about John Deere tractors and ripe Winesap apples—inferences whose correctness involves the concepts *red* and *green* without in any way constituting them.

Quine, of course, argues that one way of construing the sort of concept- (or meaning-) constitutive privilege that distinguishes the first class is defective because it does not correspond to the sort of difference in the use of the words (the practical status of the inferences) that the theory behind it entails.³⁰ There do not seem to be any inferential connections that are unrevisable in principle, immune to being undermined by suitable empirical evidence, and so a priori for those who grasp the concepts involved. But this is not to say that no pragmatic sense can be made of the intuitive difference in status between two sorts of inferences instanced above. Sellars,³¹ for instance, does not take all the materially good inferences involving a concept to be essential to it.³² He picks out the privileged concept-constitutive inferential connections as those that support *counterfactual* reasoning, and so count as having *nomological* force. This is a real practical difference; this way of drawing the line does not fall afoul of Quine's strictures, for it by no means follows that these conceptual matters are a priori—we need to investigate the world to find out what the laws are, as for any other facts. Since the laws involved are not a priori, unrevisable, or immune to factual evidence, this is not a version of analyticity. According to this view, not only claims but concepts can be correct or incorrect, depending upon whether the inferences they incorporate correspond to actual laws.

The difference between inferential connections among concepts that are counterfactually robust and those that are not is an important one, and this fact accounts for the felt difference between the two sorts of inferences mentioned above. Nonetheless, nothing is made of it here. This is partly because the notion of nomologicality and counterfactual reasoning, important though it is in other contexts, has not been reconstructed in discursive scorekeeping terms as part of this project (though the key notion required,

that of the *incompatibility* of claims, has been given a pragmatic interpretation). But neither this nor any other way of picking out a privileged subclass of concept-constitutive inferences has been appealed to in individuating concepts here, for two other reasons.

First, mastery of a special subset of distinguished inferences (for instance, the counterfactually robust ones) is not in general sufficient for grasp of a concept. For such grasp requires that one be hooked up to the *function* that takes as its argument repertoires of concomitant commitments available as auxiliary hypotheses and yields inferential significances as its values. Carrying on a conversation involves being able to move from perspective to perspective, appreciating the significance a remark would have for various interlocutors. (More is said about this below, in connection with the representational dimension of discourse.) The effect that various auxiliary hypotheses have on the inferential significance of a claim relative to a particular doxastic context cannot be determined just from the privileged inferences it is involved in (for instance, the counterfactually robust ones), unless it is assumed that the repertoire in question contains conditionals corresponding to all the other materially good inferences (for example from the ripeness of Winesap apples to their redness). Assuming that is contrary to the spirit of this enterprise: it depends on the formalist view of inference, which sees enthymematically suppressed conditionals behind every material propriety of inference. In particular, such a view would have the consequence that communities that do not yet have the expressive resources of logical vocabulary such as the conditional were precluded for that reason from counting as employing nonlogical concepts such as *red*.

The second reason that the inegalitarian attitude toward inferences is not taken in individuating concepts is that no matter how the privilege distinguishing some supposedly uniquely concept-constitutive inferences is construed (as counterfactual robustness or otherwise), endorsement even of these privileged inferences can still vary from perspective to perspective. There can be different views about what the laws of nature are, for instance, just as there can be differences about the colors of John Deere tractors and ripe Winesap apples. Failure to agree about such large-scale empirical matters does not preclude the interlocutors from nonetheless having a hold on the same concepts. This is the 'tactile Fregeanism' that explains why people can be counted as having radically false (nomologically precluded) views that are nonetheless genuinely *about*, say, arthritis.

Thus the response to the difficulty of reconciling the possibility of genuine communication with an account that individuates concepts by inferential roles comes in two parts. The first is the social-perspectival move. It allows inferential significances to vary with doxastic perspective, while conceptual content, which determines a *function* from perspective to significance, does not. But both the perspective-relative inferential significances of potential speech acts and the perspective-independent conceptual contents that deter-

mine them (in context) are thoroughly *normative* notions—consisting in *proprieties* of discursive scorekeeping.

The crucial second part of the response is accordingly the normative-interpretive move. It distinguishes the proprieties governing *correct* use in which the concepts grasped by individuals consist, on the one hand, from the dispositions to apply concepts, make inferences, and perform speech acts, in which an individual's grasping of a concept consists, on the other—and so distinguishes concepts from conceptions of them. Talk of grasp of concepts as consisting in mastery of inferential roles does not mean that in order to count as grasping a particular concept an individual must be disposed to make or otherwise endorse in practice all the right inferences involving it. To be in the game at all, one must make *enough* of the right moves—but how much is enough is quite flexible. One of the strategies that has guided this work is a commitment to the fruitfulness of shifting theoretical attention from the Cartesian concern with the grip we have on concepts—for Descartes, in the particular form of the centrality of the notion of *certainty*, that is infallibility about the content grasped, including its individuation (so long as we access it clearly and distinctly)—to the Kantian concern with the grip concepts have on us, that is the notion of *necessity* as the bindingness of the rules (including inferential ones) that determine how it is correct to apply those concepts.

Interpreting the members of a community as engaging in discursive practices is interpreting them as binding themselves by objective, shared concepts whose proprieties of use outrun their dispositions to apply them. There is no answer that could be given in advance as to how much one must be able to get right in order to be interpreted as hooked up to one concept or another. Massive individual differences in inferential dispositions among interlocutors are compatible with interpreting them all as nonetheless governed by (answerable to) the same set of conceptual proprieties. For it is compatible with interpreting them as talking about the same objects, answering to the same set of objective facts. In this way the perspectival account of propositional contents (and so conceptual contents generally) combines the intensional and extensional approaches to communication outlined above in 7.5.

6. *Three Levels of Norms*

The normative phenomenalist methodology applies a version of the stance stance to the problem of understanding normative statuses such as the proprieties implicit in discursive scorekeeping practices. It does so by focusing on when it is appropriate to adopt a certain sort of attitude—the stance of interpreting a community as engaged in inferentially articulated deontic scorekeeping practices that confer particular conceptual contents. It has been explained what it is for an interpreter to attribute to a community discursive practices that confer objective, shared, projectable conceptual con-

tents. The question that remains is, What is it that determines when it is *appropriate* or *correct* to adopt one rather than another of these interpretations, to attribute one rather than another of those sets of discursive practices? (Recall that the corresponding question that was asked without being answered above was rather what made it appropriate to adopt any such normative interpretation at all—to attribute practices rather than mere behavior.)

Once again, the issue of the origin of the warrant for employing a normative vocabulary seems just to have been put off. Norms have been appealed to at three different interpretive levels. First, talking and thinking, grasping and applying concepts, is described in terms of inferentially articulated norms; moves in the game of giving and asking for reasons are made intelligible in terms of alterations in what one is committed and entitled to at each stage. This is a normative reconstrual of the *discursive* in terms of *deontic statuses*. Second, what it is to *take* or *treat* interlocutors in practice as committed or entitled, as exhibiting deontic statuses, is explained in terms of scorekeeping practices. The norms implicit in these practices govern the alteration of *deontic attitudes*. At this stage in the account, deontic statuses are understood as instituted by proprieties of scorekeeping—of systematically altering deontic attitudes and thereby assigning pragmatic significances to performances, paradigmatically the fundamental speech act of assertion. Reference to deontic statuses is made only as the objects of deontic attitudes; the only thing one can do with a commitment is to attribute it or undertake it (perhaps, but not necessarily, by acknowledging it).

The third stage applies the methodological strategy of normative phenomenalism one more time, doing for deontic attitudes what those attitudes did for deontic statuses. The focus is now on the practices of *attributing* deontic *attitudes*—interpreting a community as engaged in implicitly normative discursive practices, as keeping deontic score by attributing and acknowledging deontic statuses. The account of deontic scorekeeping on doxastic and practical commitments explains what one must interpret a community as doing in order for it to be *talking* that one is thereby taking them to be doing. More precisely, it specifies conditions on the structure of practices a theorist attributes to a community that are sufficient for community members, so interpreted, to be treating each other as exhibiting propositionally contentful doxastic and practical commitments. Thus the relation envisaged between original intentionality and the stance of the interpreter who attributes it is analogous, at a higher level, to that obtaining between deontic statuses and deontic attitudes—for in place of a direct explanation of what commitment and entitlement are, an account of what it is to *take* someone to have such a status was offered. The phenomenalist explanatory retreat from status to attitude is applied at two levels, within the interpretation and in the relation the interpretation stands in to what is interpreted.

Norms come into the story at three different places: the commitments and

entitlements community members are taken to be attributing to each other; the implicit practical proprieties of scorekeeping with attitudes, which institute those deontic statuses; and the issue of when it is appropriate or correct to interpret a community as exhibiting original intentionality, by attributing particular discursive practices of scorekeeping and attributing deontic statuses. It is normative stances all the way down.

Regularities of communal behavior and disposition specified in nonnormative terms cannot *dictate* the attribution of scorekeeping practices that institute a particular set of normative statuses and confer a particular set of propositional contents. In adopting such a stance, the interpreter takes the interlocutors being interpreted to be *committed* to keeping score according to specific patterns, associating pragmatic significances with discursive performances that correspond to the inferentially articulated contents of the doxastic and practical commitments they express. The interpreter thereby undertakes commitments to various sorts of assessments of propriety of performance of those interpreted. Such commitments on the part of the interpreter are compatible with an indefinitely large lack of fit between the norms attributed and the actual performance of those to whom they are attributed, including their performance in assessing each other. This means that the normatively specified practices attributed by a discursive interpreter are always underdetermined by nonnormatively specified actual performances and dispositions; various sets of practices could be attributed as interpretations of the same behavior. So whenever an interpreter takes a community to be engaging in scorekeeping practices whose implicit proprieties confer one set of propositional contents on the deontic statuses they institute, there will always be alternatives, other sets of contents that could be taken to determine the pragmatic significances that scorekeepers ought to associate with discursive performances. Because of this slippage between the normative and nonnormative specifications of what community members are doing, the interpreter has considerable leeway in how to interpret them.

It remains, then, to discuss the nature of the norms that govern the choice of an interpretation of a community as engaging in one set of implicitly normative, content-conferring discursive scorekeeping practices rather than another, or rather than describing their behavior exclusively in nonnormative terms. This issue is best approached by considering the relation between the discursive scorekeeping stance adopted by the members of a linguistic community (according to an interpretation), on the one hand, and the stance adopted by the interpreter who attributes implicitly normative linguistic practices governing such scorekeeping attitudes (and so original intentionality), on the other. On the face of it, one major difference between the two stances is that discursive scorekeepers take up attitudes toward other members of their *own* communities, while an interpreter who attributes original intentionality takes the members of some *other* community to be discursive

scorekeepers. This is a misleading appearance, however. The important difference between these two sorts of norm-attributing stance is of a different sort. Indeed, under the right circumstances, the difference dissolves entirely, and the two stances coalesce. This collapse of levels provides the key both to understanding the status of the concept-articulating norms implicit in our discursive practices and also to understanding ourselves as not merely *rational*, but *logical* normative creatures, as not merely *expressive*, but *self-explicating* ones.

III. WE HAVE MET THE NORMS, AND THEY ARE OURS

1. *Original Intentionality and the Explicit Discursive Scorekeeping Stance*

The relation between the attitudes of an interpreter who attributes to a community discursive practices (and hence original intentionality), on the one hand, and the proprieties of scorekeeping implicit in those practices, on the other, is modeled on the relation between the deontic attitudes of scorekeepers and the normative statuses they attribute. What the discursive scorekeeper does *implicitly* (taking or treating others, to whom speech acts and discursive commitments are attributed, as discursive scorekeepers), the attributor of original intentionality to a community does *explicitly* (ascribing discursive scorekeeping attitudes). The underlying difference between the two stances is, not the distinction between communally *external* and *internal* attitudes or interpretations, but the distinction between *explicit* and *implicit* ones. Only a creature who can make beliefs explicit—in the sense of claiming and keeping discursive score on claims—can adopt the simple intentional stance and treat another as having beliefs implicit in its intelligent behavior. Just so, only a creature who can make its attitudes toward the beliefs of others explicit—in the sense of being able to ascribe scorekeeping attributions—can adopt the explicitly discursive stance and treat others as making their beliefs explicit, and so as having original intentionality.

Discursive scorekeeping is what the members of a community must be *doing* in order for any of their performances to have the significance (for them) of *saying* something. To take them to be a community of discursive scorekeepers whose practices confer conceptual contents, an interpreter must be capable of saying what they are doing—making *explicit* the broadly inferential proprieties that are (taken to be) *implicit* in their scorekeeping practices. For those who can adopt only the basic scorekeeping stance can attribute commitments to others (even to nonlinguistic, simple intentional systems) and can also take performances to have the significance of assertions, that is of explicit acknowledgments of discursive commitments. They

thereby implicitly recognize others as scorekeepers, and hence as attributors of commitments.

But adopting the basic discursive scorekeeping stance does not require attributing specific attributions to others; it does not require keeping score on their *attributions*, as well as their *acknowledgments* of discursive commitments. In contrast, interpreting the members of a community as engaging in discursive scorekeeping practices requires attributing to them the full range of deontic attitudes: attributing particular attributions as well as particular acknowledgments. And attributions can be attributed only by being ascribed, for it is only when made explicit in the form of propositional contents that they can be embedded in one another and so iterated. Only someone who can *say* something of the form "*S* is committed to the claim that *S*' is committed to the claim that *p*" can adopt the attitude that it makes explicit.³³

To attribute a particular conceptual content to an expression is to say something about how it is correctly used; to attribute such content to a state or status is to say something about the circumstances under which it is appropriately acquired or relinquished and the appropriate consequences of doing so. Interpreting a community as exhibiting original intentionality is taking it that the broadly inferential proprieties that articulate the conceptual contents of their expressions, performances, and states are implicit in their deontic scorekeeping practices. So one capable of adopting that interpretive stance must be able to attribute not only scorekeeping attitudes but also those implicit inferential proprieties, which relate the adoption of one scorekeeping attitude to another. Altering a deontic scorekeeping attitude is a practical doing—the sort of thing a specification of which can play the role of the conclusion of a piece of practical reasoning. So proprieties of scorekeeping can be expressed as proprieties of practical reasoning. Again, only someone who can *say* something of the form "*S* is committed to the claim that if a scorekeeper *does* attribute to *A* commitment to *p*, then the scorekeeper *should* attribute to *A* commitment to *q*" can adopt the attitude that it makes explicit.

This is to say that interpreting a community as engaging in discursive scorekeeping practices, and so as exhibiting original intentionality, requires the full expressive resources of the logical locutions whose use has been reconstructed here in scorekeeping terms. Ascriptional locutions are needed so that both essential flavors of deontic attitude can be attributed, not just adopted: *attributions* as well as *acknowledgments* of commitments. Sentential logical vocabulary, paradigmatically the conditional, makes it possible to attribute acknowledgment of specifically *inferential* commitments. Normative vocabulary is required so that endorsement of a pattern of *practical* reasoning can be attributed.³⁴ Subsentential logical vocabulary such as quantifiers and identity locutions enable the attribution of endorsements of *substitutional* commitments, and so on. The expressive power of these logical

locutions is necessary and sufficient to make possible the adoption of the *explicit* discursive scorekeeping stance.

2. *Expressive Completeness and Interpretive Equilibrium*

Of course it is not just a coincidence that foregoing chapters have explained how to introduce into the basic discursive scorekeeping model just the sorts of logical vocabulary needed to make explicit the various inferentially articulated proprieties implicit in that practice—the very proprieties in virtue of which the expressions, performances, and deontic statuses governed by them count as expressing or exhibiting nonlogical conceptual contents. One of the criteria of adequacy that has guided the project from the outset is that it be possible to elaborate the model of discursive practice to the point where it is characterized by just this sort of *expressive completeness*. This means that the model reconstructs the expressive resources needed to describe the model itself. By means of these logical resources, the theory of discursive practices becomes expressively available to those to whom it applies. What is required is just that the scorekeeping practices that confer conceptual contents on the fundamental sorts of explicating vocabulary used in stating the theory and specifying the content-conferring discursive scorekeeping practices in the first place be themselves specified within the terms of the theory. The hypothetical practitioners who play the idealized Sprachspiel of giving and asking for reasons herein described can then be understood as themselves capable of *saying* what they have been supposed to be *doing*: they can make explicit the implicit practical proprieties in virtue of which they can make anything explicit at all.

Once the expressive resources of a full range of semantically and pragmatically explicating logical vocabulary are in play, those who have mastered them can keep discursive score *explicitly*, by making *claims* about each other's doxastic, practical, and inferential commitments. They can theorize about each other's scorekeeping attitudes. The broadly inferential scorekeeping proprieties that otherwise remain implicit, in the shadows of the practical background, are brought out into the full revealing light of explicit, public, propositional awareness. Particular ascriptions of commitment and entitlement, endorsements of consequential relations among them, and acknowledgments of and failures to acknowledge deontic statuses become topics for public challenge, justification, and debate. Though all the deontic attitudes and practical inferential know-how involved in scorekeeping cannot be made explicit in the form of claims and principles at once, there is no part of that content-constitutive practice that is in principle immune from such codification—out of reach of the searchlight of explicitation. Having been all along implicitly normative beings, at this stage of expressive development we can become explicit to ourselves *as* normative beings—aware both of the sense in which we are creatures of the norms and of the sense in which they are

creatures of ours. Having been all along implicitly discursive beings, at this stage of expressive development we can become explicit to ourselves *as* discursive beings—aware both of the sense in which we are creatures of our concepts (the reasons we produce and consume) and of the sense in which they are creatures of ours.

The members of a linguistic community who adopt the explicit discursive scorekeeping stance to one another achieve thereby a kind of *interpretive equilibrium*. Each one interprets the others as engaging in just the same sort of interpretive activity, as adopting just the same sort of interpretive stance, as one does oneself. This symmetric taking of others to adopt just the same sorts of attitudes one is oneself adopting, characteristic of the discursive scorekeeping stance, contrasts markedly with the *asymmetric* relation obtaining between an interpreter who adopts the *simple* intentional stance and the nonlinguistic creature interpreted as a simple intentional system. In that case the interpreter does not take the system being interpreted to be able to do just what the interpreter is doing, namely *attributing* (as opposed to acknowledging) beliefs, intentions, and endorsement of patterns of practical reasoning. This is one of the reasons what is attributed by such interpreters deserves to be understood as a derivative sort of intentionality.

Linguistic practitioners who have not yet deployed logical vocabulary implicitly treat other interlocutors as adopting the same interpretive stance that they do—as being discursive scorekeepers. The relations between interpreter and interpreted in such basic nonlogical discursive practices are accordingly also symmetric; an interpretive equilibrium is achieved in that case as well. Their idiom is not semantically and pragmatically explicitly complete, however; they adopt attitudes they cannot make explicit as the contents of commitments that can be acknowledged by assertion. They do not attribute the sort of attitude they are adopting just by attributing propositionally contentful commitments. They can only *implicitly* treat one another as scorekeepers, by keeping score on each other.

They treat others as in the general line of business of attributing commitments (and so being scorekeepers) by treating some of their speech acts as having the force or pragmatic significance of acknowledgments of commitments. For it is a necessary condition of being able to acknowledge (and so undertake) discursive commitments in general that one can also attribute them. So the interpretive equilibrium exhibited by basic nonlogical discursive scorekeeping practices is implicit and expressively incomplete. There is still an asymmetry between the stance such scorekeepers are interpreted as adopting by one who attributes original intentionality to the community in whose practices they participate, on the one hand, and the interpretive stance adopted by the interpreter who attributes such content-conferring practices, on the other.

That gap disappears—a complete and explicit interpretive equilibrium is achieved—for a community whose members have access to the full expres-

sive resources supplied by logical vocabulary. They can adopt the *explicit* discursive stance toward one another. Each scorekeeper can explicitly take the others to be doing just what that scorekeeper is doing: attributing discursive deontic attitudes, including that very sort of attribution. Such discursive practitioners have available as topics for explicit discussion the doings that underwrite their sayings, the practices in virtue of which anything can be explicit to or for them at all, and the interpretive stance they adopt to each other.

To the Kantian dictum that judgment is the form of consciousness has been added the claim that logic is the expressive organ of *self*-consciousness. Judging has been construed here as the practical attitude of acknowledging a certain kind of inferentially articulated commitment. Logical vocabulary then supplies the expressive resources needed to make explicit—to put in judgeable form—the semantic and pragmatic bases of judgment. By its means we come to be able to talk about proprieties of *inference*, about the structures of scorekeeping attitudes within which a performance can be accorded the significance of acknowledging or undertaking a *commitment*, and about the relations between these characteristics of specifically *discursive* practice as such. The complete and explicit interpretive equilibrium exhibited by a community whose members adopt the explicit discursive stance toward one another is social self-consciousness.³⁵ Such a community not only *is* a *we*, its members can in the fullest sense *say* ‘we’.

3. *Saying ‘We’*

Such a community-constitutive ‘we’-saying attitude is also the one adopted by those external interpreters who attribute to a community both original intentionality and the use of logical vocabulary. In the weakest sense, we treat others as among us by attributing to them, and interpreting their performances in terms of, propositionally contentful practical and doxastic commitments—that is, by adopting the simple intentional stance. In a more basic sense, we treat others as among us by taking them in addition to perform speech acts. Keeping discursive score in this fuller sense is implicitly treating them as rational scorekeeping creatures who can appreciate the inferentially articulated pragmatic significance not only of their own nonlinguistic performances but also of their claims and of the actions and speech acts of others. At the next level, explicitly keeping discursive score on the members of a community—by ascribing not only acknowledgments but attributions of propositionally contentful commitments—is attributing original intentionality. This is explicitly treating the members of a community as among us, in the sense of being rational linguistic creatures. The richest sort of ‘we’-saying is then taking those others to be in addition *logical* creatures—treating them as able to adopt, toward each other and at least potentially toward us, just the attitude we are adopting toward them.

So at the highest levels of 'we'-saying, interpretive equilibrium is achieved (whether implicitly or explicitly). The interpretive stance attributed to the members of a discursive community approaches that adopted by the interpreter who attributes original intentionality to that community. Finally, the sort of scorekeeping that is—according to the interpreter outside the community—internal to and constitutive of the community being interpreted comes to coincide with the scorekeeping of the interpreter who attributes discursive practices to the members of that community. External interpretation collapses into internal scorekeeping. Thus attributing discursive practices to others is one form or another of 'we'-saying. It is recognizing them as us.³⁶

This assimilation of the external to the internal interpretive point of view means that the question of what it is to interpret the members of a community as engaged in discursive practices—what it is in this fundamental sense to say 'we' to them—has been answered by showing how the deontic scorekeeping model can be elaborated so as to make available the expressive power of logical locutions (in particular ascriptions, conditionals, and normative vocabulary). The next question dictated by the methodological strategy of normative phenomenalism about discursive norms is then, When is it *proper* or *appropriate* to adopt such an interpretive stance? When is it appropriate to say 'we' in the sense of making what others do intelligible as the acknowledging and attributing of propositionally contentful doxastic and practical commitments? When is it appropriate to *interpret* their antics, as we do for each other, rather than merely to *explain* them, as we do for nonsapients?

The collapse of the external explicit discursive interpretive stance into scorekeeping within our own expressively sophisticated practices transforms this from an abstract theoretical question into a concrete question about our own practices. Understood that way, the proper answer would seem to be latitudinarian (as suggested in the opening paragraphs of Chapter 1): one *ought* to adopt the discursive scorekeeping stance whenever one *can* adopt it. For on the one hand, the detailed requirements one must satisfy in order to count as adopting such an interpretation are stringent. Not just any group of interacting organisms can be made out to be attributing to each other commitments whose inferential and social articulation suffices to confer genuinely propositional contents on their performances. So there is little danger of such a generous policy leading to the facile or promiscuous extension of the franchise of sapience to those undeserving of it. And on the other hand, the rewards for adopting the discursive scorekeeping attitude wherever it is possible are great. Conversation is the great good for discursive creatures. Extending it increases our access to information, our knowledge, and our understanding—our semantic and pragmatic self-consciousness. Those who *can* be understood as fellow strugglers in the enterprise of making it explicit *should* be so understood.

Adopting such an inclusive demarcational attitude is saying 'we' to whoever can be understood as adopting demarcating practical attitudes—as them-

selves distinguishing by their scorekeeping a 'we' of rational agents and knowers, inhabiting a normative space of giving and asking for reasons, from an 'it' that comprises what does not live and move and have its being in such a space. Establishing entitlement to such a commitment with respect to demarcation in general would not, however, resolve the more specific issue of the status of the norms that govern the selection of one particular discursive interpretation rather than another. For the underdetermination of normative interpretation by behavior and dispositions specified in nonnormative terms means that whenever what a community does supports an interpretation of its members as engaging in discursive practices in which one set of conceptual norms is implicit, that behavior also supports rival gerrymandered interpretations of them as engaging in discursive practices in which different sets of conceptual norms are implicit. When it is possible to offer some such interpretation, how is it settled which one is most appropriate?

4. Semantic Externalism Begins at Home

The previous issue was a global one, concerning the propriety of attributing discursive scorekeeping practices at all. The present issue is a local one: assuming the global question settled in the affirmative, what is involved in choosing among various specific alternatives? Deciding to treat each of the members of some alien community as one of us (in the sense of treating them as adopting deontic attitudes, attributing and acknowledging propositionally contentful commitments) by no means settles what those contents and commitments should be taken to be. Their speech acts will typically differ in their nonnormative characteristics; they will utter different noises, make different marks (or, for all that it matters to the abstract scorekeeping model of discursive practice, turn colors, emit odors, shift voltages). What about the conceptual contents they express? What the contents of their commitments and expressions are depends on their inferential practices and on the noninferential perceptual circumstances of application and practical consequences of application implicit in their scorekeeping practices. These may differ from ours in a myriad of details and still be intelligibly interpretable. How radically different might they be?

Both the question of what makes a better discursive interpretation and the question of how different from ours the practices of the others might be taken to be before it becomes impossible to offer an intelligible interpretation of them as in the same discursive line of work as we are—as scorekeeping by changes of deontic attitude with the right social and inferential structure to confer propositional content—are questions that can be addressed only by appeal to our actual practices of interpretation in conversation. Because in a community with sufficient expressive resources the tasks of external discursive interpretation and of internal communicative interpretation are tasks of the same kind, looking at the dynamics of intralinguistic interpretation in

ordinary conversation reveals the essential features that determine also the dynamics of interlinguistic interpretation.³⁷ This is to say that there is no usefully general answer to the more specific interpretive question. The coalescence of external and internal discursive interpretation dictates a regress to the background language, to *our* discursive practices. The norms that determine the propriety of choices as to which discursive practices, and so which implicit conceptual norms, to attribute to those we take to be talkers are not available in advance as a set of explicit principles. They are implicit in the particular practices by which we understand one another in ordinary conversation.

The question the interpreter faces is to determine what discursive norms the members of a community have instituted, what conceptual contents they have conferred, by their linguistic practices and deontic attitudes. According to the scorekeeping model, two sorts of attribution are involved in such interpretation. The concepts according to which the truth of their claims and the success of their actions (and so their reliability as empirical reporters and practical agents) should ultimately be assessed are the ones they have committed themselves to (a matter of deontic status) by their dispositions to acknowledge some commitments in their linguistic and nonlinguistic behavior (a matter of deontic attitude). According to the interpreter, the conceptual contents practitioners have bound themselves by can outrun their discriminative dispositions to acknowledge their commitments. For this reason, objective, shared concepts can be understood as projecting beyond the dispositions to apply them of those whose concepts they are. The collapse of external interpretation into internal scorekeeping shows that this semantic externalism is just a special case of the sort of perspectival scorekeeping that has been in play all along: the commitments a scorekeeper *attributes* to someone outrun those that individual *acknowledges*. In acknowledging one discursive commitment, one is in general undertaking others, whether or not one knows what they are. This is the pragmatic (scorekeeping) significance of the *inferential* articulation of their semantic contents.

So the job of an external attributor of linguistic practices is just a special case of the job any discursive scorekeeper has: each must keep two sets of books, distinguishing and correlating the commitments interlocutors are disposed to acknowledge by overt performances, on the one hand, and those they undertake thereby, on the other. These correspond to two ways of specifying the contents of their claims—those made explicit in *de dicto* and *de re* ascriptions, respectively. For recall that *de dicto* specifications extract inferential consequences only with respect to auxiliary hypotheses (including those inferential commitments that would be made propositionally explicit in the form of conditional claims) the ascriptional target acknowledges as collateral commitments. *De re* specifications extract those consequences by appealing to auxiliary hypotheses (including inferential ones codifiable as conditionals) that are (according to the ascriber) *true*. Ordinary intralinguistic

communication—the ability to carry on a conversation across the most ordinary differences in doxastic perspective—requires that scorekeepers be able to move back and forth between these two sorts of specifications of the contents of the commitments they attribute. The contents of the commitments it is appropriate to attribute to another depend both on the commitment-acknowledging performances (both linguistic and nonlinguistic) the ascriptional target is disposed to perform and on how things actually are with the objects being talked about. Mastering our practices of attributing conceptually contentful commitments is learning how in particular cases to adjudicate the claims of these two sources of content. Semantic externalism—the way in which what we mean depends on how things actually are, whether we know how they are or not—is a feature of the perspectival character of propositional content.

So semantic (perspectival) externalism begins at home. The contents of the commitments attributed to others, the concepts they have bound themselves by, cannot be specified apart from reference both to what they are disposed to do and say and to what is true of what they are making claims about. For what actually follows from what (according to a scorekeeper = interpreter) depends on the facts (according to that scorekeeper = interpreter). The point that matters here is that once the task of external interpretation is recognized as a special case of internal interpretation (scorekeeping), the practical norms that govern the attribution of one set of conceptually contentful commitments rather than another can be recognized as just one more instance of deciding what others of us are talking about and what they are saying about it. Our norms for conducting ordinary conversations among ourselves are the ones we use in assessing interpretations. There is never any final answer as to what is correct; everything, including our assessments of such correctness, is itself a subject for conversation and further assessment, challenge, defense, and correction. The only answer to the question of what makes one interpretation better than another is what makes one conversation better than another. The answer is a matter of our practical norms of understanding one another here at home.

So the norms governing the use of the home idiom determine how to project the concepts used to specify the content of the stranger's attitudes (which determine how it would be proper to apply those same concepts in novel situations) in the same way they do for the ascriber's own remarks. This is so even in the case where the stranger is best made intelligible by attributing concepts that differ from those used in the home community. Thus the collapse of external into internal interpretation means that the problem caused by the existence of gerrymandered alternatives to any particular discursive interpretation of another community from the outside is displaced to the context of interpretation and projection within our own community. This regress to our own interpretive practices dissolves, rather than solves, the gerrymandering problem concerning the relation between

regularities and norms. For there is no general problem about how, from *within* a set of implicitly normative discursive practices, what we do and how the world is can be understood to determine what it would be correct to say in various counterfactual situations—what we have committed ourselves to saying, whether we are in a position to get it right or not. The account of the use of *de re* ascriptions of propositional attitude shows explicitly just what is involved in such a determination.

For our own practices come to us with the norms in; we do not just utter noises, we undertake commitments, adopt normative statuses, make pragmatically significant moves in the game of giving and asking for reasons. That there is a vocabulary, for instance any nonnormative one, that does not have sufficient expressive power to make it possible to specify our practices, make the distinctions we make, project in the way we do, has, from within our practices, no particular significance. We are always already inside the game of giving and asking for reasons. We inhabit a normative space, and it is from within those implicitly normative practices that we frame our questions, interpret each other, and assess proprieties of the application of concepts.

The account being offered is embodied in the trajectory described by attempts to answer the question, Where are the norms?

The normative first appears in the story in the guise of deontic statuses, of commitments and entitlements. Thought and talk are presented as structures of commitments and entitlements, with particular expressions having the conceptual contents they do because of the role they play in an inferentially articulated structure of such deontic statuses.

Talk of deontic statuses is then traded in, however, for talk of the attitudes of taking or treating people as committed or entitled. Deontic statuses are revealed as scorekeeping devices used for identifying and individuating deontic attitudes. In this sense the first set of norms turns out to be in the eyes of their beholders. This does not amount to a reduction of the normative to the nonnormative, however, because not only actual attitudes, acknowledgments and attributions of deontic status, but also practical *proprieties* governing the adoption and alteration of such attitudes are invoked in explaining the institution of deontic statuses by discursive scorekeeping practices.

At the next stage, these proprieties themselves are removed to the eye of the discursive interpreter, who takes a community to exhibit original intentionality by attributing to it discursive practices socially and inferentially articulated in such a way as to confer propositional contents. Once again, however, it is not only the *actual* attitudes adopted by external interpreters that must be considered but also *proprieties* governing the adoption of the discursive stance and commitment to a particular interpretation.

With the collapse of external into internal interpretation—its revelation as a special case of the sort of interpretation that goes on all the time within the practices of a discursive community—those proprieties are assimilated

to the ordinary scorekeeping proprieties in play in *our own* discursive practices. The norms turn out to be . . . here.

5. *Making It Explicit*

So the theoretical attempt to track down the 'source' of the normative dimension in discourse leads us right back to our own implicitly normative practices. The structure of those practices can be elucidated, but always from within normative space, from within our normative practices of giving and asking for reasons. That is the project that has been pursued in this work. Its aim is not reductive but expressive: making explicit the implicit structure characteristic of discursive practice as such.

The irreducibly normative pragmatics (theory of social practice) presented here is elaborated in terms of the basic deontic statuses of *commitment* and *entitlement* to commitments, and the essentially perspectival scorekeeping attitudes of *attributing* and *acknowledging* those deontic statuses. The semantics, or theory of the sorts of conceptual content that can be conferred by such deontic scorekeeping practices, takes the form of an account of the *inferential*, *substitutional*, and *anaphoric* articulation that distinguishes specifically *discursive* commitments. The result is a use theory of meaning—a specification of the social-functional roles that doxastic and practical commitments and the speech acts that express them must play in order to qualify as semantically contentful. The sorts of content addressed are those traditionally grouped together under the heading of 'intentionality'. Saying what pragmatic scorekeeping significance speech acts must have to count as assertions makes it possible to explain *propositional* contentfulness in turn as what can in that sense be made explicit—as what can in the first instance be *said* (as well as believed or meant or done). *Empirical* and *practical* contributions to such propositional (assertable, and so believable) contents are explained in terms of their conceptually articulated incorporation of the appropriate causal antecedents (in perception) and consequents (in action) of acknowledgments of discursive commitments. The *representational* dimension of propositional contents is explicated in terms of the social-perspectival character of discursive scorekeeping and the substitutional substructure of its inferential articulation. In this way it is possible to understand what is involved in assessments of judgments as *objectively* true or false—as correct or incorrect in a sense that answers to the properties and relations of the objects they are *about*, rather than to the attitudes of any or all of the members of the community of concept users.

One of the leading ideas of this enterprise is that developing an account of how semantics is rooted in pragmatics (meaning in use, content in social-functional role) is an exercise not only in the philosophy of language and the philosophy of mind but also in the philosophy of *logic*. Discursive practice is understood in terms of reasoning and representing, but above all in terms

of *expressing*—the activity of making it explicit. The expressive role distinctive of logical vocabulary is its use in making explicit the fundamental semantic and pragmatic structures of discursive practice, and hence of explicitness and expression. Pursuing the ideal of expressive completeness requires working out an account of the practices of using various particular logical locutions—paradigmatically those used to express inferential, substitutional, and anaphoric commitments and those used to ascribe discursive commitments to others.

In the end, though, this expressive account of language, mind, and logic is an account of who *we* are. For it is an account of the sort of thing that constitutes itself as an expressive being—as a creature who makes explicit, and who makes itself explicit. We are sapient: rational, expressive—that is, discursive—beings. But we are more than rational expressive beings. We are also *logical*, *self*-expressive beings. We not only make *it* explicit, we make *ourselves* explicit *as* making it explicit.