4

Perception and Action: The Conferral of Empirical and Practical Conceptual Content

> Nature to be commanded must be obeyed; and that which in contemplation is as the cause is in operation as the rule.
>
> BACON, Novum Organum

> In the beginning was the deed.
>
> GOETHE, Faust

> The true being of man is his deed; in this the individual is actual . . . What the deed is can be said of it. It is this, and its being is not merely a sign, but the fact itself. It is this, and the individual human being is what the deed is. Action simply translates an initially implicit being into a being that is made explicit.
>
> HEGEL, Phenomenology

I. ASSERTIONS AS KNOWLEDGE CLAIMS

1. Five Strategic Explanatory Commitments

Assertions are the sort of *claims* made in the standard case by uttering freestanding declarative sentences—that is, sentences whose occurrence is not embedded in the occurrence of a compound sentence. A commitment has been undertaken here to an order of explanation dictating that this principle be exploited by defining declarative sentences in terms of an account of assertion, which evidently then is required to be made available independently. This contrasts with the procedure common in formal semantics, in which the theorist leaves until later the task of getting a grip on the activity, force, or significance of assertion but provides an antecedently defined construal of sentences. The strategic commitment to treating what is expressed by the use of *sentences* (rather than what is expressed by the use of singular terms or predicates) as the fundamental sort of semantic content is an element of the present account that has been taken over from Kant.

The pragmatist strategic commitment to understanding semantics in terms of pragmatics (the *contents* associated with expressions in terms of the practices governing their *use*) is an element of the present account that has been taken over from Wittgenstein. The strategic commitment to specifying such a pragmatics in the first instance in *normative* terms is an element of

the present account that has been taken over from Kant, Frege, Wittgenstein, and Sellars. The inferentialist strategic commitment to treating the public linguistic practice of *asserting* as the fundamental activity involving such contents, rather than the private mental practice of judgment, is an element of the present account that has been taken over from Dummett. The strategic commitment to understanding asserting a sentence as a significance a performance acquires in virtue of its role in a practice of giving and asking for *reasons*, of justifying and communicating justifications, is an element of the present account that has been taken over from Sellars.

This constellation of commitments combines the *normative* articulation of the pragmatic significances of assertional performances with the *inferential* articulation of the propositional contents they express, in part by putting the issue of whether an asserter is *entitled* to the *commitment* undertaken by making an assertion at the center of the practice that institutes those significances and confers those contents. To do so is to treat the sort of claim involved in asserting as an implicit *knowledge* claim. From the point of view of the concerns that motivate the present project, this is as it should be. For the aim all along has been to elaborate a criterion of demarcation that sets us off by our peculiar susceptibility to *reasons*. It is this susceptibility that makes it appropriate to think of ourselves in terms of the categories of *knowledge* and *action*. That is why the story really begins with Kant's observation that knowings and actions are to be distinguished from other things we do by the characteristic way in which we are *responsible* for them. The notion of *discursive commitment* arises in the domain of social practice when one focuses specifically on the norms that are articulated in the form of reasons.

Absent the inferential dimension, the norms implicit in a set of social practices could be understood neither as conferring *propositional* contents nor as instituting *assertional* significances, hence not as governing genuinely *linguistic* practice. Inferential connections enter into the alterations of attitude (the scorekeeping that defines assertional practice) in three fundamental ways: one corresponding to each of the basic sorts of deontic status, and a third involving the relation between them. First, part of the significance of acknowledging an assertional commitment is that one thereby undertakes commitment as well to all those contents it entails—that is, to consequences that follow from it by commitment-preserving inferences. One who claims that a lion roared is committed thereby to a mammal's having roared. Second, part of the significance of undertaking an assertional commitment is that one thereby undertakes a conditional task-responsibility to demonstrate one's *entitlement* to that commitment, if faced with a warranted challenge. Here justificatory or entitlement-preserving inferences involving the asserted content help determine what deontic statuses are attributed to which asserters, challengers, and deferrers. Finally, part of the deontic scorekeeping practice within which performances can have the significance characteristic of claim-

ings is to withhold attribution of *entitlement* to commitments *incompatible* with a *commitment* that has been undertaken (whether by overt assertion or consequentially). This is the practice that defines incompatibility relations on the contents of deontic states; two claims are incompatible if commitment to one precludes entitlement to the other. Connections of these three sorts consist in proprieties that govern the alterations of deontic attitudes by which interlocutors keep discursive score. Deontic statuses then count as having inferentially articulated contents because of the pragmatic scorekeeping significance of the performances that express the acquisition of those statuses.

2. *Knowledge as a Complex Hybrid Deontic Status*

That assertions have the default status and significance of implicit knowledge claims is to be understood in terms of these inferentially structured interactions between the two *modally* distinct deontic statuses, *commitment* and *entitlement*, and the two *socially* distinct deontic attitudes, *attributing* and *undertaking* deontic statuses. The status one attributes in attributing knowledge is traditionally understood according to the tripartite structure: justified true belief (JTB).[1] One of the leading ideas of the present approach is that the notion of normative status can be made to do much of the theoretical and explanatory work that the notion of intentional state has heretofore been called on to do. In the social practice model, talk of *belief* is replaced by talk of assertional or doxastic *commitment*. According to the JTB approach, attributing knowledge is attributing a special kind of belief. So attributions of knowledge are to be rendered here in terms of the deontic attitude of *attributing commitments*—specifically commitments of the sort that can be undertaken or acknowledged by performing a speech act that has the significance characteristic of assertions.

As Plato had already pointed out, there is more to attributing knowledge than attributing belief. There is also the issue spoken to by the demand for justification, or as Plato has it, for an account. According to the canonical tripartite understanding, knowledge is not just belief but *justified* belief. Clearly what corresponds to this condition in the deontic version is the demand that the one taken to be a knower not only have a commitment but be *entitled* to that commitment. Making an assertion, it has been said, is making a knowledge claim. Assertional performances as modeled here have the significance not only of undertaking commitments but of defeasible claims to entitlement to those commitments. So one is not attributing knowledge to someone unless one not only attributes a commitment but also attributes a corresponding entitlement. Most classical epistemological problems are really problems concerning this deontic status—justification for believing, or more generally entitlement to believe.

Before considering entitlement to believe, however, a few words are in

order about the third limb of the tripartite rendering of what one is doing in taking someone to have the status of a knower. In attributing knowledge one is not just attributing justified belief—a commitment of the sort that can be undertaken by asserting it and an entitlement to that commitment of the sort that can be demonstrated by justifying it. One must also take the belief to be *true*. What is the social-deontic attitude corresponding to the truth condition on attributions of knowledge?

The attitude of taking-true is just that of *acknowledging* an assertional commitment (the attitude that grounds consequential undertakings of such commitments). A theory of asserting and assertional commitment is a theory of taking-true. Evidently this principle can be exploited according to two different orders of explanation: moving from a prior notion of truth to an understanding of asserting (or judging) as taking, treating, or putting forward *as* true, or moving from a notion of asserting to a notion of truth as what one is taking, treating, or putting forward a claim *as*.[2] The latter line of thought accords 'true' an expressive role, in permitting us to *say* something about assertion, rather than an explanatory role, as something that can be understood in advance of understanding assertion and used to advance to such an understanding. This approach is the one pursued in Chapter 5.

In taking someone to be a knower, one *attributes* a *commitment, attributes entitlement* to that commitment, and *acknowledges commitment* to the same content oneself. Undertaking the commitment is part of what the asserter authorizes others to do—not only to attribute the commitment but also to undertake it, on the asserter's authority. That authority depends on the asserter's entitlement to the commitment, so the asserter is implicitly claiming that entitlement as well. That is why assertions in the basic practices described here have the significance of claims to knowledge. For others to take those claims to be successful is for them to attribute the commitment undertaken, in addition to attribute entitlement to it, and finally to endorse the claim themselves. These correspond, in the model of linguistic practices in terms of scorekeeping with deontic attitudes, to taking to believe, taking to be justified in that belief, and taking the belief to be true.[3]

According to this way of understanding knowledge claims, the significance of the truth condition on attributions of knowledge lies in the fundamental difference in social perspective between *attributing* a commitment (or other deontic status) to another and *acknowledging* it oneself. Knowledge is a *complex* deontic *status*, in the sense that it involves both commitment and entitlement. But attributions of knowledge (and so claims of knowledge) are also *hybrid* deontic *attitudes*. So knowledge can be called a hybrid deontic status. Attributions of knowledge (the attitudes in terms of which that status is to be understood) are hybrid deontic attitudes in the sense that they involve both attributing and acknowledging commitments. These attitudes are perfectly intelligible in the context of the model presented here of the social practices that institute assertional force.

But it is also clear in those terms that there is a great danger of misinterpreting what one is doing in calling a claim true. The danger is in misunderstanding taking-true solely in terms of *attributions* of status (rather than as essentially involving also the *undertaking* of one), by assimilating that attitude too closely to the attributions of commitment and entitlement involved in the other dimensions involved in treating a claim as a bit of knowledge. If the hybrid nature of that attitude is overlooked, it will be thought of as consisting only in attributions; attributions of knowledge will be taken as comprising attributions of belief, attributions of justification, and attributions of truth. One then looks for the property or status one is attributing to a claim in taking it as true.[4] The property or status projected by misconstruing undertaking a commitment as attributing some property or status to it is bound to be "queer."

In the deontic scorekeeping model of inferentially articulated linguistic social practices, asserting is making a knowledge claim. The attitudes in terms of which the hybrid deontic status of knowledge is understood are just those in terms of which the significance of assertions is specified. Assertional practice is accordingly a version not only of the game of giving and asking for reasons but also of the game of making and defending claims to knowledge. Practitioners who can produce and consume assertions are *linguistic* beings. Practitioners who can produce and consume reasons are *rational* beings. Practitioners who can produce and consume knowledge claims are *cognitive* beings. On the account presented here, these are three ways of talking about the same practices and the same capacities.[5]

Underlying all of them is the inferentially and socially articulated notion of *discursive commitment*. It is the topic in which philosophy of language, philosophy of mind,[6] and epistemology are alike rooted. What epistemology studies is a deontic status that is implicitly in play in any practices involving propositional contents—whether or not those practices include the expressive resources provided by words like 'knowledge', which can be used to make attitudes toward that status explicit. For making and defending what are implicitly claims to knowledge is an essential feature of discursive practice as such.

On this account, prizing and searching for knowledge are not specialized intellectual virtues, appropriate only to a sophisticated, culturally late-coming elite. They are built into what we fundamentally are. The complex hybrid deontic status of knowledge defines the *success* of assertion. Treating an assertion as expressing knowledge—attributing to the asserter entitlement to the commitment undertaken thereby and endorsing that commitment oneself—is the response that constitutes the practical recognition of the authority implicitly claimed by the assertion. For that is the authority to license undertakings of commitment to that same claim by those in the audience, in virtue of the asserter's entitlement to the commitment. For a scorekeeper fully to accept the authority implicitly claimed in the making of an assertion

is just for that scorekeeper to treat it as having the status of knowledge. So the aspiration not only to *truth* but to *knowledge* is built right into the normative structure of assertional practice. (And it should be noticed that the sense in which the status of knowledge provides the *point* of assertion can be specified in advance of any consideration of the *intentions* of the practitioners.) Knowledge is on this account an ideal projected by the very possibility of saying anything at all.

3. *Justifying and Being Entitled*

In this context it is useful to look a bit more closely at the structure of attributions of entitlement to assertional commitments, from a more traditionally epistemological perspective. According to the tripartite analysis, to take a claim to express knowledge is to take it to express a justified true belief. The justification condition on knowledge will be misunderstood if one does not distinguish between two senses in which a belief can be said to be justified. In one sense, to call a belief justified is to invoke its relation to the process of *justifying* it. To *be* justified in this sense is to *have been* justified—exhibited as the conclusion of an inference of a certain kind. In another sense, to call a belief justified is to attribute to it what might be called *positive justificatory status.* Positive justificatory status is just what has been talked about here in terms of *entitlement* to a claim.

The relation between possession of such status and the activity of justifying may be quite indirect. In particular, justifying a claim is only one way in which it can acquire positive justificatory status. Indeed, as has already been pointed out, to avoid embarking on a foundationalist regress it is necessary to acknowledge that a commitment may have a positive justificatory status without having been justified (indeed, without that entitlement having been defended in any way, whether *intra*personally by inference or *inter*personally by deference). Since any activity of justifying—even if that term is understood broadly, as entitling (so as to include deferring as well as inferring)—is a mechanism making possible the *inheritance* of entitlements, there must be some at least prima facie entitlements available to get the process off the ground. If dogmatism is to be avoided, such entitlements must not be immune to criticism; there must be mechanisms for bringing them into question. The combination of prima facie entitlements and ways of criticizing and undermining them is what was called the *structure of default and challenge.* It characterizes a dynamic process of acquisition and loss of entitlements by various commitments on the part of various interlocutors (kept track of in the attitudes of claiming and attributing entitlements) and of withholding such claims and attributions.

Classical foundationalism considers only justifying in the narrow sense of an inferential activity, not in the broader sense of vindication that includes the communicational dimension appealed to by deferential entitling (the

authority of testimony). This is unfortunate, for if the analysis just offered of what one is doing in calling something knowledge is correct, the hybrid deontic status of knowledge is incomprehensible in abstraction from the social distinction of perspective distinguishing the deontic attitudes of attributing and undertaking commitments. One of the centerpieces of the present account is its attention to the interaction of the two dimensions of the practice of giving and asking for reasons for commitments to inferentially articulated contents: the intracontent, interpersonal communicational dimension and the intercontent, intrapersonal justificatory dimension. Both the individuation of the contents individuals are responsible for and the individuation of the individuals responsible for them are to be understood in terms of this structure. Equally important, as Chapter 8 shows, an inferential understanding of the *representational* dimension of conceptual content depends upon an appreciation of the *social* articulation of inferential practice. None of this is accessible from the point of view of the one-dimensional approach that ignores the significance of communication for justification. Even within that narrower compass afforded by exclusive attention to intrapersonal, intercontent entitling, however, the consequence of insisting that positive justificatory status can be the result only of justifying is a dual regress—one regress on the side of entitlement to premises, and another on the side of entitlement to inferences.

To illuminate the default and challenge response to the threat of a foundationalist regress on the side of premises, consider its twin on the side of inferences. If entitlement to a commitment to q is at issue, and that commitment is justified by asserting p, the vindication might be unsuccessful either because the commitment to p is not one the interlocutor is entitled to or because the inference from p to q is not correct (in this case, not entitlement-preserving). A regress on the side of the inferences results if one insists that each inference is, to begin with, in need of support or justification. Endorsing the propriety of an inference is brought into the game of giving and asking for reasons in a new way by making the inference explicit in the form of a conditional, which can be endorsed, challenged, and defended like any other assertible. The demand is then for an explicit rule or principle to warrant the propriety of every inferential transition appealed to in justifying a claim.

But pragmatism maintains that to demand this is to view things the wrong way around. One must start with a notion of taking or treating inferences as correct in practice. Without such a practice, there is no game of giving and asking for reasons to bring inferences into in the form of explicit assertions. Once the game is under way, the practical inferential attitudes it involves can then, on suitable occasions, be made explicit in the form of endorsements of conditionals. But what those conditionals express is intelligible only in terms of the underlying inferential practice.

If it is insisted instead that no move be treated as entitled or entitlement-

preserving until its entitlement has been demonstrated or justified, a new premise is introduced corresponding to every inference, and also a new inference employing that premise. Then the regress ramifies, as the entitlements to those new premises and those new inferences must themselves be secured. Within the resulting regress can be discerned Kant's and Wittgenstein's regress of rules—and where in addition the goodness of inference is identified with *formal* goodness of inference, Lewis Carroll's regress of conditionals and detachment from conditionals as well. Carroll's multiplication of premises standing behind inferences should be halted by acknowledging primitive *rules* of inference; the multiplication of conditionals explicitating implicit "enthymematically suppressed" premises should be halted by acknowledging primitive *material* rules of inference; and the multiplication of rules should be halted by acknowledging primitive material-inferential *practices.*

These moves have all already been considered. The default-and-challenge structure of assertional entitlement just amounts to extending to the case of assertions the policy that underwrites these ways of thinking about inferences. What is fundamental in each case is the practical attitude of taking or treating as correct moves in the game of giving and asking for reasons. Though such entitlements can be brought into question later, one initially is entitled to whatever one is in practice taken or treated as entitled to; deontic statuses must be understood in terms of practical deontic attitudes. It makes sense that this way of construing the proprieties of inference that articulate the propositional content of assertional speech acts and the commitments and entitlements they involve should extend as well to the proprieties that govern those assertional performances and deontic statuses. For asserting and inferring are two sides of one coin; neither activity is intelligible except in relation to the other. Undertaking an assertional commitment involves a commitment to the propriety of inferences from the circumstances of application to the consequences of application of the concepts in terms of which its content is articulated. If claiming is to be possible at all, some of those content-constitutive implicit inferential proprieties must in practice be taken for granted, treated as prima facie in order—not as innocent until *proven* guilty, but at least as innocent until indicted on the basis of reasonable suspicion. In the same way, sometimes a defeasible presumption that the application of those concepts in an assertion or judgment is appropriate must be in order.

II. RELIABILITY

1. Reliabilism and Entitlement

When the justification-as-entitlement of a belief is decoupled to this extent from the activity of inferentially (or, for that matter, deferentially) justifying it, the question arises whether the latter notion need be taken to

play any role whatever in the understanding of the status of being a justified belief that is appealed to by the tripartite analysis of knowledge. It is generally agreed that *some* sort of entitlement to a claim is required for it to be a candidate for expressing knowledge. But it is not obvious that inferring in the sense of justifying is at all fundamental to that sort of entitlement. When examples of the sort that motivate the tripartite analysis in the first place are examined more closely, it appears that what is forbidden is that it be merely *accidental* that one has a true belief. Someone who flips a coin to decide which is the correct road to Athens may by accident pick the right one, and may somehow or other come to believe that the one chosen is the correct road. But such a true belief does not qualify as knowledge. One way of showing that the belief is not merely accidental is indeed to provide an account, to offer reasons for the belief.

It has been suggested, however, that this is merely *one* way, and by no means the most basic, in which the belief could be shown to have credentials beyond those provided by happenstance and coincidence. In particular, the correctness of the belief is not merely fortuitous if it is the outcome of a generally reliable belief-forming mechanism. Epistemological *reliabilists* claim that this is the sort of entitlement status that must be attributed (besides the status of being a true belief) for attributions of knowledge.[7] The perceptual mechanisms underlying entitlement to empirical claims provide the most important and persuasive examples. This line of thought is sometimes extended to an analysis of justification as consisting simply in the demonstration of the reliability of a belief-forming mechanism. The version that is of interest here, however, claims only that reliability can do *all* the work that inferential justifying is taken to do in the standard tripartite analysis (and, by extension, in the assertional practices of making and defending knowledge claims described here). The further Procrustean reductive assimilation of justifying to the paradigm of indicating reliability can be put to one side at this point, for it is the difficulties that arise already with the weaker thesis that are most instructive.

Reliability accounts of entitlement to assertional commitments and *regularity* accounts of the correctness of such commitments are species of one genus. They share a common strategy for naturalizing the different norms they address. In tracing the relation between them, it is helpful to keep in mind the basic case in which making an assertion consists in noninferentially applying a ground-level empirical concept in a particular situation. Regularity theories are attempts to determine the boundaries of concepts—which determine the difference between correct and incorrect application—by appealing to regularities or patterns in the actual applications of the concept and the dispositions to apply that concept that are exhibited by an individual or a community. The sort of *correctness* of application of concepts that such theories aim to explicate is what is assessed by judgments of the *truth* of the resulting assertion.[8] What it is for *A* to take *B*'s claim that

something is a porcupine to be a correct application of that concept is for *A* to take the claim to be true, that is, for *A* to endorse it, for *A* to undertake or acknowledge commitment to that same content. This is a different deontic attitude from *A*'s attributing to *B entitlement* to the commitment undertaken by that claim, and so to that application of the concept; different sorts of normative status are involved.

The concept of *reliability* in making a claim or applying a concept presupposes, rather than analyzes, such a notion of *correct* claiming or concept-application. For a reliable performer is just one who generally produces a correct performance; assessments of reliability are assessments of the probability of correctness. Thus the issue of reliability cannot be raised until the question of correctness that regularity theories address has been answered. Nonetheless, these different theories share an approach. For reliability theorists offer an account of entitlement that appeals only to patterns or regularities of *correct* claiming or application of concepts. The reliabilist idea is that *entitlement* to a particular claim or application of a concept—a derivative sort of correctness of claiming—can be understood entirely in terms of *dispositions* to produce *correct* performances of that kind. The regularist idea is that such correctness of claims or applications of concepts can be understood entirely in terms of *dispositions* (in some variants, those that are in some sense communal dispositions) to produce performances of that kind.

One of the major difficulties raised for the strategy of construing the correctness of discursive performances in terms of regularities or dispositions specified in nonnormative terms was the gerrymandering problem. There is no single pattern or regularity exhibited by any set of actual or virtual performances; where there is one, there are many—indeed an infinite number. No matter what a candidate performance whose correctness is at issue is like, and no matter what the history to which it must answer is like, there is *some* way of specifying the pattern exhibited by those prior performances so as to include the candidate as just what is required to continue that pattern "in the same way." The attempt to identify the normative distinction between correct and incorrect performances with the naturalistic distinction between regular and irregular performances fails because no performance is simply irregular (even relative to a specified set of performances with respect to which its coregularity is to be assessed), and so none would be counted as incorrect by such a criterion. Appeal to regularity and irregularity can do normative explanatory work only if there is some way of *privileging* some regularities over others—some way, in other words, of saying what the *correct* regularity is. The problem of sorting performances into correct and incorrect is transformed by the regularist strategy into the problem of sorting regularities into the relevant and the irrelevant, the ones that *ought* to be taken account of in assessments of correctness, and those that ought not. From this vantage point, regularity theories appear as merely putting off the normative issue, moving the bump in the carpet around rather than smoothing it out.

2. Barn Facades

Reliability theories share with regularity theories the same fundamental strategy for explaining in naturalistic terms the normative statuses involved in discursive practice (although the phenomena they address are at different levels). So it might be expected that the possibility of gerrymandering would raise similar difficulties for a pure reliability strategy for construing entitlements to claims that it does for a pure regularity strategy for construing the correctness of those claims. This is indeed the case. A striking illustration of how the gerrymandering considerations get a grip on assessments of entitlement in terms of reliability is provided by Goldman's barn-facade example.[9] The example is forwarded to show the inadequacy of an account that seeks to ground the cognitive authority of noninferential reports exclusively in features of the causal chain leading from the reported state of affairs to the perceptual reporting of it. The leading idea of such causal theories is that a true belief, paradigmatically one acquired perceptually, counts as knowledge just in case it is caused in the right way by what it is about. To see that such a particular causal connection is not sufficient to make a true belief qualify as knowledge (and so cannot by itself perform the explanatory job assigned to the entitlement condition by the tripartite analysis), Goldman suggests comparing two cases that are alike as far as the causal chain leading to a claim is concerned, but unlike in the status of that claim as knowledge.

In each case the subject is in ideal circumstances for visual perception and is confronted by what is in fact a barn. In each the subject responds to the visible presence of the barn by confidently reporting the presence of the barn. The causal chains in each case are entirely standard, the barns reflecting light, which travels undisturbed to the subject's retina, and so on. Yet one of the subjects is, and the other is not, without knowing it located in Barn-Facade County. The local hobby in that county is building incredibly lifelike trompe l'oeil facades of barns. In this county, 99 percent of what appear to be barns are actually such facades. Each subject would in fact, if confronted with such a facade (and not alerted to the special practices of the natives), confidently report the presence of a barn. Goldman's plausible claim is that the claim of the subject who is *not* in Barn-Facade County *does* express knowledge (is a claim that subject is in the relevant sense entitled to), while the claim of the subject who *is* in Barn-Facade County does *not* express knowledge.

For the first point: the fact that there are *some* circumstances under which a subject could be fooled does not in general preclude the subject from having knowledge in the case in which that subject is not fooled. As Austin argued, the fact that it is possible to make a replica of a sparrow so cunningly contrived that I cannot tell it from the real thing does not mean that I cannot see a sparrow and know that it is a sparrow.[10] The mere possibility of hyperbolic doubt does not entitle anyone to it and does not undermine

entitlements in ordinary cases. For the second point: even though the subject who has the cognitive bad luck to be in Barn-Facade County is in fact looking at a barn, it is in an important sense just an *accident* that that is the case. The county is rife with perceptual situations in which the subject would with equal confidence and warrant *falsely* report the presence of a barn. Under these circumstances, the belief just *happens* to be true, and the subject should not be taken to *know* that there is a barn present.

Goldman claims first that this sort of example shows that one must look beyond the particular causal antecedents of a belief in order to determine its status as one the believer is entitled to in the sense relevant to assessments of knowledge. For in this case what distinguishes the two subjects is not the causal chains connecting them to the barns but only the incidence of barn facades in the vicinity, which is causally irrelevant to their perceptual transactions with the barns they are in fact looking at. His second claim is that the way in which that difference of causally irrelevant circumstance makes a difference to the assessment of entitlement and hence of knowledge can be understood in terms of the variable *reliability*, in those different circumstances, of the belief-forming mechanism that leads to the true belief in each case. The same differential responsive dispositions, the same noninferential reporting capacity, is in play in both cases.

The difference is that in Barn-Facade County it is not a reliable mechanism, while in the rest of the (largely barn-facadeless) world it is. How reliable a belief-forming mechanism is, how likely it is to yield a true claim, a correct application of a concept, depends on the circumstances in which it is exercised. My inability to tell sparrows from cunning duplicates does not disqualify me from being a reliable reporter of sparrows, so long as my environment is quite unlikely to confront me with a ringer. If such duplicates were to become common, the reliability of my differential responsive dispositions would degrade (and with it my capacity to acquire knowledge thereby in the cases where all goes well), even though the way in which that mechanism would respond to each possible case remained the same throughout. The probability of being correct in a particular case depends on the actual incidence of indistinguishable phonies. Thus the notion of reliability of belief-forming mechanisms provides just what is wanted to explain the barn-facade cases.

3. *Gerrymandering and the Problem of Reference Classes*

Goldman's argument is decisive against exclusively causal theories of knowledge, and it shows how assessments of reliability can function in assessments of entitlement—particularly entitlement to commitments acquired as a result of noninferential reporting capacities. But (though he does not make the point) it also underscores the possibility of gerrymandering, and hence the inadequacy of construing cognitive entitlement exclu-

sively in terms of reliability. In the case of regularity theories of the correctness of the application of a concept, it is the boundaries of the concepts that can be gerrymandered in such a way as to preclude assessments of irregularity, and hence of error. In the case of reliability theories of entitlement, it is rather the boundaries of the reference class with respect to which reliability is assessed that can be gerrymandered in such a way as to preclude assessments of unreliability, and hence of lack of entitlement.

Goldman's idea is that reliability is an objective affair, determined by the objective probability of a correct judgment, given one's circumstances. But such probabilities vary with the specification of those circumstances. Given a reference class of relevantly similar cases, frequencies of success define objective probabilities. The question remains how a privileged reference class is to be determined. What is the *correct* reference class with respect to which to assess such probabilities?

If the reference class is restricted to the actual case of the perceptual judgment that a barn is present, even in Barn-Facade County (since in the case being considered by hypothesis one is actually looking at a barn) the frequency of correct judgments is 1. So relative to that quite restricted reference class one is totally reliable. If the reference class is widened to the whole county, the frequency of correct judgments is reduced to 1 percent. So, relative to that less restricted reference class, one is quite unreliable. But since the customs of Barn-Facade County are quite parochial, the relative frequency of barn facades in the country as a whole is quite low. Relative to the nation as a whole, one is quite a reliable noninferential reporter of the presence of barns. Relative to the state, one's reliability will fall somewhere in between. One of the nice things about this example is that here the metaphor of *boundaries* is made concrete, and the difficulty of selecting the proper boundary is literally geographic.

Focusing on the relativity of reliability to decisions about where to draw these boundaries makes it evident that the question "Reliable or not?" is underdetermined in exactly the same way that the question "Regular or not?" is underdetermined. There are always *some* regularities that are being instantiated, and (in the case where the claim one is making is true) there are always *some* reference classes with respect to which one is reliable. Using these naturalistic notions to stand in for genuinely normative assessments works only relative to some way of privileging regularities or reference classes. The notions of regularity and reliability cannot do all the work by themselves. For the concept of regularity cannot discriminate between regularities, and that of reliability or probability of success relative to a reference class cannot discriminate between reference classes. In the sense in which, given a regularity, there is an objective matter of fact as to whether a further performance continues it, there is no objective matter of fact as to which of the various regularities exhibited by a given history of actual or dispositional performances is the right one to assess correctness with respect to. In the

sense in which, given a reference class of relevantly similar cases, there is an objective matter of fact as to what the probability that a certain skill exercised in those circumstances will yield a correct performance, there is no objective matter of fact as to which of the various possible reference classes to which the case in question might be assimilated is the right one to assess reliability with respect to. An objective or naturalistic theory of cognitive entitlement cannot be derived solely from considerations of reliability, any more than an objective or naturalistic theory of the correct application of concepts can be derived solely from considerations of regularity.

4. Taking or Treating as Reliable

The general strategy of this work is to supply what is wanting in regularity theories of correct concept-application by appealing to the social (in an *I-thou* sense) and practical deontic attitudes of taking or treating a performance as correct or incorrect. The paradigm is taking or treating an assertion as correct in the sense of endorsing it, undertaking that commitment oneself, which is taking what it says to be true. It is these attitudes on the part of interpreters, of the deontic scorekeepers who attribute discursive commitments, that privilege some regularities over others and give a sense to the notion of *correct* use of expressions and so applications of concepts. The deontic *status* of *being* a correct application of a concept is to be understood in terms of the deontic *attitude* of *taking* or *treating* such an application *as* correct. That attitude (endorsing a claim, undertaking an assertional commitment) cannot be understood apart from its role in the essentially social practice of giving and asking for reasons, making and defending knowledge claims. The norms implicit in the application of concepts are social and perspectival, not (to begin with) objective and naturalistic.[11]

Regularity theories attempt to naturalize the normative status of correct claiming or concept-application. The countervailing idea pursued here is to explain that status by saying what it is for a performance to be taken or treated in practice as having such a significance. This is to focus on the deontic attitudes of acknowledging conceptual norms by attributing normative statuses and significances. Reliability theories attempt to naturalize the normative status of *entitlement* to the commitment undertaken by making a claim or applying a concept. The corresponding countervailing idea to be pursued here is accordingly to explain that status by saying what it is for a performance to be taken or treated in practice as having such a significance. This is to focus on the deontic attitudes that acknowledge that status and attribute that significance.

What in practice privileges some of the reference classes with respect to which reliability may be assessed over other such reference classes is the attitudes of those who attribute the commitment whose entitlement is in question. Each interpreter implicitly distinguishes between reference classes that are relevant and those that are irrelevant to the assessment of reliability,

and hence of entitlement to claims, by the circumstances under which that interpreter accords cognitive authority to those claims. The sort of authority in question here is not that acknowledged by the interpreter's own endorsement of the claim—that is, taking it to be correct in the sense of taking it to be *true* (which is the sense of correctness addressed by regularity theories rather than by reliability theories). The sort of authority in question is rather that of having an inheritable entitlement: the sort that supports successful deferrals by others (potentially including the interpreter). It is the scorekeeping social practices that actually govern the use of an expression (in particular the acknowledgment of entitlement to the commitments undertaken by its assertional use) that supply what is missing from pure reliability theories.

It is tempting, from the point of view of such theories, to think of the choice of reference class as a merely pragmatic matter—in a sense of 'pragmatic' that restricts it to what concerns the interests and goals of those performing speech acts. So it might be thought that for some purposes and in some contexts I should be counted as knowing that a sparrow is in front of me, even though I would believe that also if a sufficiently lifelike replica were there instead, while for other purposes and in other contexts (for instance where the stricter standards appropriate to discussions of principled skepticism are in force) I should not. No doubt there is such a variation in standards of entitlement depending on what is taken to turn on the issue, and it may be particularly acute in connection with the word 'know'. But the contribution made by interpreters (those who attribute commitments and entitlements to commitment) to the determination of the boundaries with respect to which reliability is assessed are not "merely pragmatic" in this sense.

They make a fundamental contribution to the semantic content of empirical concepts. Indeed, this is one of the situations in which traditional ways of distinguishing semantic from pragmatic concerns can be seen to be inapposite. In particular, as will appear, what an interpreter takes to be the circumstances under which an expression can appropriately be used in noninferential reports—that is, when interlocutors are *entitled* to commitments because the acknowledgment of the commitment arises through the exercise of dispositions to respond differentially to various aspects of their environment—is an important feature of the empirical content the interpreter associates with that expression. The sort of authority accorded to noninferential reports, and the way the model of assertional practice can be extended to incorporate it, is discussed further below.

III. OBSERVATION REPORTS AND NONINFERENTIAL AUTHORITY

1. Knowledge, Entitlement, and Understanding

The topic of reliability theories of cognitive entitlement was introduced in connection with the thought that once the notion of entitlement

or positive justificatory status that matters for attributions of knowledge has been broadened by the recognition that a belief, claim, or commitment can in this sense *be* justified without having *been* justified—that justify*ing* is not the only way that status can be acquired—the way seems open to dispensing entirely with inferential justifying in explaining the deontic status of entitlement. It was then pointed out that identifying the entitlement of a commitment with its being the output of a reliable process has the same sorts of difficulties with gerrymandering that plague its relatives that identify the correctness of a claim or application of a concept with its being in accord with a regularity exhibited by other claimings or applications of the concept. But these difficulties concern only one way of following out the original thought about the in-principle dispensability of inferential justification in explaining the status of knowledge claims. The deep mistake involved in completely decoupling justifying as giving reasons from cognitive entitlement has to do rather with the sort of *understanding* that is presupposed by claims to and attributions of knowledge.

An assertion, even if true, is not taken to express knowledge unless the one making it *understands* the claim being made. A practical grasp of the significance of making the claim is inseparable from an appreciation of its role as possible reason for other claims, and as something that reasons can in turn be offered for. It is being caught up in this way in the game of giving and asking for reasons that makes a performance the undertaking of a commitment (the making of a claim) in the first place. Unless one accords one's own performance such a significance (treats it as a move in that game), one is not making a claim, not undertaking a commitment that is eligible for the status of knowledge.

It is on this basis that Sellars objects to construals of cognitive entitlement exclusively in terms of reliability. Reliable differential responsive dispositions are only a necessary condition for observational knowledge. Parrots and thermometers can have such dispositions and so can be used by us in effect as measuring instruments to acquire knowledge. But what *they* have is not knowledge. For they do not understand the significance of their responses; they do not take those responses as reasons for further claims; and they do not understand claims as potentially in need of reasons. To decouple entitlement from reason-giving entirely is to jettison the inferential articulation in virtue of which the performances and commitments one is entitled to can be understood as *propositionally* contentful. It is to discard precisely what makes responses, however reliably produced, have the significance of undertaking *discursive* commitments. What is left is not a *cognitive* affair at all.

The most serious objection to a pure reliability theory accordingly is presented not by the general Wittgensteinian strand of thought concerning the significance of gerrymandering for attempts to construe norms as regularities, which Kripke expounds so forcefully. The most serious objection to reliabilism stems rather from the more particular Sellarsian insight concern-

ing the essential role played by the specifically inferential articulation even of noninferential reports. Sellars insists first that for a performance elicited by a reliable differential responsive disposition to be a candidate for expressing knowledge, it must count as an *endorsement* by the reporter of some claim, as the undertaking of a *commitment.* Furthermore, he recognizes that the identity and content of such commitments depend on their role in inference and justification, in giving and asking for reasons. He sees further that being capable of endorsing a claim requires grasping the role of that claim in inference and justification—that the official tripartite analysis of knowledge implicitly involves *understanding,* as part of what is required for *belief.*

Unfortunately, motivated by these insights, Sellars stakes out far too strong an antireliabilist position concerning the role of inferential justifying in entitlement to claims to observational knowledge. There is accordingly a danger that where the various strands of thought are not carefully sorted out, distaste for the epistemological internalism about cognitive authority that Sellars endorses will obscure the important lessons that ought to be drawn from his account. His basic point is that a noninferential reporter must be "in the space of giving and asking for reasons," in addition to having the right differential responsive dispositions. That space is, for Sellars as here, articulated by relations of *authority,* and to be in that space one must be able to recognize or acknowledge the authority of claims. It is, in other words, in virtue of one's capacity to adopt practical deontic attitudes, to take or treat something as having cognitive authority, that one counts as moving in the space of giving and asking for reasons.

2. *Sellars on the Authority of Noninferential Reports*

In the passages (from "Empiricism and the Philosophy of Mind") that follow, Sellars is concerned with the nature of the authority (or as he sometimes puts it, "credibility") of noninferential reports (Carnap's *Konstatierungen*), which are claims to observationally acquired knowledge. The distinctive feature of such reports is that "the credibility of such tokens as 'express observations' [is] a credibility which flows from tokens to types."[12] This contrasts with the credibility of nonobservational claims such as "Dogs are mammals," which is attached to tokens in virtue of their being of types that are credible. The reliability approach then recommends itself as offering a simple and natural account of the source and nature of the credibility of sentence tokenings that report empirical observations: "An overt or covert tokening of 'This is green' in the presence of a green item is a *Konstatierung* and expresses observational knowledge if and only if it is a manifestation of a tendency to produce overt or covert tokens of 'This is green'—given a certain set—if and only if a green object is being looked at in standard conditions."[13]

What does such reliability have to do with *authority?* "The first hurdle to

be jumped concerns the *authority* which, as I have emphasized, a sentence token must have in order that it may be said to express knowledge. Clearly, on this account the only thing that can remotely be supposed to constitute such authority is the fact that one can infer the presence of a green object from the fact that someone makes this report."[14] This is an important move. The authority of reliability consists in its underwriting a propriety of *inference* (what might be called "the reliability inference"). The noninferential undertaking of a commitment by a reliable reporter can *inferentially authorize* another to undertake a commitment with that content. To take or treat someone as a reliable reporter (in certain circumstances) is for a scorekeeper to endorse the propriety of the move from *attributing* to the reporter a noninferentially acquired doxastic commitment to the scorekeeper's *undertaking* of a corresponding commitment (and taking others to be similarly entitled). This notion of how reliability fits into the giving of reasons is the key to understanding the special sort of authority characteristic of noninferential reports, which in turn is essential to the notion of *empirically* contentful claims.

As is by now familiar, Sellars has already taken issue with the sort of foundationalism that sees empirical knowledge as an inferential superstructure raised on an autonomous noninferential base. The target of his criticism

> is the idea that there is, indeed *must be*, a structure of particular matter of fact such that
>
> (a) each fact can not only be noninferentially known to be the case but presupposes no other knowledge either of particular matter of fact, or of general truths; and
>
> (b) the noninferential knowledge of facts belonging to this structure constitutes the ultimate court of appeals for all factual claims—particular and general—about the world.[15]

Sellars, inferentialist and antifoundationalist though he is, does not deny either (b) or the first half of (a). His quarrel is with the second half of (a).

There are particular instances of believing or being committed that are noninferential in the sense that their acquisition was not the conclusion of an inferential process. There are no beliefs or discursive commitments that are noninferential in that what is expressed by a sentence can be *understood* without mastering inferential relations that content stands in to others. So a bit of knowledge (belief) can, and indeed all of it does, presuppose other knowledge (belief), even though it is not inferred from that other knowledge or belief. This possibility was not seriously examined by the classical epistemological tradition. It is a certain hierarchical picture of *understanding* (at this level a necessary condition of believing) that Sellars rejects. He does *not* object to a hierarchical picture of empirical *justification*, once that has been suitably disentangled from bad foundationalism concerning understanding. His claim that the authority that accrues to noninferential reports in virtue

of their being the results of reliable reporting or belief-acquiring mechanisms is a broadly inferential authority is in no way inconsistent with understanding observational knowledge to be authoritative in virtue of the reliable *noninferential* differential responsive dispositions that produce it. Inference need not be involved in the process that leads to a tokening of 'This is green'; but it is involved in grasping the type of authority that such noninferentially produced tokenings have, and so in understanding such tokenings, and so in their being potential expressions of knowledge.

The question is just what the relation is between mastery of this inference and possession by a tokening of the sort of authority characteristic of ground-level observational knowledge. Sellars's claim is that "to be the expression of knowledge, a report must not only *have* authority, this authority must *in some sense* be recognized by the person whose report it is."[16] The notion of claims *having* cognitive authority is indeed intelligible only in connection with practical attitudes of *taking* or *treating* claims *as* having such authority. Sellars has suggested that the authority distinctive of observational knowledge should be understood in terms of the correctness of an inference, from the making of a report such as 'This is green' by one whose differential responsive dispositions are taken to be reliable to the undertaking of a commitment to the effect that there is something green there.

He concludes: "In other words, for a *Konstatierung* 'This is green' to express observational knowledge, not only must it be a *symptom* or *sign* of the presence of a green object in standard conditions, but the perceiver must know that tokens of 'This is green' *are* symptoms of the presence of green objects in conditions which are standard for visual perception."[17] So Sellars's view is that the reliable reporter can count as being entitled to a noninferentially acquired commitment, and so the assertion by which that commitment is acknowledged can be cognitively authoritative in licensing or entitling others by the standard assertional mechanism of communicative entitlement inheritance, *only* if the reporter can inferentially *justify* the noninferential claim. Such a justification consists precisely in exhibiting the inference whose premises are the reliability of differential responsive dispositions to make such claims and responsive elicitation of the claiming in question and whose conclusion is another tokening of the claim itself.

3. *Attributing Reliability Is Endorsing an Inference: An Inferentialist Middle Way between Justificatory Internalism and Reliabilist Externalism*

There are two problems with this conclusion. First, Sellars takes it that for the claim of the reliable observer to *be justified*, the observer must be able to *justify* it inferentially—to offer reasons by displaying premises from which it follows. Second, he assumes that such justification must involve explicit invocation of reliability, that is, that a claim of reliability

must be one of the premises. Thus Sellars claims that "observational knowledge of any particular fact, e.g. that this is green, presupposes that one knows general facts of the form *X is a reliable symptom of Y.*"[18]

This latter is an odd move for Sellars to make. He, after all, is the one who urged that material proprieties of inference not be everywhere traded in for suppressed premises. Why should it not be that reliability underwrites a propriety of inference, without the claim of reliability having to appear as a premise in the inferences so underwritten? Here Sellars overreaches himself. He is right to insist that reliability matters because it warrants inferences of a certain form. He is wrong to insist that this warranting must be understood in terms of endorsement of an explicit claim that can serve as a premise in inference.

The first claim also seems too strong. Securing entitlement to a claim need not always be assimilated to inferential justifying of the claim. The possibility of vindication of a commitment by deference rather than inference—by the invocation of communicational mechanisms depending on intracontent interpersonal testimony rather than intrapersonal intercontent justifying—shows that much. It would be a mistake to assimilate deferential entitlement inheritance to inferential entitlement inheritance by insisting that the one invoking the authority of another's assertion be able to produce an explicit *argument* in which a claim as to the informant's reliability would appear as a premise. Rather, deferring involves an *implicit* claim as to the reliability (more particularly the entitlement in this case) of the informant. One who accepts the deferring as successful (and so attributes the claimed entitlement to the one deferring) thereby implicitly endorses the propriety of a permissive inference from the informant's claiming that *p* to *p*—which is just the inference that Sellars picks out as corresponding to reliability. But at the ground level, all of this can be made sense of as implicit in what is done in practice. Because it can, it is possible to explain the expressive role of the locutions that can be introduced at a later stage to make these attitudes explicit. Assimilating entitlement to the commitments acknowledged by noninferential reports, and therefore their authority, to that secured by explicit inferential justification, as Sellars does, is a mistake of the same sort.

A symmetrical mistake would be to assimilate the authority of noninferential reports to that of testimony, by understanding the invocation of such authority as a kind of deference to a "world-asserter."[19] The structure of authority exhibited by noninferential reports is sui generis, to be reduced neither to that of inferential justification nor to that of testimony. These three are mutually irreducible—none can take over the function of any of the others. One of the primary explanatory aims of this work is to explain how commitments that are implicit in the fundamental practices that confer assertible conceptual content can eventually themselves be made explicit and assertible, expressed in a form in which reasons can be given and asked for them. The implicit attitudes that can in this way be explicitly expressed

once suitable vocabulary has been introduced include those involved in invocations and recognitions of the authority of both testimony and reports of observations. But the locutions that play these explicitating roles can themselves be made intelligible only by understanding first the implicit structures they bring out into the assertional light of day.

In fact, Sellars's insight concerning the irreducible role played by inferential justification does not require insisting that noninferential reporters can be authoritatively entitled to their claims only if they can justify them. As just indicated, the inference from the noninferential undertaking of a commitment as the result of a reliable differential responsive disposition to endorsement of the claim thereby made may be implicit in practical attitudes, rather than explicit in claims offered as justifications for that endorsement. Nor is it necessary that the one who makes an observation report endorse the propriety of that inference, even in this implicit practical sense. Reliability may entitle the reporter to the knowledge claim, may qualify it as knowledge, even if the reporter does not even implicitly endorse the inference that is the practical acknowledgment of the authority of reliability. This is the primary insight that stands behind the justificatory *externalism* of reliability epistemologies, in contrast to the justificatory *internalism* Sellars exemplifies.

Suppose that Monique has been trained reliably to discriminate hornbeams by their leaves. As a result of the training, she is often disposed to respond to the visibility of leaves of the right sort by noninferentially reporting the presence of a hornbeam. She understands what it means to claim that something is a hornbeam and, in circumstances appropriate for such reports, actually comes to believe that there is a hornbeam present. She may still be uncertain of her discriminatory capacity long after she has in fact become reliable.[20] In such a situation she may have a true belief that there is hornbeam in front of her, yet be completely unable to justify that claim (for instance, by citing features distinctive of hornbeam leaves), and even deny that she is a reliable noninferential reporter of hornbeams.

Yet, the reliabilists point out, it can be entirely in order for one who does take her to be a reliable reporter of them, not only to come to believe that there is a hornbeam present on the basis of her report, but to cite her report (at least deferentially) as what warrants that belief. This is treating the claim as authoritative in just the way that is required for knowledge. Someone who thus takes her to be reliable can accordingly attribute to Monique the knowledge that there is a hornbeam in front of her, in spite of her protestations to the contrary. What makes her claim knowledge (according to the attributor) is the fact of her reliability (according to the attributor), regardless of her attitudes toward that reliability.[21] The status of her claim as knowledge is accordingly external to her attitudes—not only because of the truth condition on knowledge, but also because of the entitlement condition. Sellars is committed to withholding the attribution of knowledge in the absence of the

candidate knower's capacity to justify the claim, and so is committed to disagreeing with reliabilists about examples like this. Yet on this point the reliabilists are surely correct.

Sellars, however, is right that for a reliably elicited differential response to be a candidate for knowledge, the one making the knowledge claim must be in the space of reasons, must be capable of understanding the claim, and so must have some grip on its role in reasoning, hence on its use as a premise and conclusion of inferential justifications. Requiring this general capacity, of course, falls short of requiring that on each occasion the reporter must be able to justify the claim for it to count as the expression of observational knowledge. Furthermore, while reliabilism about cognitive entitlement and so cognitive authority is clearly correct that knowledge can be attributed even where the one to whom it is attributed cannot demonstrate entitlement to the claim inferentially, by providing a justification that appeals to other claims the putative knower endorses, it does not follow from this observation that reliability by itself is enough for entitlement and cognitive authority, apart from all consideration of attitudes of taking or treating the knower as reliable, as a thoroughgoing externalism about entitlement would have it. Sellars is also right to insist that attributions of knowledge require not just reliability but at least implicit endorsement of the inference that is the practical acknowledgment of reliability—the inference namely from the occurrence of a report, or the noninferential undertaking of the commitment such a report expresses, to the endorsement of the claim.

Where Sellars is wrong, as the sort of example just considered shows, is in thinking that the one who endorses this inference must be the one who *undertakes* the claim to observational knowledge. It has been pointed out that attributing knowledge is a hybrid deontic attitude involving not only the attribution of commitments but the undertaking of them. Not only does the attributor of knowledge take the candidate knower to endorse a claim; the attributor also must endorse that claim.

It is likewise the *attributor* of observational knowledge who must attribute reliability to the knower. Attributing such reliability is endorsing exactly the general form of permissive inference that Sellars points to—treating as appropriate the inference from the noninferential undertaking of a commitment (of the right sort, and in circumstances of the right sort) by the observer to the endorsement by others of the claim so elicited. Taking a report that is the outcome of a particular differential responsive disposition as entitling others to the claim (for instance accepting as entitling their deferrals to the reporter on such issues) just *is* treating the reporter in practice as reliable about such matters. Monique need not, *pace* Sellars, take herself to be a reliable reporter of hornbeams in order for her to count as knowing observationally that there is a hornbeam in front of her. But the one who attributes such knowledge must take her to be reliable. And adopting that practical attitude is endorsing the pattern of permissive inference that connects the

attribution by others to Monique of a noninferentially acquired belief about the visible presence of hornbeams with their *undertaking* of a commitment to the visible presence of hornbeams in Monique's vicinity.

Just as the truth condition on knowledge requires that the attributor of knowledge undertake, as well as attribute, commitment to the content of the knowledge claim, so satisfying the entitlement condition by mere reliability requires that the attributor of knowledge undertake (but not necessarily attribute) commitment to the propriety of the reliability inference. Where the language is rich enough to include the expressive resources necessary to make the reliability inference explicit (conditionals and 'claims that . . .' or 'believes that . . .'), attributors of knowledge can be challenged and called on to defend their endorsement of the conditional "If Monique claims (sincerely, responsively, and in appropriate conditions) that a hornbeam is visibly present, then (probably) a hornbeam is present." At this point, reliability could be invoked to *justify* the belief that there is a hornbeam present. But this is a sophisticated, late-coming possibility, built on the implicit acknowledgments already described. So full-blown reliabilist externalism about cognitive entitlement is mistaken in ignoring the necessity of such inferential attitudes on the part of attributors of knowledge, while full-blown Sellarsian internalism about cognitive entitlement is mistaken in insisting that the knower must have such attitudes. These are complementary ways of misunderstanding the essentially social structure of the cognitive deontic attitudes, in terms of which the status of a claim as knowledge must be understood.

4. *Observational Knowledge and Empirical Conceptual Content*

The noninferential authority possessed by claims issuing from the exercise of reliable differential responsive dispositions—although reducible neither to the sort of interpersonal authority invoked by deferring nor to the sort of intercontent authority invoked by inferring—is not fundamental in the way that those structures of authority are. In the model of assertional practice that has been put on the table, communication and justification are two aspects of the game of giving and asking for reasons; neither is intelligible except in the context of the other. They are intelligible, however, in the absence of noninferential responsive authority. Practices that do not involve according any knowledge claims the significance of observation reports can nonetheless be understood as instituting specifically *assertional* significances, and so as conferring specifically *propositional* contents.

What is missing from such practices is claims with *empirical* content. Discourse recognizable as mathematical can be like this: reasons are given and demanded; claims communicated, challenged, and justified; and regresses of entitlement inheritance halted by appeal to axioms, free moves that anyone is treated as entitled to at any point in the conversation. Our

discourse is not in general like this, however, and the sorts of contents our claims have cannot be conferred by assertional practices that do not acknowledge some claims as having empirical authority stemming from their status as reports of observations. Indeed, it is essential to the contents of the ordinary concepts in terms of which we conduct our lives that they stand in inferential relations both to the acknowledgments of commitments resulting from what Sellars calls "language entry transitions," in perception, and to the acknowledgments of commitments that result in what Sellars calls "language exit transitions," in intentional action.[22] (The contribution of the latter *practical* empirical structure is discussed in the second half of this chapter. Attention is restricted here to the *cognitive* empirical structure.)

The practical significance characteristic of claims to observational knowledge is best understood in terms of the role they play in the default-and-challenge structure of entitlement. Noninferential reports can function as unjustified justifiers: claimings that are treated as having a defeasible default status as entitled. Properly made claims to observational or perceptual knowledge can accordingly provide entitlements that can then be inherited either inferentially or communicationally. So observation provides regress-stoppers, and in this sense a foundation for empirical knowledge. This is what stands behind Sellars's endorsement of the claim (quoted above) that "noninferential knowledge of facts . . . constitutes the ultimate court of appeals for all factual claims—particular and general—about the world."[23]

Default entitlements are of two sorts, depending on whether the entitlement attaches to a commitment in virtue of the type it instantiates or in virtue of the circumstances in which it is tokened.[24] There are sentence types that would require a great deal of work for one to get into a position to challenge, such as "Red is a color," "There have been black dogs," "Lightning frequently precedes thunder," and similar commonplaces. These are treated as "free moves" by the members of our speech community—they are available to just about anyone any time to use as premises, to assert unchallenged. Noninferential reports, by contrast, have their default entitlement status as a result of the way in which the report tokening, or the particular acknowledging of the commitment that would be expressed by such a tokening, is elicited through the exercise of a reliable differential responsive reporting disposition. Treating such a claim as one the reporter is entitled to involves an implicit commitment on the part of the attributor to the actual circumstances being among those in which the reporter is responsively reliable concerning the sort of matters reported.

There will typically be some sorts of reports such that under appropriate reporting conditions (the same for all), essentially all the members of the linguistic community are reliable. Almost anyone can, under suitable circumstances, tell whether it is a warm day out or whether the marble one is holding is approximately round. Other sorts of reports involve not only more specialized circumstances but specialized training. Particle physicists are

trained reliably to respond noninferentially to the presence of mu-mesons in a bubble chamber by reporting the presence of mu-mesons. Not all of us can do this reliably. It is only someone who is taken not only to be looking at a bubble chamber but also to be properly trained to be reliable about these matters (and who has the right sort of collateral beliefs) whose reports will be accorded noninferential entitlement and the corresponding authority. As Quine says, what is observable varies from community to community.[25] He understands the status of being an observation report for a community (perhaps a proper subset of the whole linguistic community) in terms of what that community can agree on under concurrent stimulation, that is, in the same standard reporting circumstances.

The authority of noninferential reports requires the collaboration of both dimensions into which Sellars analyzes them: not only that they arise from the exercise of reliable differential responsive dispositions but that the response is to endorse a claim, to acknowledge a commitment, with a certain content. What makes it a mu-meson that the physicist is reporting rather than the hooked vapor trail that also forms part of the reliably covarying chain of events culminating in the report is to be understood not in terms of the differential responsive dispositions but in terms of the inferential role of the claim being made. The consequences that can be inferred from the presence of a mu-meson are quite different from those that can be inferred from the presence of a hooked vapor trail, covariant and concomitant though these phenomena may be. For instance, mu-mesons are much smaller, and move much faster, than the vapor trails they produce (see further at 7.1.6 below). As Quine argues further, it is important to understand that under the appropriate circumstances, which include the presence of a bubble chamber or similar device, and for the right community of observers, mu-mesons are literally observable—noninferentially reportable in much the same sense in which red things are for the rest of us. It is a mistake to think that what is really noninferentially observed is only the vapor trail and that the presence of mu-mesons is only *inferred.* Such an inference can be made, and learning to make it might be part of the training process that leads to becoming a reliable observer of mu-mesons (in bubble chambers). But coming to be disposed reliably to respond to the vapor trail, and *hence* to the presence of mu-mesons, by asserting or acknowledging a commitment to the presence of a mu-meson is learning to *observe* mu-mesons, to report them *non*inferentially. And this is so even if one is not totally reliable, in that there are circumstances in which one would *mistakenly* report the presence of a mu-meson *because* of the presence in the chamber of a vapor trail indistinguishable from those one has learned to respond to noninferentially by reporting mu-mesons (just as the fact that one can be fooled by a cunning replica does not preclude one from seeing a sparrow in the cases where one is not being fooled).

The claim is, then, that one is *directly* observing mu-mesons, in the sense

of *noninferentially* coming to be aware of them, to make claims about them, to know something about them—rather than indirectly, inferentially coming to a conclusion about mu-mesons on the basis of an *inference* (perhaps unconscious or implicit) from the presence of a vapor trail with a certain shape. This claim may seem implausible in light of the common practice of retreating, under certain suitable sorts of challenge, from the claim that a mu-meson is present to the claim that a hooked vapor trail is present. Is not such a retreat to be understood as relinquishing commitment to an inference, and therefore to its conclusion, while continuing to defend its (genuinely noninferential) premise? No.

Such cases ought to be understood as retreats (given a credible challenge to the effect that this might be one of the cases in which the exercise of a generally reliable capacity nonetheless leads one astray) to a claim that is *safer.* Being safer in this sense, however, is not a matter of withdrawing endorsement of an inference. One retreats to a different report with respect to which one is *more* reliable, as measured for instance by percentage of correct differential responses in the relevant circumstances, or by the same percentage of correct responses within a wider range of circumstances, or by the size of the community that does not share esoteric theoretical beliefs but does share the differential responsive disposition and corresponding capacity to make noninferential reports. Doing this can amount to offering an inferential justification of the original noninferential belief, by explaining how one was able to see it. An analogous case would be explaining that there was a mirror, not apparent to the audience assessing the authority of the claim, but apparent to the reporter, in order to explain how one was noninferentially able to report something that a challenger has pointed out is around a corner and so ought to be invisible.

The possibility of such an inferential justification of a claim on the basis of a safer claim does not show that the original claim should be understood as itself the product of a process of inference, any more than the capacity of sophisticated reporters to offer justifications of their claims to observational knowledge by citing their reliability and appealing to the reliability inference shows that their original claim was arrived at as a result of an inference from that premise. Nor does the fact that the capacity to make certain sorts of noninferential reports depends on collateral beliefs show that those reports are really inferences from something more basic, together with those collateral beliefs. One must have many beliefs about mu-mesons in order to be able to understand and so to make *any* claims about them, noninferential or otherwise. That does not preclude one from coming to be able to observe them.

The basis of observational knowledge, then, is that it should be possible to train individuals reliably to respond differentially to features of their environment by acknowledging doxastic commitments. Those commit-

ments are inferentially related to others that not only play inferential roles but also are themselves appropriately elicited noninferentially by features of the environment. These cross-connections put constraints on endorsements of inferences relating expressions whose circumstances of appropriate application include noninferential ones. Both oranges and orange things can be noninferentially reported, so someone who reports the presence of an orange after tasting but not seeing it and then infers from its being an orange to its being orange in color is liable to be challenged by another who is in a position to report it noninferentially as purple. For the commitment entitlement to which was acquired noninferentially is incompatible with that entitlement to which was acquired inferentially. Either the identification of the orange by taste, which formed the premise of the inference, or the identification of its color as purple might itself bear further challenge and investigation; but if these stand up, the reliability of the inference from being an orange to being orange in color will be impugned. In this way the possession of noninferential circumstances of appropriate application of some concepts imbues them with empirical content—recognizable as conceptual content in virtue of its inferential articulation and as empirical in virtue of its dependence on the noninferential acquisition of commitments to those contents (and of entitlements to those commitments).

Similarly, the inferences from circumstances to consequences of application (which are implicit in conceptual contents) are subject to empirical criticism in virtue of inferential connections among the contents of commitments that can be acquired noninferentially. So it may happen that one uses the term 'acid' in such a way that a substance's tasting sour is a sufficient condition for applying it, and that it will turn litmus paper red is a necessary consequence of applying it. Finding a substance that both tastes sour and turns litmus paper blue shows that such a concept is inadequate. Conceptual contents can accordingly be criticized, groomed, and developed empirically in a way parallel to the sort of Socratic process discussed in Chapter 2. In virtue of their inferential connections to concepts that can be used to make reports, even purely theoretical concepts (those whose *only* circumstances of appropriate application are inferential) inherit empirical content and have the inferences they are involved with constrained by the commitments and entitlements actually thrown up by what is responded to noninferentially. That the reliable differential responsive dispositions underlying this structure of noninferential authority are dispositions to acquire *commitments* and *entitlements* to those commitments, that is, to alter deontic status, means that the practices they appear in must include corresponding practical deontic *attitudes*. Something practitioners can do must be the taking or treating of performances as having the significance of noninferential reports, the recognition of the status of some claims as deriving their entitlements from their being expressions of reliable differential responsive dispositions to ac-

knowledge commitments. For one cannot make sense of normative significance, even the normative significance of reliability, apart from consideration of its uptake or attribution.

5. *Attributing Observational Entitlement*

It is straightforward to extend the model of assertional practice as outlined so far to encompass the structure of authority in virtue of which claims can have and be treated as having the significance of noninferential reports. What is primarily required is to say what it is for one interlocutor to attribute noninferential or observational authority to the claim of another, thereby recognizing or acknowledging it as having a special sort of entitlement. The authority involved is entitlement heritable by the usual intrapersonal intercontent inferential and interpersonal intracontent communicational pathways. What is distinctive of observational authority is that such authority is accorded to particular *tokenings* of acknowledged commitments (rather than to their types) and the way in which that authority depends on a special combination of content-based and person-based features. For the imputed reliability of an observer varies from content to content within each observer, and from observer to observer—someone who is taken to be able reliably and noninferentially to discriminate mu-mesons in bubble chambers may not be taken to be able to do so for '52 Pontiacs in traffic.

Furthermore, if the topic is fixed (the concepts used in the reports being assessed are specified), imputed reliability still varies depending on circumstances. For each particular observable, there will be an associated set of appropriate circumstances of reporting, according to the one attributing reliability and so observational authority. The authority of a reliable reporter is conditional on the obtaining of those appropriate circumstances. Even a lookout who is in general a reliable reporter of whales must be facing in the direction of what is being reported, cannot see well in the direction of a horizon-hugging sun, is less reliable if there are large walruses about, and so on. These appropriate circumstances of reporting, associated with the observable content (and perhaps the individual reporter), figure as commitments *undertaken* by the one attributing or assessing the responsive authority of a claiming.

So associated with each sort of noninferential authority a given interlocutor grants to another (the product of a person and a kind of content, for example, reports of the presence of whales, or mu-mesons), there is a set of enabling conditions (looking in the right direction, looking in a bubble chamber) and a set of defeating conditions (presence of many walruses in the vicinity, physicist drunk and woozy), and it is the interaction of these, according to the commitments undertaken by the one assessing noninferential authority, that determine whether responsive entitlement is attributed or

not. If the assessor undertakes commitment to a suitable range of the enabling reporting conditions and does not undertake commitment to any of the defeating reporting conditions, the reporter's claim is treated as having a default status as entitled. In this way the empirical authority of some attributed commitments, on the basis of implicit inferential acknowledgment of reliability under suitable circumstances, is distinguished from the type-based default status of "Red is a color" and "There have been black dogs," which do not exhibit a similar relativity to person, content, and the environing conditions as they are taken to be by the assessor. Observational authority is accordingly another hybrid deontic status: attributing it involves not only *attributing* commitments and entitlements but also *undertaking* or *acknowledging* them by endorsing reliability inferences.

Once the attitude of taking or treating someone's performance as having the significance of a noninferential report whose authority is grounded in the local and conditional reliability of the observer is in place, it is possible to introduce a type of performance that is the claiming of or petitioning for such authority by an observer. A certain sort of noise or gesture (perhaps a shrug) can come to have the significance of *invoking* observational authority. Then if a report is challenged, it can be vindicated (its entitlement demonstrated) by invoking observational authority rather than by deferral or inferential justification. But there is no strict need for practices encompassing empirical conceptual contents to include a speech-act kind with this significance. It is enough if interlocutors sometimes accord such authority, and thus take the commitments acknowledged by noninferential reports in some circumstances to be vindicated (implicitly, according to the assessor's attitudes) by the fact of the reporter's reliability. Where there is such a speech act, it would implicitly mean something like "I *see* it (for example, that it is red)." An explicit assertion to this effect can be introduced as well, but just how will not be clear until Chapter 8, where pragmatically explicitating locutions such as 'believe that' and 'claim that' are officially introduced into the model of assertional practices. The significance of an invocation of observational authority does not depend on any assertionally explicit content that the invocation might have. (Compare Wittgenstein's suggestion that if challenged to say how one knows that the thing in front of one is red, one might say simply, "I speak English.")

6. *Expressions of Belief That Are Not Claims to Knowledge*

The account here of assertions as claims to *knowledge* turns on the implicit obligation to vindicate the commitment undertaken by demonstrating one's *entitlement* to it. The foregoing discussion of observation and reliability focused on the importance for the *attributor* of observational knowledge claims of implicitly attributing reliability. Adopting that attitude requires endorsing the inference from the attribution of a noninferentially

responsively elicited acknowledgment of a *commitment* (under suitable circumstances and for a qualified observer) to the attribution of *entitlement* to that commitment. It is in the context of concern with entitlement to assertional commitments that the complaint was levied against reliabilists that they ignore the inferentially articulated attitude in which recognition or attribution of entitlement consists. The corresponding objection to Sellars was that, while appreciating the significance of that hybrid practical deontic attitude, he inappropriately insists that the reliability inference it involves be endorsed by the one *making* the observation rather than the one attributing or assessing it.

But this concern with entitlement can seem out of place in a discussion of a sort of discursive commitment that is intended to do the sort of explanatory work characteristically performed by a notion of *belief.* If belief is to be understood in the first instance as the state or status expressed by assertional speech acts, it seems wrong to treat assertions as also involving a claim to *knowledge.* For expressing a belief and claiming to know are different.

When an idiom is developed to the point that it has the expressive resources provided by the English locutions 'believes that' and 'knows that'—which make the pragmatic status being attributed or undertaken explicit as part of the *content* of what is claimed—it becomes possible to say of someone else, "He believes that Arnauld did not write *The Art of Thinking,* but he does not know it." The case has already been considered where what is expressed is the attitude of a scorekeeper who attributes commitment to a claim but does not endorse the attributed claim—that is, does not take it to be true. It is also possible, however, to distinguish expressions of mere belief from claims to knowledge in the first-person case, in which the claim *is* being endorsed or taken-true. In such cases, the social-perspectival distinction between attributions of knowledge and attributions of belief cannot get a grip.

For although undertaking an assertional commitment *is* taking-true the claim, a difference can arise precisely over the issue of entitlement or justification. The attribution of knowledge may be withheld by a scorekeeper who attributes a commitment without attributing a corresponding entitlement. Indeed, sometimes we make claims while fully aware that they may legitimately be challenged and that we are not in a position to vindicate them by demonstrating our entitlement to them. This is the implicit attitude that becomes assertionally explicit in claims such as "I believe that Arnauld did not write *The Art of Thinking,* but I don't claim to know it." For this sort of reservation can concern not the truth of the belief but my capacity to justify it. I may continue to take the claim to be true, to endorse it, to acknowledge the commitment it expresses, and yet not be prepared to shoulder the justificatory burden associated with a knowledge claim. This might be because I have forgotten the source of my conviction, or it might be an expression of a claim's having a ground-level status for me as an unjustified justifier that I do not take to be widely shared—I just believe that people with beards cannot be trusted, or that house cats are dangerous.

The speech acts that express such attitudes are what might be called *bare* assertions,[26] corresponding to *mere* beliefs, without the implicit claim to entitlement that is demonstrable should someone become entitled to challenge it (paradigmatically by expressing an entitled commitment to a claim incompatible with it). Does not the possibility of such bare expressions of commitment without claim of entitlement, of conviction without warrant, show that it is a mistake to understand claims on the model of claims to *knowledge*? No. Such claims are intelligible only as exceptions against a background of practices in which claims typically have the significance of claims whose authority *is* redeemable by demonstration of warrant. The possibility of bare assertion is parasitic on the possibility of assertions that implicitly involve undertaking a conditional task-responsibility to demonstrate the asserter's entitlement to the commitments undertaken by the performance of speech acts of that kind.

For bare assertions and the commitments they express would be completely idle if they could not figure as premises in inference and could not be passed along in communication. This is to say that bare assertions involve something of the *authority* of full-blooded assertions, while disavowing the corresponding *responsibility*. Yet that authority (licensing inferences by the asserter to commitments with other contents and the undertaking of commitments with the same contents by other interlocutors) makes sense only in a context in which inferential and deferential invocation of such authority can be demanded. What assertions are *for* is justifying other assertions. To accept someone's bare assertion is to take it to be a claim from which conclusions can be drawn. But giving reasons presupposes the possibility of asking for them, or at least the possibility that claims often stand in need of reasons. A game of giving and asking for reasons cannot consist exclusively in the exchange of speech acts that are accorded the significance of bare assertions. Within the broader context of full-blooded assertions (which do involve the demonstration of entitlement by inference and deference), however, it is possible to make sense of treating some claims as having the significance of bare assertions. Assertional commitments essentially involve the dimension of entitlement. Assertions are paradigmatically knowledge claims, and the sort of belief they express is unintelligible except in relation to the possibility of assessing beliefs for their status as knowledge, as warranted and true.

IV. RATIONAL AGENCY

1. *Methodological Constraints on the Conception of Practical Rationality*

Beliefs make a difference both to what we say and to what we do. They manifest themselves both linguistically, in assertions, and practically,

in actions. A basic criterion of adequacy for any theoretical account of this fundamental sort of intentional state is that it explain both of these ways in which beliefs can be expressed in behavior, and the relation between them. The methodologically parsimonious idea that one or the other of them ought to be accorded explanatory priority is the motivation common to both of what Stalnaker (in the rough-and-ready botanization alluded to in the previous chapter) distinguishes as the "linguistic" and the "pragmatic" approaches to intentionality. Theories of the sort he calls "linguistic" construe believing by analogy to claiming: as a kind of inner asserting of sentences. They are accordingly obliged, first, to explain assertions without appeal to their role as expressions of belief and, second, to explain the norms that determine the role of belief in rational agency in terms of the proprieties that govern the public use of sentences. Theories of the sort he calls "pragmatic" (such as the one Stalnaker himself endorses), in contrast, take the role of belief in intentional action to be primary. They then owe both a nonlinguistic explanation of rational agency and an account of speech acts, paradigmatically assertion, in terms of intentional states so understood.

The approach pursued here takes belief to be intelligible only in the context of social-linguistic practice. But it is a relational, rather than a reductive, linguistic theory. Although doxastic commitment (the sort of deontic status corresponding to the intentional state of belief) cannot be made sense of apart from the possibility of expressing such commitments by performing speech acts that have the significance of assertions, neither can assertional significance be made sense of without reference to the commitments such speech acts undertake and acknowledge. As regards asserting and believing, the theory is even-handed; it accords explanatory priority to neither one. It nonetheless deserves to be called a *linguistic* account of intentionality (in a sense broader than Stalnaker's[27]) inasmuch as it does accord explanatory priority to the linguistic manifestation of belief in assertion over its practical manifestation in action.

As Dennett and Davidson have emphasized, attributing propositionally contentful intentional states such as beliefs to a creature is taking it to be *rational.* Thus Kant uses the rubrics of *theoretical* and *practical* rationality to distinguish the sort of normative competence manifested in giving and asking for reasons for *claims* or *judgments* from the sort of normative competence manifested in giving and asking for reasons for *actions*—judgments and actions being picked out precisely as the sorts of things reasons can be given for and for which reasons can be asked. Being rational is understood here generically as being able to play the game of giving and asking for reasons, which is to engage in a specifically *linguistic* social practice. For one cannot give reasons unless one can make claims. Doing so requires mastery of the normative dimension of inference: a practical grasp of the notion of *right* reasoning, of the distinction between correct and incorrect inference. Assessing performances as correct or incorrect is adopting normative atti-

tudes that are intelligible only in a context of interpersonal scorekeeping—even though in such a context it is possible to make sense both of self-assessment and of assessments of objective correctness, for which no one's scorekeeping attitudes are counted as decisively authoritative.[28] First-person deliberation is the internalization of such third-person assessment.

To take this line—identifying the rationality that qualifies us as sapients with being a player in the normative game of offering and assessing, producing and consuming reasons—is to deny two widely held reductive conceptions of rationality: one that identifies rationality with *logical* competence, and another that identifies it with prudence or *instrumental* competence. All parties can agree that to be rational is to distinguish good inferences from bad inferences. The disagreement concerns whether 'good inference' in this formula can be restricted to *logically* good inferences, or again to *instrumentally* good inferences—ones whose correctness is determined by their utility in satisfying desires or maximizing preferences. Logical competence is mastery of the use of locutions by means of which inferential proprieties are made *explicit* as the contents of claims. This theoretical ability to codify practices as principles accordingly presupposes prior practical mastery both of *implicit* inferential proprieties and of the use of the ordinary, nonlogical claims they articulate and govern. When this expressive role of logical vocabulary is appreciated, the identification of rationality in general with its manifestation as logical manipulation is unmasked as another form of the intellectualism that insists on discerning a propositionally explicit principle underlying every implicit propriety of practice—a form of platonism whose remedy is a complementary pragmatism.

Identifying rationality in general with the sort of instrumental rationality manifested in rational agency also inverts the proper order of explanation.[29] For the propositional contents of the intentional states appealed to in practical reasoning presuppose assertional-inferential proprieties, and hence linguistic social practices. (Though to say this is not to deny that proprieties of practical inference also contribute to the propositional contents of the states and expressions caught up in them.) To make out this claim it is necessary to say something about the practical reasoning that is implicitly attributed in interpretations of individuals as rational agents. In particular, just as it was shown how the capacity for logical reasoning is to be made intelligible in terms of (as the explicitation of) a conceptually prior capacity for nonlogical reasoning, it must also be shown how the capacity for practical reasoning incorporates and depends upon a conceptually prior capacity to give reasons for claims, rather than for actions.

Only an account of assertion of the sort introduced in Chapter 3 leaves room for the pursuit of such an order of explanation. Everyone ought to agree that asserting is putting forward a sentence as true. Following Davidson's lead, it has been suggested that distinguishing practical attitudes as taking or treating something *as true* requires a specifically linguistic social context of

mutual interpretation—that is, attribution of doxastic commitments, of the sort that has been elaborated as assertional-inferential scorekeeping. The next chapter develops the idea that this principle is best exploited by starting with an antecedent notion of assertional significance and then moving via that principle to an understanding of what is involved in talk of truth.

Commitment to understanding rational agency in terms of linguistic practice, rather than the other way around, strongly constrains the construal of the putting-forward portion of the principle that asserting is putting forward a sentence as true. For the claim that our first grip on the paradigmatic intentional state of belief (taking-true) is as what is expressed by assertions—rather than as what makes certain nonlinguistic performances intelligible in a way that is made explicit by exhibiting a piece of practical reasoning—is evidently incompatible with understanding asserting instrumentally, as a means intentionally adopted by a rational agent in order to achieve certain desired ends. If the linguistic practice of making and assessing claims (the game of giving claims as reasons and demanding reasons for claims) is an essential element of the context required to make sense of the notion of *propositional* intentional content (assertible, believable contents, which in English can be made explicit by the use of 'that' clauses), then what has been called "agent semantics" is not entitled to the conceptual raw materials it employs. In particular, one may not appeal to the *intentions* of the asserter—for instance intentions to say something true, or to make the audience believe that what is said is true, or to make the audience believe that it is uttered with the intention of saying something true or of engendering the corresponding beliefs. For what one is attributing can be identified as intentions to bring about various states of affairs only in virtue of the role such states play in a larger practical whole—one that includes the possibility of attributions of beliefs that the corresponding states of affairs obtain.

Less obviously, this order of explanation also precludes appeal to *conventions*, at least as commonly understood. The influential account offered by Lewis, for instance, takes a convention to be a social regularity that is sustained in a special way by the beliefs, intentions, and desires of the parties to the convention.[30] They are required not only to conform to the regularity but, among other conditions, to believe that others do so, to conform themselves because of that belief, to prefer that everyone conform, and to believe that everyone else has such beliefs and preferences. The present view is that on such a construal of convention, as Davidson concludes, "philosophers who make convention a necessary element in language have the matter backwards. The truth is rather that language is a condition for having conventions."[31] Construing the putting-forward bit of the principle that for a sentence to have assertional significance is for it to be put forward as true in terms of social conventions rather than individual intentions is also not an acceptable move according to this explanatory strategy. Certainly conventions of the sort that Dummett tries out—conventions to the effect that one

is to be understood as trying to utter sentences only with the intention of uttering true ones—are of no avail in the context of these explanatory commitments.[32] That is why it was necessary to move beyond explaining asserting in terms of *intentions* or *conventions* to explaining it instead in terms of *practices*, which themselves can be explained without appeal to intentions or conventions.

The next task is to show how that account of practices can be extended so as to encompass deontic statuses corresponding to the other sorts of intentional states that figure in the giving of reasons for *non*linguistic performances: the intentions and desires that play an essential role in the *practical* reasoning implicitly attributed by interpretations of individuals as rational *agents*. The aim is to provide a broadly Kantian account of the will as a *rational* faculty. By exploiting the analogy between discursive *entry* transitions in *perception* and discursive *exit* transitions in *action*, the rational will can be understood as no more philosophically mysterious than our capacity to notice barns or red things. A scorekeeping account can pick out performances (largely nonlinguistic ones) as intentional (under some specification) and hence as actions (under any specification) insofar as they are expressions of deontic attitudes—acknowledgments of a certain kind of commitment. Practical reasoning can then be understood as leading to performances with this sort of deontic significance. And on that basis, the expressive role of distinctively *normative* vocabulary can be specified in terms of its role in making explicit the endorsement of patterns of practical reasoning.

2. *Acting and Perceiving*

The general claim is that there are two species of discursive commitment: the cognitive and the practical. Acknowledging commitments of these two sorts is adopting deontic attitudes that correspond to the intentional states of *believing* and *intending*, respectively. A practical commitment is a commitment to *act*. The content of such a practical commitment is to *making*-true a claim. These commitments and their contents are intelligible only in a context that includes also the *taking*-true of claims. For it is in terms of such assertional taking-true that the success of actions, the fulfillment of practical commitments, must be understood. The category of cognitive discursive commitments accordingly enjoys a certain explanatory priority over that of practical discursive commitment. Each is essentially something that reasons can be given for and for which reasons can be asked, and one cannot give reasons unless one can acknowledge doxastic commitments by making claims.

The practical dimension of discursive practice can be understood by exploiting two ideas. The first is that practical commitments are like doxastic commitments in being essentially inferentially articulated. They stand in

inferential relations both among themselves and to doxastic commitments. The second idea is that the noninferential relations between acknowledgments of practical commitments and states of affairs brought about by intentional *action* can be understood by analogy to the noninferential relations between acknowledgments of doxastic commitments and the states of affairs that bring them about through conceptually contentful *perception.* The causal dimension of acting for reasons—acknowledging practical commitments by acting on them—involves the exercise of reliable differential responsive skills on the *output* side of the game of giving and asking for reasons, just as perception does on the *input* side. Elaborating the first idea involves examining the sense in which practical reasons are *reasons;* elaborating the second idea involves examining the sense in which practical reasons are *causes.*

Adding practical commitments to the model of discursive practice enriches the propositional contents that such practice can be understood to confer on states and their expressions in a way analogous to the enrichment provided by including the empirical authority of observationally acquired doxastic commitments. In each case the general category of assertional commitments and their contents can be understood in advance of the enrichment. The three structures of authority that the model of assertion, as presented thus far, comprises are mutually irreducible, but not all are equally fundamental. The inferential authority invoked by justification and the testimonial authority invoked by deference are intelligible apart from the default authority of noninferential reports; but inferential and deferential practice are two sides of one coin, apart from which the authority of noninferential reports is not intelligible. Thus *empirical* content represents an enrichment of the generic sort of propositional content specifiable in abstraction from the contribution of observation. Similarly, *practical* content represents an enrichment of the generic sort of propositional content specifiable in abstraction from the contribution of action. The empirical and practical involvements of claims—even those that are purely theoretical in the sense that they are only inferentially connected to claims that have direct empirical and practical significance—make a fundamental contribution to their contents. Only a model that incorporates both of these not purely inferential dimensions of discursive articulation has any prospect of generating propositional contents that resemble those expressed by the declarative sentences of natural languages.

The best way to understand the place of action in the deontic model of discursive practice is to exploit the analogy between action and perception. Sellars divides the "moves" that can be made in a language game into three kinds: intralinguistic moves, language entry moves, and language exit moves.[33] The first kind consists of inferential moves. These are moves in which a position within the language game (paradigmatically the endorsement of a claim) is responded to by the adoption of another such position.

The second kind consists of noninferential reports of observations. These are moves in which a nonlinguistic situation is responded to by the adoption of a position within the language game (paradigmatically the endorsement of a claim). The third kind consists of deliberate actions. These are moves in which a position within the language game (for instance, endorsement of a plan) is responded to by bringing about a nonlinguistic situation.

Following Sellars's lead, language entry moves have been analyzed in the first three sections of this chapter in terms of two components in their content: their inferential articulation and their noninferential elicitation. In virtue of the former they are *conceptually* contentful, and in virtue of the latter they are *empirically* contentful. These components and their interaction have been elaborated here in the idiom of deontic scorekeeping, into which Sellars's framework has been transposed. In that idiom, noninferential reports count as *entries* in the sense that they are responses that consist in changes of deontic scorekeeping attitude, elicited by stimuli that do not themselves consist in changes of deontic score. As such they contrast with inferential moves, in which an alteration of deontic attitude—for instance the undertaking or attributing of a commitment—has as a scorekeeping consequence another alteration of deontic attitude. The language or discursive scorekeeping *exits* (intentional *actions*) are to be understood by analogy to these *entries* (perceptual *observations*). In action, alterations of deontic attitude, specifically acknowledgments of practical commitments, serve as stimuli eliciting nonlinguistic performances.

Observation depends on reliable dispositions to respond differentially to states of affairs of various kinds by acknowledging certain sorts of commitments—that is, by adopting deontic attitudes and so changing the score. A competent observer under suitable circumstances responds to the visible presence of a red ball by coming to acknowledge a commitment to the claim that there is a red ball present. The content of the commitment responsively undertaken is jointly determined by the chain of reliably covarying events that culminates in its acquisition and by its inferential connection to other contents (including those empirical conceptual contents that themselves incorporate a responsive observational component). Action depends on reliable dispositions to respond differentially to the acknowledging of certain sorts of commitments (the adoption of deontic attitudes and consequent change of score) by bringing about various kinds of states of affairs.[34] A competent agent under suitable circumstances responds to the acquisition of a commitment to flip the light switch by flipping the light switch. The content of the commitment so expressed is jointly determined by the chain of reliably covarying events that its acknowledgment initiates and by its inferential connection to other contents (including both other contents that themselves incorporate a practical component and those empirical conceptual contents that incorporate a responsive observational component).

In any given situation, interlocutors can be taught to be reliable noninfer-

ential reporters of only certain sorts of circumstances. Unaided by special instruments, we cannot reliably discriminate the presence of X rays, and we cannot tell automobiles that will at some point in the future be painted green from those that will not; we can reliably discriminate the presence of loud noises, and we can tell automobiles that are now painted green from those that are not. Similarly, interlocutors can be taught to be reliable performers of only certain kinds of acts. Unaided by special tools, we cannot reliably produce X rays, and we cannot make an automobile have been painted green some time in the past; we can reliably produce loud noises and paint automobiles green. What can be noninferentially reported varies from reporter to reporter and from situation to situation. Only a properly trained physicist can noninferentially observe the presence of a mu-meson, and then only with a bubble chamber; only a properly trained pianist can noninferentially produce a performance of the Moonlight Sonata, and then only with a piano.

Observation requires reliable responsive dispositions to acquire acknowledged commitments, while action requires reliable responsive dispositions to fulfill acknowledged commitments. Reliability in the first case concerns the relation between the state of affairs responded to and the content of the commitment acknowledged. Reliability in the second case concerns the relation between the content of the commitment acknowledged and the state of affairs brought about. In each case, assessments of reliability require some independent access to the eliciting or the elicited state of affairs—assessments of the *truth* of the claim the perceiver has noninferentially come to make and of the *success* of the performance the agent has noninferentially come to produce. Attributions of reliability consist in endorsements of scorekeeping inferences *from* commitments *attributed* to reporters or agents *to* commitments *undertaken* by the attributor of reliability (commitments concerning the state of affairs reported or produced). Thus my noninferentially acquired doxastic commitment to the effect that there is a red thing in front of me is, under appropriate conditions, a good reason for others inferentially to acquire a doxastic commitment to the effect that there is a red thing in front of me. My acknowledging a practical commitment to the effect that I will raise my arm in the next minute is, under appropriate conditions, a good reason for others to undertake a doxastic commitment to the effect that I will raise my arm in the next minute.

In observation, the elicited commitment-acknowledgment is an attitude toward a *doxastic* discursive deontic status. In action, the eliciting commitment-acknowledgment is an attitude toward a *practical* discursive deontic status. The first sort of attitude corresponds to believing or taking-true—in *one* sense of believing, namely the causally relevant sense that depends on what one would *acknowledge* commitment to, not the ideal sense in which if p entails q, then believing that p *is* believing that q, whether one knows it or not. The second sort of attitude corresponds to intending or making-true—in *one* sense of intending, namely the causally relevant sense that depends on what one would *acknowledge* commitment to, not the ideal sense in

which if doing *A* entails doing *B*, then intending to do *A* *is* intending to do *B*, whether one knows it or not.

The wider ideal senses of 'believe' and 'intend' correspond to the deontic *statuses* of doxastic and practical commitment, rather than to the deontic *attitudes* of acknowledging them. These senses are to be understood in terms of the fundamental scorekeeping principle that *undertaking* a commitment (to begin with, by acknowledging it) licenses others to *attribute* it, and the attributions that are thereby authorized can outrun what one is disposed to acknowledge. The only function of the concept *deontic status* in the idiom in which the model of discursive practice is formulated is its use in keeping score. It is a creature of the activity of scorekeeping on deontic attitudes; deontic statuses figure only as the objects of attitudes, as what is undertaken and attributed.

Understanding practical discursive commitments (commitments to act) is accordingly a matter of understanding their pragmatic significance: the way they depend on and influence the deontic score interlocutors keep by acquiring and relinquishing attitudes toward their own and others' deontic statuses. Practical commitments, like doxastic or assertional commitments (including noninferentially acquired empirical ones), are discursive or conceptually contentful commitments in virtue of the *inferential* articulation of their pragmatic significance. The scorekeeping significance of practical commitments is analogous to that of doxastic commitments—indeed the inferential and incompatibility relations that the contents of practical commitments stand in are largely inherited from those of corresponding doxastic commitments, except for their role in the sort of practical reasoning that connects them inferentially with doxastic commitments proper. Thus one practical commitment can have others as consequences; a commitment to drive to the airport today entails a commitment to go to the airport today, because the inference from '*X* drives to *s*' to '*X* goes to *s*' preserves *doxastic* commitments. In the same way, one practical commitment can be incompatible with another, as are a commitment to drive to the airport today and a commitment to spend the day snoozing in a hammock under a shade tree—again because of the incompatibility of the corresponding doxastic commitment contents.

The instrumental inferences corresponding to the principle "Who wills the end wills the means," like inferences generally, come in two flavors: committive and permissive. Some instrumental inferences (those whose premises specify goals one is committed to and whose conclusions specify the necessary means to those ends) are also commitment-preserving. If cutting down a tree is the only way to get across the ravine, then undertaking or attributing a commitment to getting across the ravine has as a scorekeeping consequence undertaking or attributing a commitment to cutting down a tree. But some means-end reasoning is permissive in nature; there may be more than one way to skin a cat.

Inferences whose premises express commitments to secure certain ends and whose conclusions express sufficient (but not necessary) means to those

ends are entitlement-preserving rather than commitment-preserving. One who is entitled to a practical commitment to secure an end is entitled thereby to a practical commitment to performances that would (according to the one whose scorekeeping is being elaborated) bring about that end—in the absence (as is always the proviso with permissive inferences) of collateral commitments incompatible with such a commitment. Entitlement to a practical commitment to achieve some end may simultaneously entitle one to each of a set of mutually incompatible alternative means; entitlement to a commitment to cross the ravine may instrumentally entitle one to cut down the tree at the edge of the ravine, and it may entitle one to anchor a rope bridge to the top of that tree, even though doing one of these things precludes doing the other. In the same way, permissive inferential relations (paradigmatically inductive ones) among the contents of doxastic commitments can result in entitlement to each of a set of incompatible conclusions. In each case, choosing one, committing oneself to a conclusion or a means, relinquishes entitlements to those incompatible with it. For incompatibility is a relation involving both deontic statuses: two contents, whether doxastic or practical, are incompatible in case commitment to one precludes entitlement to the other. Entitlement to both without commitment to either is not ruled out. Neither, of course, is commitment to both; making this possibility straightforwardly intelligible is one of the cardinal advantages deontic normative construals of belief have over causal-functional ones.

3. *Asymmetries between Practical and Doxastic Discursive Commitments*

It is in their relations to their corresponding entitlements that practical discursive commitments differ most markedly from doxastic discursive commitments. The significance of undertaking a doxastic commitment, paradigmatically through its overt acknowledgment by assertion, was explained in terms of the interactions between the coordinate dimensions of authority and responsibility. The responsibility involved is to vindicate the commitment, by demonstrating or displaying one's entitlement to it, if it is brought into question by a suitable challenge (an incompatible assertion with an equal, prima facie claim to entitlement). Default entitlements aside, this responsibility can be discharged by appeal to the authority of other doxastic commitments; credentials for the commitment are secured by displaying its entitlement as inherited from that attached to other commitments.

The authority of the commitments undertaken by assertion exhibits a dual structure, corresponding to two different sorts of routes by which entitlement can be passed on for use in discharging the responsibility associated with other commitments. On the one hand, a doxastic commitment to which one interlocutor is entitled licenses further commitments (with different

contents) by that same interlocutor. These are its *inferential* consequences (committive and permissive). This sort of authority is invoked to vindicate those consequential commitments by presenting a *justification*, which appeals to the authorizing claims as premises. On the other hand, a doxastic commitment to which an interlocutor is entitled licenses further commitments with the same content, by other interlocutors. This is its authority as *testimony*. It is invoked to vindicate the commitments it authorizes, by *deferral* to the one whose testimony is relied upon.

The first way in which the structure governing the attribution of entitlements to practical discursive commitments differs from that governing the attribution of entitlements to doxastic ones is that there is nothing corresponding to the authority of testimony in the practical case. The issue of entitlement can arise for practical commitments, as for all discursive commitments. But the (conditional) responsibility to vindicate such commitments is, in the practical case, exclusively a *justificatory* responsibility. Default entitlements aside, it is only by exhibiting a piece of reasoning having as its conclusion the practical commitment in question that entitlement to such commitments can in general be demonstrated or secured.

This feature of the deontic scorekeeping model of discursive commitments reflects a fundamental asymmetry between expressing a belief by making a claim and expressing it by performing an action. What I take-true I thereby, ceteris paribus, authorize you to take-true. Though there can be various complications about the transfer of title (because of differences in collateral beliefs), in general what serve me as good reasons for belief can serve you also as good reasons for that same belief. What I (seek to) make-true, however, I do not thereby in general authorize you also to (seek to) make-true. What serve me as good reasons for action may or may not be available to you as good reasons for action, even bracketing differences in collateral beliefs. For you and I may have quite different ends, subscribe to different values, occupy different social roles, be subject to different norms. That I have good reasons to drive to the airport today does not imply that you do. If you form a similar intention, you cannot in general show that you are entitled to it by deferring to me ("Well, *he's* going"). Only some kinds of reasons that entitle me to an intention and action are automatically available to you. You might be in a position to make the same argument I can, but if so, that in general is independent of my being in a position to use that line of thought; there is no general (even defeasible) presumption of heritability.[35]

Committing oneself to a claim is putting it forward as *true*, and this means as something that everyone in some sense *ought* to believe (even though some unfortunates will for various reasons not be in a position to do so and need not be blameworthy for that failure). Committing oneself to a course of action need not be like this. It need not (though in special cases it can) involve putting it forward as something that everyone else ought to do (even subject to the recognition that some unfortunates will for various reasons not

be in a position to appreciate this, and need not be blameworthy for that failure). Some kinds of reasons for actions, paradigmatically moral ones, have a permissive or committive force that is independent of interpersonal differences. But reasons for action in general do not have this kind of force. What an agent has reason to do can depend on what that agent wants (or on what institutional role that agent occupies). Differences among agents as to desires and preferences (or institutional roles) need not have the significance of indications of normative failures. Whenever two believers disagree, a diagnosis of error or ignorance is appropriate for at least one of them. Though agents with differing practical commitments can also be criticized on the grounds of error and ignorance, mere difference of desire or preference is not sufficient in general to make them liable to such criticism.

We come with different bodies, and that by itself ensures that we will have different desires; what is good for my digestion may not be good for yours; my reason to avoid peppers need be no reason for you to avoid peppers. Our different bodies give us different perceptual perspectives on the world as well, but belief as taking-true incorporates an implicit norm of commonality—that we should pool our resources, attempt to overcome the error and ignorance that distinguish our different sets of doxastic commitments, and aim at a common set of beliefs that are equally good for all. Talk about belief as involving an implicit commitment to the Truth as One, the same for all believers, is a colorful way of talking about the role of testimony and challenge in the authority structure of doxastic commitment—about the way in which entitlements can be inherited by others and undercut by the incompatible commitments they become entitled to. The Good is not in the same way One, at least not if the focus is widened from specifically moral reasons for action to reasons for action generally, so as to include prudential and institutional goods. Desires and preferences can supply reasons for actions (can entitle agents to practical commitments) in the sense of 'entitle' that corresponds to that at stake in the discussion of doxastic commitments, and desires and preferences can vary from individual to individual. That there is no implicit normative commitment that plays the same role with respect to desire (and therefore intention and action in general) that truth plays with respect to belief consists simply in the absence (in the structure according to which entitlements to practical commitments are inherited) of anything corresponding to the interpersonal dimension of testimony and vindication by deferral.

It is of course possible to add an interpersonal dimension of practical authority as a superstructure to the basic game of giving and asking for reasons for actions. Where within a certain sphere of practical activity the performance of one individual licenses or compels performances by others, there exists an authority relation of superior to subordinate. In a practice in which reasons can be given (and so asked for) at all—that is, a linguistic practice, one in which some performances are accorded the significance of

assertions—the authorizing performances can be speech acts with the significance of imperatives and permissives. The superior issues an order, which specifies what the subordinate is obliged to do (what the subordinate thereby acquires a commitment to do), by displaying the assertion that must be made-true, the assertible content of the doxastic commitment that anyone must be entitled to undertake (perhaps observationally) upon completion of the task. Or the superior offers permission in the form of a licence, which specifies what the subordinate is entitled to do by displaying the assertion that can be made-true.

This sort of practical authority structure is like that of testimony in some ways. When the issue of the agent's entitlement to a practical commitment (perhaps claimed implicitly by deliberate action) is raised, rather than the entitlement being inherited from reasons that could be cited for it—either by the agent in terms of doxastic and practical commitments undertaken or by other scorekeeping assessors in terms of doxastic and practical commitments attributed to the agent—that entitlement can be inherited from the superior who ordered or permitted it. Such authority can be invoked by deferring to the issuer of the command or license (a mode of vindication codified in the legal doctrine of respondeat superior). So besides intrapersonal entitlement inheritance invoked by inference, there can be a mode of interpersonal entitlement inheritance invoked by deference, in the practical as well as the doxastic case.

There are many disanalogies between these two cases as well, however. First, the licensing is restricted as to subject matter and the interlocutors involved, to those situations in which a prior superior/subordinate authority relation has been established. The employer can authorize or compel only certain sorts of performances, and only on the part of certain individuals. Perhaps this difference does not go very deep. The limitation is characteristic of a society in which such authority relations are established and limited by explicit contracts. In a society based on status rather than contract, the superior/subordinate relations are fixed once and for all in advance, appearing as part of the nature of things, and need not be restricted as to subject matter at all. Furthermore, such restrictions can arise, de facto or even de jure, in the case of assertional authority as well: the teacher of secret doctrines may not authorize their repetition or answer for them except to favored students. And there are what amount to hierarchies of assertional authority regarding technical topics such as mu-mesons and quarter horses.

One difference that does go deep, however, is an asymmetry between the authorizing performance and the authorized performances, in the case of commands and the issuing of licenses. The asserter licenses members of the audience to perform speech acts with just the same content and significance as the original assertion. They are authorized to authorize others in the same sense in which they are authorized. Assertion, at least as it is construed in the ideal Sprachspiel presented here, is an egalitarian practice in a sense in

which commanding and giving permission is not. Only in very special cases does the practical license one is given authorize the further issuance of such licenses; only in very special cases does the command one is given compel or permit one to offer such commands to others.

The structure of default entitlement and calling to account by challenging entitlements to action are similarly asymmetric in the case of interpersonal practical authority structured by superiors and subordinates. There is in general nothing corresponding to assessments of reliability underlying the default authority of superiors (though analogs exist for special cases). Entitlement to challenges must similarly be relativized to superior/subordinate relations, if commands and licenses are to have any significance. The point of rehearsing these asymmetries is just that the fundamental differences between doxastic and practical structures of authority and entitlement inheritance remain even in the case where the normatively significant social status of individuals as superior or subordinate is widely or universally instituted by the attitudes of those keeping score on commitments and entitlements. If the asymmetries characteristic of superior/subordinate relations are removed, making interpersonal practical authority look more like assertional authority as here conceived, nothing remotely resembling the issuing of orders or the giving of permission results.

In the interests of simplicity, the deontic scorekeeping model of assertional significance has been talked about as though assertional authority is always made universally available throughout the community and is always universally recognized. Where testimony has this sort of catholic significance, the community can be thought of as engaged in the search for a single common body of truths, for anyone's entitlement to any claim is open to challenge from any quarter. Doxastic practice need not be so monolithic, of course. There may be many subcommunities, distinguished precisely by what sorts of authority they acknowledge, and so what sorts of challenges to entitlements they take to be in order. Specialists may recognize the authority only of other specialists. Members of one speech community may be divided into competing schools of thought on various topics and may not recognize the entitlements or therefore the challenges of those from other groups, as regards claims concerning those topics. Yet within those subcommunities it is essential that the authority granted by an assertion include a reassertion license—a license to do just what the asserter did. This feature makes it possible for the claims of one interlocutor to have the significance of challenges to the claims of another.

The importance of this structure is particularly evident in the case of empirical practice, for it is by testimony that observations by one interlocutor can be assessed and adjusted by confrontation with the observations of another. The notion of entitlement to a doxastic commitment depends on the in-principle heritability of interpersonal authority. Because an assertion that would be defended by appeal to testimony can have just the same

entitlement status as one that would be defended by providing a justification or by the invocation of noninferential responsive reliability (observational prowess), the credentials of each claim do not need to be traced back to their source before it can be treated as having the significance of a prima facie challenge to incompatible claims. This is a basic feature of the assertional default-and-challenge structure. That a status or performance whose entitlement is inherited from another should have just the same authority as the status or performance that authorized it (according to the subcommunity that recognizes such authority) accordingly distinguishes doxastic discursive commitments from practical discursive commitments. For as has been pointed out, if subordinates have the same authority as their superiors in virtue of being commanded or licensed by them, the entitlement of a superior to issue a command would be subject to challenge by commands issued by subordinates, not just by other superiors, and the hypothesized asymmetry between superior and subordinate would disappear.

V. PRACTICAL REASONING: INFERENCES FROM DOXASTIC TO PRACTICAL COMMITMENTS

1. Acting for Reasons and Acting Intentionally

Giving and asking for reasons for *actions* is possible only in the context of practices of giving and asking for reasons generally—that is, of practices of making and defending *claims*. The structure of entitlement characteristic of practical discursive commitments is not autonomous but presupposes that of doxastic ones. This dependence appears in two ways. On the side of the circumstances of acquisition of practical deontic statuses, it appears in the role of practical reasoning; practical reasoning requires the availability of doxastic commitments as premises. On the side of the consequences of acquisition of practical deontic statuses, it appears in the essential role that propositional (= assertible) contents play in specifying conditions of success—that is, what counts as fulfilling a commitment to act. With regard to this latter role, it has already been pointed out that practical commitments inherit some of their inferential relations from the propositional contents that specify their conditions of success. If doxastic commitment to p has as a scorekeeping consequence doxastic commitment to q, then a practical commitment to make-true p has as a scorekeeping consequence a practical commitment to make-true q. Understanding what one has committed oneself to by undertaking a practical commitment to bring it about that p accordingly requires mastery of the inferential role p plays in doxastic discursive practice.

The relation between doxastic and practical commitments that is most important for extending the deontic scorekeeping account to include both

species of discursive status, however, is that exhibited in practical reasoning. Because interpersonal inheritance of entitlements is not an essential part of the scorekeeping structure that institutes practical deontic discursive statuses, the conditional responsibility to demonstrate entitlement that is part of undertaking a commitment to act is a specifically *justificatory* responsibility. Justifying a practical commitment is exhibiting a suitable piece of practical reasoning in which it figures as the conclusion. It is in terms of practical inferences that we give reasons for action, make our own and each other's conduct intelligible, exhibit it as rational. Practical reasoning accordingly forms the core of intentional explanations of nonlinguistic deportment.

In what follows, an abstract account of practical reasoning is sketched in the deontic scorekeeping terms familiar from the treatment of theoretical reasoning concerning doxastic commitments.[36] The explanatory framework in which the notion of practical reasoning is to function is the Kantian one, according to which to treat a performance as an *action* is to treat it as something for which it is in principle appropriate to demand a *reason*. Not everything an agent does is an action. If I am walking along the top of a cliff and stumble and fall off, stumbling and accelerating at 32 feet per second per second are both things I do (in the sense that they are bits of my behavior), but they are not actions of mine; walking and grabbing a bush as I topple over the edge are. Actions are the things agents do *intentionally*. In the terms to be employed here, acting intentionally is noninferentially producing a performance that either is the acknowledgment of a practical commitment (in the case of intentions in action) or results from exercising a reliable differential disposition to respond to such an acknowledgment (in the case of prior intentions). The acknowledgment of the practical commitment can be thought of as the intention with which the performance is produced.

One can act *with* a reason, but unintentionally (for instance in a case in which one is unaware of the commitments that supply the reason that an attributor might cite). But only what is done intentionally can be done *for* a reason—though one can act intentionally but without a reason.[37] Only rational beings can be agents, but there are such things as irrational actions: for instance where one acts intentionally, but on impulse rather than according to what one has reason to do. In the deontic framework, such irrational actions are *intentional* in that they are acknowledgments of practical *commitments* (or arise from the exercise of reliable noninferential dispositions to respond differentially to them), and they are irrational in that the practical commitment in question is not one the agent is *entitled* to by a good practical inference from premises that agent is committed and entitled to—either because one has *no* reason or because one has an overriding reason to do something incompatible with what one in fact does. Since to be so entitled requires having a *reason* for performing the action, practical commitments, and therefore actions (intentional performances), are attributed only to those who are in the space of giving and asking for reasons—that is, to those who are (treated as) rational.

Undertaking any discursive commitment involves a conditional responsibility to demonstrate entitlement to it. In the case of practical commitments this takes the form of a specifically justificatory responsibility. Only against the background of a general capacity to comprehend and fulfill such a justificatory responsibility—to assess and produce reasons for practical commitments—can what one does have the significance of an acknowledgment of a practical commitment, that is, the significance of acquiring or expressing an intention. Given such a general capacity or status as rational, however, one can in particular cases undertake practical commitments to which one is not entitled by reasons, and so act irrationally. Intentional but irrational actions are perfectly intelligible within the deontic framework, in the same way and for the same reasons that, on the side of doxastic rather than practical discursive commitments, incompatible beliefs are—namely as commitments lacking the corresponding entitlements. These phenomena cause explanatory difficulties for other sorts of accounts (for instance those that construe intentional states exclusively in terms of causal-functional role), difficulties that simply do not arise when those states are construed in terms of deontic statuses instituted by scorekeeping attitudes.

To be entitled to a practical commitment is to have suitable reasons for it. Practical inferences—as distinct from the doxastic inferences that have been considered in previous chapters—are those that have practical commitments as their conclusions.[38] Reasons for such commitments, and hence for the actions elicited by the acknowledgment of such commitments, are the premises of good practical inferences. It has already been pointed out that intentions can serve as reasons for other intentions—the intention to bring it about that *p* serving as a reason for intending to bring it about that *q* if that *q* is true is necessary or sufficient for bringing it about that *p*. What about reasons for commitments to act that are not themselves commitments to act? Facts, as acknowledged in doxastic commitments, can provide reasons for practical commitments.

2. *Three Patterns of Practical Reasoning*

Consider the following three bits of practical reasoning:

(α) Only opening my umbrella will keep me dry,
so I shall open my umbrella.

(β) I am a bank employee going to work,
so I shall wear a necktie.

(γ) Repeating the gossip would harm someone, to no purpose,
so I shall not repeat the gossip.

'Shall' is used here to express the significance of the conclusion as the acknowledging of a practical commitment.[39] The corresponding doxastic commitment would be acknowledged by a standard assertion using 'will'.

The role of the speech acts performed by uttering sentences of this form can be understood in scorekeeping terms from their fundamental pragmatic significance as acknowledging a practical commitment, together with the inferentially articulated content that results from combining its involvement in the inferences deriving from the corresponding 'will' statements with its involvement in the sort of basic practical inferences of which (α), (β), and (γ) are examples.[40]

There are two ways to think about inferences like these, which move from doxastic premises to practical conclusions, from beliefs to intentions. What is perhaps the standard approach is that taken by Davidson.[41] He defines a *primary* reason as the pair of a belief and what he calls a pro-attitude. He allows that sometimes one or the other is cited by itself as a reason, but insists: "In order to understand how a reason of any kind rationalizes an action it is necessary and sufficient that we see, at least in essential outline, how to construct a primary reason."

In other words, inferences such as those exhibited by (α), (β), and (γ) are *enthymemes*, in which a premise necessary for the correctness of the inference has been suppressed or omitted. In the first inference, what is missing is some such premise as:

(a) Let me stay dry,

an expression of a desire, preference, or pro-attitude that would be explicitly self-ascribed by something like "I desire (prefer) to stay dry."[42] The second inference might be underwritten by something like:

(b) Bank employees are obliged (required) to wear neckties.

In the third case, the suppressed premise is something such as:

(c) It is wrong to (one ought not) harm anyone to no purpose.

Each of these supplies the missing pro-attitude required to make the premises into primary reasons.

As appears in these examples, the notion of pro-attitude encompasses not only wants, desires, and preferences but also more general evaluative attitudes. This assimilation represents an important insight, which will be exploited below. In fact, Davidson thinks that all pro-attitudes are expressed by sentences that are in a broad sense normative or evaluative.

> There is no short proof that evaluative sentences express desires and other pro-attitudes in the way that the sentence "Snow is white" expresses the belief that snow is white. But the following considerations will perhaps help show what is involved. If someone who knows English says honestly "Snow is white," then he believes snow is white. If my thesis is correct, someone who says honestly "It is desirable that I stop smoking," has some pro-attitude towards his stopping smoking.

> He feels some inclination to do it; in fact he will do it if nothing stands in his way, he knows how, and he has no contrary values or desires. Given this assumption, it is reasonable to generalize: if explicit value judgments represent pro-attitudes, all pro-attitudes may be expressed by value judgments that are at least implicit.[43]

Davidson thinks of evaluative expressions as expressing something like desires, but the connection can equally well be exploited in the other direction. For conversely, one who desires or prefers p to q (say desires that one eat pears rather than that one eat peaches) thereby attaches some value to p over q, takes p to be preferable to, or more desirable than, q. What is important is to see that normative claims and expressions of desire and preference are species of a genus defined by the role they play in completing primary reasons.

Pro-attitudes must be included in primary reasons, on this account, to bridge the gap between what one believes and what one decides to do. My preference to stay dry makes my belief that I can stay dry only by opening the umbrella relevant to the practical issue of whether to open the umbrella. The fact that bank employees are obliged to wear neckties makes my working at the bank relevant to the practical issue of whether to wear a necktie. And the negative value of causing pointless harm (the fact that it is wrong or that one ought not to do it) makes the consequence of gossiping relevant to the practical issue of whether to gossip. In the context of different pro-attitudes, those same beliefs would provide reasons for quite different intentions and actions.

3. *Normative Vocabulary Makes Explicit Material Proprieties of Practical Reasoning*

There is another way of construing the relation between (α), (β), and (γ), on the one hand, and (a), (b), and (c), on the other hand. That relation could be modeled on the relation between *materially* good inferences and the *conditionals* whose addition as premises would turn them into *formally* (logically) good inferences. In that case (considered in Section IV of Chapter 2), it turned out to be a fruitful strategy to consider the apparently enthymematic inferences as in order just as they stood, and to treat the conditionals not as suppressed *premises* but as making *explicit* (expressing in the form of a claim) what is *implicit* in the endorsement of the inferences. Part of the payoff from considering things this way around is an understanding of the expressive role played by conditionals; they can be understood as making *inferential* commitments propositionally explicit (= assertible). What makes that approach possible is an account of proprieties of inference as deontic social statuses instituted by scorekeeping attitudes, so that commitment to a material propriety of inference can be understood in terms of what it is to

take or treat an inference as correct in keeping score by attributing and acknowledging assertional commitments.

There is no bar to understanding (α), (β), and (γ) as materially good inferences in this sense. The fact that endorsement of claims incompatible with (a), (b), and (c) would void these inferences does not show that they function as suppressed premises—any more than the fact that endorsement of $p \rightarrow \sim q$ would void the inference from p to q shows that the conditional $p \rightarrow q$ is a suppressed premise in the material inference from p to q. The claims (a), (b), and (c) might, like $p \rightarrow q$, be understood rather as codifying material-practical inferential commitments. The payoff from doing so would be making it possible to understand the expressive role played by the broadly evaluative words (such as 'prefer', 'obliged', and 'ought') used to express these pro-attitudes, in a way analogous to the understanding suggested for conditionals. According to such an account, although the sort of practical inference instanced by (α), (β), and (γ) does not need supplementation to be correct, in a language with sufficient expressive resources it is possible to make the inferential commitments that are implicit in endorsing such inferences explicit in the form of claims.

The benefits of doing so are the familiar benefits of propositional explicitness: once expressed in the form of claims, these commitments are themselves subject to challenge and justification, rather than simply being accepted or rejected. Two interlocutors who disagree about the correctness of an inference such as (γ) can now argue about whether (c) is true, challenge entitlement to such a claim, and offer counterclaims to it. A new venue is opened up for resolving disagreements about what follows from what, about which claims rationalize which actions. Davidson's view (transposed into the deontic idiom) that in order to see how a doxastic commitment can rationalize a practical commitment we must be able to see ("at least in essential outline") how to construct a primary reason is correct in the following sense.

Once the expressive resources provided by terms such as 'prefer', 'obliged', and 'ought' are available, it must be possible to use them to make explicit the implicit practical inferential commitment underlying bits of practical reasoning such as (α), (β), and (γ). But there is nothing incoherent about an idiom that lacks those expressive resources. Practical reasoning can still take place in it, and there is still a perfectly serviceable distinction between good and bad inferences available within such an idiom. It is by comparing the more primitive practices of giving and asking for reasons for action to the sophisticated ones made possible by the introduction of inference-explicating locutions such as 'ought' that we can understand (in terms of deontic scorekeeping) the expressive role those locutions play.

The broadly normative or evaluative vocabulary that Davidson understands as expressing the pro-attitudes needed to turn the incomplete reasons offered as premises in (α), (β), and (γ) into complete reasons is actually used to make explicit in assertible, propositional form the endorsement of a *pat-*

tern of inferences. Different patterns of inference correspond to different sorts of norms or pro-attitudes. Consider someone assessing the propriety of the practical inference in (α), in a primitive idiom that does not yet permit the formulation of (a). Suppose that the scorekeeper who assesses the practical reasoning attributes to someone commitment to the premise of (α), and also entitlement to that commitment. The question is whether entitlement to the doxastic commitment serving as the premise is inferentially heritable by the practical commitment serving as the conclusion.

To take it that it is, for a particular interlocutor, just is implicitly to attribute a desire or preference for staying dry. If the inferential commitment that underwrites this piece of practical reasoning is as expressed by (a), then (α) is just one of a whole family of inferences that stand or fall together. For instance, an attributor who takes (α) to be entitlement-preserving will also take the following two inferences and a host of similar ones to have that status.

(α′) Only standing under the awning will keep me dry,
so I shall stand under the awning.

(α″) Only remaining in the car will keep me dry,
so I shall remain in the car.

To attribute a preference for staying dry to an individual is just to take inferences of this form to be entitlement-preserving, for that individual.

4. *Varieties of Prima Facie Reasons for Action*

Of course there can be competing entitlement-preserving inferences, corresponding to other desires. For recall that permissive inferences generally, whether doxastic or practical, can entitle one to incompatible conclusions[44]—though once an interlocutor endorses one of them, the undertaking of that commitment removes any entitlements that may hitherto have been available for competing claims. So endorsing this pattern of inferences as entitlement-preserving for an individual—which is implicitly attributing the preference that one could explicitly attribute either by attributing commitment to (a) or by undertaking commitment to the ascriptional claim "*A* desires to stay dry"—does not require attributing to that individual the practical commitment expressed by the conclusion in case commitment to a premise of the proper form is attributed. This is another way of saying that even in the presence of the desire, the belief need not lead to the formation of an intention, for there may be competing desires or other considerations in play. The notion of entitlement-preserving inferences accordingly provides a pragmatic analysis, in deontic scorekeeping terms, of the notion of *prima facie reasons* (whether doxastic or practical).

That a scorekeeper treats inferences of the form common to (α), (α′),

(α″) . . . to be entitlement-preserving for interlocutor *A* does not involve any commitment to treating them as entitlement-preserving for interlocutor *B*, even apart from any consideration of the attribution of incompatible commitments (doxastic or practical, or either species of inferential). Treating these inferences as permissively good for *A* but not for *B* is just what attributing the relevant preference to *A* but not to *B* consists in. This is not how endorsement of doxastic inferences (even permissive ones) works. Endorsing a doxastic inference (one whose premises and conclusions are claims, that is, expressions of possible beliefs), treating that inference as entitlement- or commitment-preserving for *one* interlocutor, involves treating it as good for *all* interlocutors—subject, as always, to disqualification by commitment to incompatible claims, and with the proviso that differences in collateral doxastic commitments can make a difference in what premises are available as auxiliary hypotheses in such inferences.

This difference in generality is a fundamental difference between doxastic inference and this sort of practical reasoning. Desire is multifarious and different from individual to individual, but truth is one; so, according to each scorekeeper, the inferences that can be described unofficially as good in the sense of truth-preserving are one, while those practical inferences that are underwritten by desires are many. Of course, beliefs may differ from individual to individual as much as desires, and with it the endorsement of inferences whose propriety is underwritten by particular doxastic commitments, although the social institution of the status of objective information by the interpersonal dimension of assertional authority and the justificatory responsibility to respond to challenges incorporates an implicit norm of common belief that has no analog for desire. The difference being pointed to here is rather that attributions of *conative* commitments are construed here as fundamentally a kind of inferential commitment, linking doxastic and practical commitments, while *cognitive* or doxastic commitments and practical commitments are inferentially articulated and inferentially significant, but not themselves inferential commitments. That (in informal terms) desires vary from individual to individual, as beliefs do, is accordingly reflected in a structure of inferential commitments in the conative case that differs from that of the cognitive case.

Permissive proprieties of practical reasoning, endorsement of which is implicitly attributing—or in the reasoner's own case, acknowledging (which is self-attributing)—desires or preferences (pro-attitudes in a strict sense, as represented by the example of (α)), are, however, only one species. Those represented by the example of (β) need not be understood as having anything in particular to do with desires or preferences. The norm, rule, or requirement that bank employees wear neckties is what makes going to work into a reason for wearing a necktie, for bank employees. Taking it that there is such a norm or requirement just is endorsing a pattern of practical reasoning—namely, taking (β) to be an entitlement-preserving inference for anyone

who is a bank employee. This inferential pattern is different from that exhibited by (α) in two ways. First, there need not be for each interlocutor for whom (β) is taken to be a good inference a set of other inferences corresponding to (α), (α'), (α''). Second, the scorekeeper will take (β) to be a good inference for any interlocutor *A* such that the scorekeeper *undertakes* doxastic commitment to the claim that *A* is a bank employee. Thus the way in which the scorekeeper's endorsement of the inference (β) as permissively good is distributed across various interlocutors who might reason that this way is different from that of (α).

Here the norm implicitly underwriting the inference is associated with having a certain status, as employee of a bank, rather than with exhibiting a certain desire or preference. Whether one has a good reason to wear a necktie just depends on whether or not one occupies the status in question. This pattern—where what matters is the scorekeeper's undertaking of a commitment to *A*'s occupying the status, rather than *A*'s acknowledgment of that commitment—corresponds to an objective sense of 'good reason for action' (according to the scorekeeper). In this sense, that *A* is preparing to go to work can be a good reason for *A* to wear a necktie, even though *A* is not in a position to appreciate it as such. The scorekeeper might take it that *A* is entitled to a practical commitment to wear a necktie, even though *A* could not justify it by producing the reasoning in (β). In wearing a necktie, *A* would be acting *with* a reason, even if not *for* a reason. This corresponds to taking a reliable noninferential reporter to be entitled by that reliability to various observations, even in the case where the reporter is not in a position to appeal to that reliability in justifying those claims.

In another sense, of course, for the norm that the scorekeeper takes to be in force to supply a reason *for A*, the claim that *A* is a bank employee must also be acknowledged by *A*. For *A* to be able to justify a commitment to wearing a necktie by rehearsing the reasoning of (β), *A* must also endorse the pattern of inference codified explicitly in (b). For a scorekeeper to take *A* to have a good practical reason in this stronger sense—that not just the scorekeeper, but *A* could produce it—requires that the scorekeeper attribute to *A* endorsement of an *inference*. In the model as presented thus far, this can be done only by attributing commitment to a claim codifying that inference. Depending on the expressive resources available, this might either be (b) or a set of corresponding conditionals.

One final stronger sense of *reason for A* is sometimes invoked by philosophers who insist that even (b) together with *A*'s acknowledgment of being a bank employee fall short of providing one unless supplemented by *A*'s desire to do what is required as a bank employee. It is indeed always possible, by supplying "suppressed" premises as needed, to assimilate all practical reasonings to the form of (α) + (a) (assimilating them to belief-desire reasons), so that norms and evaluations appear only in the role of objects of preference, as *staying dry* does in (α) + (a). And it remains true that the role of (β) would

be quite different if (b) were conjoined with some claim *incompatible* with attributing to *A* such a desire. But this is just another instance of the phenomenon mentioned above as motivating but not warranting an enthymematic view. The inference from *p* to *q* is also undercut by the *denial* of the conditional $p \rightarrow q$, but this does not show that the conditional is a suppressed premise in the original argument. The present approach requires no such instrumental reductionism, however, for it is possible to say what it is for scorekeepers to *treat* various other (from this point of view truncated) forms of practical reasoning *as* correct, as entitling agents to their practical commitments. Doing so makes it possible to explain how various normative vocabulary works (what it expresses)—including the vocabulary of commitment and entitlement that is employed in laying out the deontic scorekeeping model of discursive practice.

Endorsement of practical reasoning of the sort of which (γ) is representative, codified in the form of a normative principle by (c), corresponds to an inferential commitment exhibiting a pattern different from those involved in either (α) or (β). For a scorekeeper who takes (γ) to be entitlement-preserving for *A* takes it to be entitlement-preserving for *anyone*—regardless of desires or preferences and regardless of social status. Inferential commitments displaying this pattern are made explicit by *unconditional* 'ought's, whereas those displaying the other two patterns are made explicit by *prudential* 'ought's (in the case of (a)) and *institutional* 'ought's (in the case of (b)). Unconditional 'ought's, which correspond to this agent- and status-blind pattern of endorsement of practical inferences as entitlement-preserving, are one candidate that has been proposed as a good thing to mean by "moral 'ought'." Some thinkers insist rather that to treat reasons as moral reasons requires treating them as overriding; this amounts to saying that the 'ought' in (c) is a moral 'ought' only if (γ) is not only entitlement-preserving but also *commitment*-preserving—that anyone committed to the doxastic premises is thereby committed to the practical conclusion. It is not the point of this discussion to take a stand on how to distinguish specifically moral norms. Nor is the point to try to provide an exhaustive catalog of the sorts of norms (or "pro-attitudes," in the broad sense) there can be. The point is just to show that various important sorts of norms (or pro-attitudes) can sensibly be thought of in deontic scorekeeping terms as corresponding to different patterns of endorsement of practical *inferences*.

To endorse a practical inference as entitlement-preserving is to take the doxastic premises as providing reasons for the practical conclusion. To exhibit a piece of good practical reasoning whose conclusion is a certain intention is to exhibit that intention, and the action (if any) that it elicits, as *rational*—that is, as reasonable in the light of the facts cited and the commitments exhibited in the premises. So *all* of the 'ought's that make explicit species of practical reasoning taken as examples here (the prudential or preferential 'ought', the social or institutional 'ought', and the unconditional

'ought') are different kinds of *rational* 'ought'. Being rational is just being in the space of giving and asking for reasons, and being a rational *agent* is being in the space of giving and asking for reasons for what one *does*. When the proprieties of practical inference that articulate that space are made explicit in the form of claims, they take the form of norms—of rational 'ought's. Rationality is the genus to which all these species of 'ought's belong.

There is no a priori reason to identify the rational with some one of the species of practical reasoning (for instance the prudential) and cut and paste the rest into suitable shape to be assimilated to the favored one. Being a reason is to be understood in the first instance in terms of what it is for a community to *treat* something in practice as such a reason, on the practical side of reasons for action just as on the doxastic side of reasons for claims. In neither case is this approach to normative status (what one is *really* entitled or committed to) through normative attitude (what one is *taken* to be entitled or committed to) incompatible with making eventual sense of *objective* norms, which underwrite the possibility that everyone's attitudes toward them are wrong. But understanding what is meant by such objective proprieties—what is *really* a good reason, as opposed to just what is *treated* as one—comes at the end of the story.[45] It is not something that can be understood a priori and imposed as a constraint at the outset.

VI. INTENTIONS

1. Reasons and Entitlement to Practical Commitments

Exhibiting a piece of practical reasoning *rationalizes* the practical commitment or intention that is its conclusion. It displays reasons for that intention, offers a rational justification for it, shows how one might become rationally entitled to it. Accepting a practical inference as entitling someone to a practical commitment in this sense requires endorsing the *inference* as permissively good (and so only as providing a prima facie case for commitment to the conclusion, defeasible by incompatible commitments) for the agent whose conduct is being assessed. It does not require that the inference be accepted as one that would be (permissively) good in the scorekeeper's own case; the scorekeeper need not share the desire, preference, or institutional status that is implicitly attributed by treating some practical inferences as good for some agents. Nor does it require that in all cases the scorekeeper assessing that entitlement endorse the *premises*; a requirement of that sort picks out the special sense of objective entitlement. As Davidson says about prudential or preferential practical reasoning (the only kind he acknowledges): "When we talk of reasons in this way, we do not require that the reasons be good ones. We learn something about a man's reasons for starting a war when we learn that he did it with the intention of ending all wars [for Davidson this is equivalent to 'because he desired to end all wars'],

even if we know that his belief that starting a war would end all wars was false. Similarly, a desire to humiliate an acquaintance may be someone's reason for cutting him at a party though an observer might, in a more normative vein, think that that was no reason. The falsity of a belief, or the patent wrongness of a value or desire, does not disqualify the belief or desire from providing an explanatory reason."[46]

An agent can fulfill the justificatory responsibility involved in undertaking a practical commitment by exhibiting a piece of practical reasoning in which a commitment with that content serves as the conclusion. In the doxastic case, what is an entitling justification for one is an entitling justification for all, except for disqualifications due to commitment to claims incompatible with the premises or the conclusion. In the practical case, entitling justifications need not be portable across agents in this way. Displaying an intention as the conclusion of a piece of practical reasoning that is good in this sense makes it intelligible by showing reasons that could entitle the agent to it. To secure an attribution of entitlement to a practical commitment in this way, the practical reasoning in question may be offered as a justification by the agent, perhaps upon being challenged to do so. But it also may be attributed by the scorekeeper, who constructs the practical argument from premises already attributed to that agent, according to patterns of practical inference the scorekeeper endorses for that agent. These might be patterns of inference the scorekeeper endorses for everyone (unconditional 'ought's), or endorses for the agent on the basis of status (institutional 'ought's), or endorses only for the agent, thereby implicitly attributing idiosyncratic desires or preferences (prudential 'ought's), or of some other kind.

A scorekeeper who in this way takes an agent to be entitled to a practical commitment on the basis of its being the conclusion of a practical inference taken to be good for the agent, and who also attributes to that agent commitment to the premises of that inference, need not take it that acknowledgment of the practical commitment actually arose as the result of a process of inference by the agent from acknowledgment of those premises. As Davidson says "We cannot suppose that whenever an agent acts intentionally, he goes through a process of deliberation or reasoning, marshals evidence, and draws conclusions. Nevertheless, if someone acts with an intention, he must have attitudes and beliefs from which, had he been aware of them at the time, he *could* have reasoned that his action was desirable (or had some other positive attribute)."[47] For Davidson, acting intentionally and acting for reasons are the same thing. From the present point of view, this position involves conflating the two deontic statuses of practical commitment and entitlement to such a commitment. An act is intentional if it is (or is, as the exercise of a reliable differential responsive disposition, noninferentially elicited by) the acknowledgment of a practical commitment. To act for reasons is to be *entitled* to that practical commitment. One can in particular cases act intentionally but without reasons, even though there is no making sense of inten-

tions apart from their liability to the demand for reasons. But Davidson's point survives this confusion. An intention can be rendered intelligible as rational (a practical commitment can be displayed as one the agent is entitled to) by displaying reasons for it (premises from which it *could* legitimately have resulted as the conclusion of a good practical inference), even in cases where it was not in fact arrived at by such a process.

Contemporary thought about action begins with Anscombe's insight, developed with great force and clarity by Davidson, that the difference between actions and other performances—the answer to Wittgenstein's challenge to explain the difference between my raising my arm and my arm going up—is that actions are performances that are intentional under some description.[48] Any performance can be specified in many ways: Davidson moves his finger, flips the switch, turns on the light, alerts the burglar, causes a short-circuit in the wiring, starts a fire. These are all things he *does;* the different descriptions are different ways of specifying one action he performs. Not all of these are specifications under which what he does is intentional. But they are all specifications of an action, so long as what he does is intentional under *some* specification, for instance turning on the lights. In Davidson's slogan, being an action is an *extensional* property of an event (whether a given event is an action or not is not sensitive to how the event is specified), while being intentional is an *intensional* property of an event (whether a given event is intentional or not is sensitive to how the event is described). The very same event is intentional as turning on the lights but unintentional as alerting the burglar, causing a short-circuit, and starting a fire. The extensional property of performances, *being an action,* is defined in terms of the intensional property of performances of *being intentional* by existential quantification over descriptions or specifications of the performance: if it is intentional under any one of them, it is an action under all of them.

By this strategy the problem of explaining what privileges some (but not all) of an agent's performances as actions is reduced to the problem of explaining what privileges some (but not all) descriptions or specifications of an action as ones under which a performance is intentional. Davidson's solution to this problem in turn is that a performance is intentional under a description if that description figures as the conclusion of a piece of practical reasoning that exhibits the agent's reasons for producing it. These two moves together—the account of actions as performances that are intentional under some description, and the account of performances as intentional under just the descriptions that appear as the conclusions of practical inferences that rationalize those performances by giving reasons for them—reduce the problem of explaining what is special about action to that of explaining the giving of practical reasons. Davidson's account of primary reasons as pairs of beliefs and pro-attitudes is then offered to explain what it is for reasons to rationalize a performance according to a practical inference.

The Davidsonian explanatory structure provides a recipe, then, for turning

an account of practical reasoning into an account of action. In such a context, the present account of practical reasoning in terms of deontic scorekeeping on inferentially articulated practical commitments and entitlements to such commitments has some advantages over the one Davidson himself endorses. As originally presented, Davidson's theory eschews intentions entirely, in favor of beliefs and desires: "The expression 'the intention with which James went to church' has the outward form of a description, but in fact it is syncategorematic and cannot be taken to refer to an entity, state, disposition, or event. Its function in this context is to generate new descriptions of actions in terms of reasons; thus 'James went to church with the intention of pleasing his mother' yields a new, and fuller, description of the action described in 'James went to church'."[49] The account offered here, by contrast, explains deontic statuses corresponding to beliefs and intentions and defines those corresponding to desires (and other 'expressions of pro-attitudes') in terms of them.

Acknowledging intendings as full-fledged intentional states (or attitudes toward deontic statuses) avoids at the outset a difficulty that forced Davidson to modify his earlier account. For there are cases where someone has an intention (not just a reason for action), but no action arises from it. The possibility of intention without action is a symptom of the limited scope of Davidson's original discussion. As he puts the point in his introduction to the collection of his essays on this topic: "When I wrote ['Actions, Reasons, and Causes'] I believed that of the three main uses of the concept of intention distinguished by Anscombe (acting with an intention, acting intentionally, and intending to act), the first was most basic. Acting intentionally, I argued . . . was just acting with some intention. That left intending, which I somehow thought would be simple to understand in terms of the others. I was wrong . . . Contrary to my original view, it came to seem the basic notion on which the others depend; and what progress I made on it partially undermined an important theme of ['Actions, Reasons, and Causes']—that 'the intention with which the action was done' does not refer to an entity or state of any kind."[50]

2. *Two Sorts of Intention*

Explaining intentional action requires only what Searle calls *intentions in action*. Explaining pure intending requires also what he calls *prior* intentions. The distinction is motivated by the fact that "I can do something intentionally without having formed a prior intention to do it, and I can have a prior intention to do something and yet not act on that intention."[51] Pure intendings are special cases of prior intentions. In the deontic idiom, both sorts of intentions are (acknowledgments of) practical commitments—that is, commitments to act. Cases of intentions in action without prior intentions are those in which the performance that is accorded by a scorekeeper

the significance of an acknowledgment of a practical commitment is just the action itself. Consider Searle's example: "Suppose I am sitting in a chair reflecting on a philosophical problem, and I suddenly get up and start pacing about the room. My getting up and pacing about are clearly intentional actions, but in order to do them I do not need to form an intention to do them prior to doing them. I don't in any sense have to have a plan to get up and pace about. Like many of the things one does, I just do these actions; I just act."

In scorekeeping terms, undertaking a commitment is doing anything that makes it appropriate for it to be attributed. This may involve a distinct explicit acknowledgment, as in the case of asserting; it may be consequential, as in undertaking commitment to the consequences of a claim that is asserted; or it may be a default matter. Scorekeepers will take one not only to be entitled but to be committed to the claims that orange is a color and that there have been black dogs, unless these have been overtly disavowed (and even this may not be sufficient, if collateral commitments that are not disavowed entail these claims). The sort of case Searle considers concerns actions that have the default significance of intentional actions, as ones accompanied by a commitment for which the question of entitlement by reasons is in principle in order. It might turn out that what is, when described in a suitably impoverished vocabulary of motions of limbs, exactly similar behavior is not intentional at all but automatic, involuntary, compulsive behavior, triggered ineluctably by pheromones. In that case the default attribution of a practical commitment would be defeated. But the undertaking or acknowledging of a commitment to act need not be a performance distinguishable from the act one is committed to perform.

Nonetheless, in many cases, it is a separately datable event. In such cases of prior intention, acknowledging the commitment antedates fulfilling it (or not, as the case may be). I can *now* acknowledge a commitment to get on the bus *when* it arrives. My mastery of the practical conceptual (because inferentially articulated) content of that commitment includes my mastery of the *noninferential* differential disposition to respond to it and the joint fulfillment of its condition (the bus arriving) by getting on the bus. When the bus arrives, the condition is fulfilled. My getting on the bus is an acknowledgment of a practical commitment to get on the bus *now*—an intention in action. If I have the reliable noninferential differential disposition to respond to the acknowledgment of a practical commitment to do *A* when (if) *C* by doing *A* when (if) *C*, then my prior intention to get on the bus when it arrives will *mature* into a corresponding intention in action (marked above by the 'now' in the linguistic expression of the intention).

Prior intentions must specify the actions one is committing oneself to perform in *general* terms; they would be expressed by statements of the form

I shall get on the bus when it arrives,

in which the action is specified by a description. Intentions in action can be thought of as specifying the action one is committing oneself to perform in *demonstrative* terms; they could be expressed by statements of the form

I shall *now* do *this*.

Only an intention in action can be directed at a particular unrepeatable action. The process by which a prior intention ripens into an intention in action—the exercise of the practical skill of reliably responding to the undertaking of a commitment to bring about a state of affairs of such and such a description when such and such conditions obtain by doing so—involves recognizing when doing *this now* would fulfill such a commitment. (See further discussion below in 8.5.2.) What are here called "prior intentions," Sellars calls simply "intentions," and what are here called "intentions in action," Sellars calls "volitions":

> A simple case of the relation of intending to volition can be illustrated by considering Jones, who has formed the intention of raising his hand in ten minutes. Suppose that no alternative course of action recommends itself to him. Then we may picture the situation as follows:
>
> I shall raise my hand in *ten* minutes.
>
> . . .
>
> I shall raise my hand in *nine* minutes.
>
> . . .
>
> I shall raise my hand *now*.
>
> (which culminates in action, if Jones happens not to be paralyzed).[52]

In Sellars's idiom, a volition is an intention whose time has come.

Sellars takes the capacity respond reliably to prior intentions whose time has come by the formation of intentions in action (here, acknowledgments in the form of suitable performances) to be part of grasping the meaning of what is expressed by 'shall', and of the practical content of the particular concepts that articulate the content of those intentions. As in the account endorsed here, he takes these capacities to be part of the "important similarity between learning to make the language-entry transition of responding to presented red objects by saying 'This is red,' and learning the language-departure transition" involved in exercising those capacities.[53] Just as in the case of language entries or noninferentially elicited but inferentially articulated doxastic commitments, the existence of reliable differential responsive dispositions is compatible with making mistakes, so in the case of language exits or performances noninferentially elicited by inferentially articulated practical commitments, the existence of reliable differential responsive dispositions is compatible with failure. Mistakes of observation are diagnosed by scorekeepers by comparison of the contents of the doxastic commitments *attributed* to the observer and those *undertaken* (whether noninferentially,

inferentially, or by testimony) by the scorekeeper—for instance when the scorekeeper attributes a commitment to the claim that the ball is red and undertakes a commitment incompatible with the claim that the ball is red. Failures of action are diagnosed by scorekeepers by comparison of the contents of the practical commitments *attributed* to the agent and the contents of the doxastic commitments *undertaken* by the scorekeeper.

For instance the scorekeeper who attributes a commitment that could be expressed (Sellars-wise) as

Shall [The basketball goes through the hoop]

may also be obliged (perhaps observationally) to acknowledge a separate commitment incompatible with the claim that the basketball goes through the hoop (for example that expressed by "The ball missed the hoop"). A practical commitment may also remain a "pure" intending, eventuating in no action, successful or unsuccessful, for one of two reasons. It might be that the condition of maturation of a prior intention into an intention in action never is satisfied—the bus never arrives, the commitment expires. Or it may be that the commitment is relinquished, perhaps in favor of one with an incompatible content—the agent undergoes a change of mind.

3. *Acknowledgments of Commitments Can Cause and Be Caused*

When prior intentions are made explicit, they specify in general descriptive terms the performance the agent is committed to produce. When intentions in action, which are implicit in suitable performances, are made explicit, they specify those performances demonstratively. The way in which prior intentions elicit suitable performances, and so intentions in action, is a causal process. The analogy between action and observation (between language-exit transitions and language-entry transitions) is intended to illuminate the nature of the process involved.

Davidson's original essay famously endorses the claim that reasons are causes. The difference between a commitment's being *a* reason for an agent's action and its being *the* reason for that action must be explained in terms of differences in the causal roles played by various states. Primary reasons, conceived of as pairs of a set of beliefs and a set of pro-attitudes, rationalize actions (which accordingly count as intentional) first by providing reasons for them and second by serving to bring them about ("in the right way"). The account offered there has been criticized here for running together the notion of being *committed* to act in a certain way and being *entitled* to do so by reasons; even though the first deontic status cannot be made sense of apart from the second (any more than doxastic commitments can be made sense of apart from practices of giving and asking for reasons entitling interlocutors

to them), nonetheless these are distinct statuses, and one can be committed without being entitled, can act intentionally without having reasons for doing so. This distinction between practical commitments and entitlements to such commitments also, it has been claimed, opens up a space for the notion of *prior* intention (besides that of intention in action), of which pure or unconsummated intendings are a species.

What becomes of the doctrine that reasons are causes, when intentional states are construed in terms of social scorekeeping on deontic statuses and the deontic statuses of commitment and entitlement are appropriately distinguished? The account of the social practices that institute deontic statuses appeals to such statuses only as scorekeeping devices. The significance of being committed to a certain claim or assertible content is normative. It has to do with what else one is committed or entitled to. It is articulated by *proprieties* of scorekeeping and consists of the *proper* antecedents and consequences of that status. In the same way, that there are two strikes on a batter is a status properly acquired by various performances (just which depending on the antecedent score), a status that alters the significance various further performances have for the subsequent score.

Any effect that such elements of the score have on what performances are *actually* produced is indirect, mediated by the *attitudes* of those who keep score. The score determines only what *ought* to be done, what would be *proper*. What ought to be done and what is proper affect what players do only insofar as they are trained to respond in various ways to *taking* a certain course of action to be proper. The only access that deontic statuses have to the causal order is through the deontic attitudes of the scorekeeping practitioners.

Inferential relations among propositional contents are a matter of normative relations among deontic statuses: commitment to the claim that lions are mammals entails commitment to the claim that lions are vertebrates. Inferring, by contrast, is a causal process that relates deontic *attitudes: acknowledging* (and equally, attributing to someone else) commitment to the claim that lions are mammals will, under various circumstances and in those well versed practically in the inferential relations among deontic statuses, have as a causal consequence acknowledging (or, correspondingly, attributing) commitment to the claim that lions are vertebrates. Unless the members of a linguistic community are pretty good at keeping score by altering their attitudes as they ought to according to the contents associated with the deontic statuses in terms of which they keep score, there is no point in interpreting them as engaging in the practices specified by those proprieties of scorekeeping. Nonetheless, normative status is one thing, the attitudes of attributing and undertaking those statuses, the alteration of which is what scorekeeping consists in, is another.

As it is with the inferential articulation of the conceptual contents conferred on states, attitudes, performances, and expressions by deontic score-

keeping practices, so it is with their noninferential involvements, which confer empirical and practical conceptual contents on them. What observable states of affairs causally elicit in perception, according to reliable differential responsive dispositions, is in the first instance deontic *attitudes* rather than statuses: *acknowledgments* of doxastic commitments. What in action causally elicits the production of performable states of affairs (by the exercise of reliable differential responsive dispositions) is in the first instance deontic *attitudes* rather than statuses: *acknowledgments* of practical commitments. ("In the first instance" because acknowledging a commitment is one way of undertaking one, so those deontic attitudes have scorekeeping consequences for the deontic statuses of those whose attitudes they are.) That a particular doxastic commitment was elicited by the exercise of such dispositions is another way of putting the condition on perception that the belief not only be caused by the state of affairs reported but be caused by it "in the right way." That a particular practical commitment elicits a performance by the exercise of such dispositions is another way of putting the condition on action that the performance not only be caused by the intention but be caused by it "in the right way." Mastering the two sorts of reliable differential responsive dispositions connecting noninferentially acquired acknowledgments of doxastic commitments to their appropriate causal antecedents and noninferentially efficacious acknowledgments of practical commitments to their appropriate causal consequents is part of grasping, in one's scorekeeping practice, the empirical and practical components of the contents of concepts employed in observation and action, and of those theoretical concepts inferentially related to these.

Thus just as 'belief' is ambiguous in scorekeeping terms, referring sometimes to a deontic status and sometimes to a deontic attitude (sometimes to doxastic commitment and sometimes to acknowledgment of such a commitment), so 'intention' is ambiguous in scorekeeping terms, referring sometimes to a deontic status and sometimes to a deontic attitude (sometimes to practical commitment and sometimes to acknowledgment of such a commitment). Believing in the sense that entails one's readiness to avow what one believes and to act on it corresponds to acknowledging a doxastic commitment. Intending in the sense that entails one's readiness to act on it (and, should the expressive resources for doing so exist in the linguistic practices in question, to avow it with a 'shall' claim) corresponds to acknowledging a practical commitment. In this sense of 'intention', then, intentions are causes, for in the properly trained agent, acknowledgments of practical commitments reliably causally elicit performances. In this sense of 'belief', when beliefs provide reasons that entitle one to a practical commitment, they *may* function also as causes. They do just in case the acknowledgment of the practical commitment in fact arose by inferring it from an acknowledgment of the belief playing the role of premise in practical reasoning.

One, however, may have intentions without reasons, practical commit-

ments to which the agent is not entitled by doxastic commitments suitably related as premises of practical inferences. In that case there can still be action, but it will not be caused by reasons. Again, the agent may be entitled to the practical commitment, according to a scorekeeper, by doxastic commitments that are attributed by the scorekeeper but not acknowledged by the agent. This might happen where the scorekeeper takes the agent to have undertaken corresponding doxastic commitments as inferential consequences of others that are acknowledged, but where the agent has never been through the process of inference that would lead to acknowledging those consequences. Commitments of this sort could still entitle the agent to the practical commitment, even though only the scorekeeper, and not the agent, would be in a position to exhibit the practical reasoning that secures that entitlement. In such cases, too, the agent's reasons for the action would not be functioning as causes. So once the deontic statuses of commitment and entitlement to commitments are properly sorted out, it turns out that a nonlinguistic performance can have at least three different sorts of scorekeeping significance:

1. an agent's acting intentionally—that is, acknowledging a practical commitment by producing a performance or exercising a reliable noninferential differential disposition to respond to acknowledgments of practical commitments by producing a performance,
2. an agent's having reasons for action or acting *with* reasons—that is being entitled to a practical commitment, and
3. an agent's acting *for* reasons, the action being caused by (attitudes toward) reasons for action—that is the acknowledgment of the practical commitment having arisen by a process of inference from acknowledgment of the commitments that provide the entitling reasons.

The first does not entail the second (nor vice versa), nor does the second entail the third, though they are all compatible; one can act intentionally either with or without reasons, and one may or may not act for the reasons one has.

4. Acknowledging Commitments Need Not Be Modeled on Promising

Davidson's own view about intentions (once he comes to countenance them at all) identifies them as all-things-considered judgments, in the light of all the agent's primary-reason-providing beliefs and desires, that an action of a certain kind is desirable, good, or ought to be performed.[54] From the present point of view this is an unsatisfactory conclusion, both because of its appeal to unanalyzed notions of desirability, good, and what ought to be done, and because it does not say what it is for the attitudes these locutions express to become explicit in the form of an evaluative judgment.

He has not laid out the practices that could confer such a content on a judgment, so he has not explained how to understand fully the implicit commitment that is being made propositionally explicit by the use of this vocabulary. Such a complaint is of exactly as much interest as the concrete alternative account that is recommended; it is better to light a candle than to curse the darkness. The deontic scorekeeping account of practical commitments, built on that of doxastic commitments, is meant to supply this want. In order to recommend that idiom over the one Davidson endorses, however, it is necessary to confront the argument against construing intentions as a sort of commitment that he offers along the way to his identification of intentions as a special kind of judgment.

The leading idea of the present account is that acting intentionally is doing something that has the deontic scorekeeping significance of acknowledging a practical commitment (in the case of intentions in action), or noninferentially producing a performance by exercising a reliable differential disposition to respond to the acknowledgment of a practical commitment (where prior intentions are involved). Intentions are identified with such acknowledgments of commitments, and the reasons for or with which an agent acts with the attitudes or facts that entitle that agent to those practical commitments, according to the role they play as premises in practical inferences. This normative (more specifically deontic) approach to intention and action is rooted in Sellars's discussion of the giving and asking for reasons for action, which has been elaborated along different lines by Castañeda.[55] Although the details of their accounts are different, the overall approach is very similar. Sellars never actually talks about intentions in terms of commitments, but this way of putting it is implicit in his account.[56] If there is something wrong with thinking about intentions in terms of commitments, then this whole approach is broken-backed. So it is of the first importance to consider Davidson's arguments against it.

Davidson begins by considering theories that focus on the speech act of expressing an intention (the speech act that Sellars regiments using 'shall'). He observes that "saying, under appropriate circumstances, that one intends to do something, or that one will do it, can commit one to doing it; if the deed does not follow, it is appropriate to ask for an explanation."[57] The suggestion that forming an intention is performing a speech act of this sort (perhaps addressed to oneself)—a performative theory of intention—is rejected because "the performative character of commands and promises which makes certain speech acts surprisingly momentous depends on highly specific conventions, and there are no such conventions governing the formation of intentions."[58] Indeed it seems enough to observe that, although for Davidson as for the deontic scorekeeping account, one must be able to talk in order to have intentions (because it is only in the context of linguistic practices of giving and asking for reasons that anything could be accorded the significance of an intention), there is no necessity that there actually be a term 'shall' that

overtly marks a special speech act that attaches to propositional contents the significance of explicit undertakings of practical commitments. The present explanatory strategy demands rather that that implicit force first be explained in scorekeeping terms so that a clear sense can be made of the introduction of locutions whose expressive function is to make that force explicit in the form of an assertion. The present account does not understand the undertaking of a practical commitment as requiring a special speech act. What is important is the attitude of acknowledging a commitment to act; any connection with special performative speech acts comes later.

But Davidson objects as well to the invocation of commitments in this context: "Promising involves assuming an obligation, but even if there are obligations to oneself, intending does not normally create one. If an agent does not do what he intended to do, he does not normally owe himself an explanation or apology, especially if he just changed his mind; yet this is just the case that calls for explanation or apology when a promise has been broken. A command may be disobeyed, but only while it is in force. But if he does not do what he intended because he has changed his mind, the original intention is no longer in force."[59] There are a number of points being made here; they turn on disanalogies between forming an intention and making a promise, which serves for Davidson as the paradigm of the undertaking of a commitment. To begin with, he offers the implicit suggestion that there may be problems with the notion of making a promise to oneself. Promises (like commands) are made *to* someone, while no one else is typically addressed by the formation of an intention. Then the central objection is presented, that there seems to be no *sanction* associated with failure to perform as one is committed to perform. If there is a commitment, then fulfilling it or failing to fulfill it ought to make some sort of difference, as it does in the case of failing to fulfill a promise or to carry out an order from a suitable authority. Yet once an intention has been formed, it can be withdrawn without penalty—the agent can have a change of mind. Promises would not be promises, would not involve the undertaking of *commitments*, if they could be canceled at the whim of the promiser. How could sense be made of a commitment that was in force only as long as the one committed decided to keep it in force but that could be relinquished without penalty at any time? Davidson concludes that the disanalogies are too great and that forming an intention cannot sensibly be conceived as undertaking a commitment.

These disanalogies between intending and promising, even promising oneself, should be acknowledged. But the conclusion follows only if there is no other model of acknowledging or undertaking commitments available besides that of promising. The deontic scorekeeping account of acknowledging *assertional* or doxastic commitments shows that this is far from being the case. Assertional commitments, after all, can be withdrawn without penalty by the asserter who undergoes a change of mind. Commitments of this sort

are put in force by the performance of a speech act and, except for special cases (where one is also consequentially committed to the claim in virtue of other, unwithdrawn assertions), can be relinquished without penalty by another speech act. It is true that nothing resembling promising could work this way, but the model of asserting shows that there are other ways to conceive the undertaking and acknowledging of commitments. While a doxastic commitment is in force, that fact has consequences; the undertaking of such a commitment has a significance for the deontic score. Commitment to one content entails commitment to others and precludes entitlement to yet others.

It is the same with practical commitments as here presented. Undertaking one is not without significance simply because it can be voided, withdrawn, or overridden essentially at the whim of the agent. For when such a commitment is in force (according to a scorekeeper who attributes it), it is significant. It entails various further commitments and precludes various entitlements. It can license the attribution of doxastic commitments (standing in for beliefs) that would warrant it, according to an attributed piece of practical reasoning. Like doxastic commitments, practical commitments involve a (conditional) justificatory responsibility to vindicate the commitment by demonstrating entitlement to it (upon suitable challenge). This forms part of the significance of these commitments, on the side of antecedents rather than of consequents, for it determines the circumstances under which it is appropriate to acquire these statuses. The disanalogies between promising and undertaking doxastic commitments do not make the latter sort of status unintelligible as a species of commitment, and the same disanalogies between promising and the undertaking of practical commitments, which are modeled closely on doxastic ones, do not make that sort of status unintelligible as a species of commitment.

It might be objected that the disanalogies between doxastic and practical commitment reinstate the difficulty. For on the one hand, doxastic commitments are like those undertaken by promising, and unlike intendings, in that they are intelligible only in terms of a speech act that has the significance of an overt public acknowledgment of them. And on the other hand, a scorekeeping sanction for failing to fulfill the justificatory responsibility associated with undertaking a doxastic commitment is the loss (in the eyes of the scorekeeper who attributes the failure of entitlement) of its authority, its capacity to license commitment by others to that same content. But the lack of this sanction is precisely one of the important points of disanalogy between doxastic and practical commitments.

Each of these points might have force if practical commitments were conceived as autonomous—that is, as statuses that could be instituted by practices that did not also institute doxastic deontic statuses. This sort of autonomy is claimed only for doxastic commitments and entitlements, however, not for their practical counterparts. In this context, the disanalogies

between doxastic and practical deontic statuses do not reinstate Davidson's objections. The first point just shows that because of the essential role played by the overt public acknowledgment of doxastic commitments by the performance of speech acts accorded the significance of assertions, doxastic commitments are more like those undertaken by promising than are practical commitments. For practical commitments need stand in no such intimate relation to speech acts accorded the significance of acknowledgments of them. Explicitating locutions permitting the production of speech acts of this sort can be introduced, but practical commitments are intelligible even in their absence. Practical commitments as here conceived are unintelligible apart from all reference to the overt undertaking of commitments by speech acts; that is why they are an essentially linguistic phenomenon. But as here described, the only sort of speech act they presuppose is assertion, the acknowledgment not of practical but of doxastic commitments.

The second point was that intentions do not implicitly claim the sort of interpersonal authority that assertions do, so that the consequences of attributing commitment without entitlement cannot be in the practical case what they are in the doxastic case, namely the undercutting of that authority. But this does not show that entitlement to practical commitments is unintelligible, only that it is different in detail from entitlement to doxastic commitments. Entitlement to practical commitments still has an intrapersonal significance, for instance in connection with the incompatibility of practical commitments (which is linked to that of doxastic commitments), and so with their permissive entailments. Again, this feature does not threaten to make practical commitments unintelligibly private, both because of the irreducibly social character of the deontic scorekeeping, in terms of which such statuses are explained by theory and sustained by practice, and because of the connection with doxastic commitments via practical reasoning.

Entitlements aside, treating a performance as (or as elicited by) an acknowledgment of a practical commitment—that is, treating it as intentional—has scorekeeping consequences not only for the attribution of further practical commitments but also for the attribution of doxastic commitments. It is often possible to infer what an agent believes from what that agent does. Committing oneself to act in a certain way may be committing oneself to taking the world to be a certain way, in the eyes of a scorekeeper who attributes a suitable background of other commitments (of both discursive species). The doxastic commitments an agent is taken consequentially to have undertaken in this way may be incompatible with other doxastic commitments the scorekeeper attributes, in which case entitlement to all of them is undercut. So in part in virtue of the intimate connections between them, the asymmetries between practical and doxastic commitments do not threaten the intelligibility of the deontic scorekeeping significance of the former.

5. *'Should' and 'Shall'*

Once these points are appreciated, it becomes clear that Davidson's considerations provide no reason not to understand forming an intention as acknowledging a commitment, provided commitment is properly understood according to the practical deontic scorekeeping model of doxastic commitment rather than on the model of promising. But it also becomes clear that there is surprisingly little difference between his construal of intentions as all-out evaluative judgments and the deontic scorekeeping construal of them as acknowledgments of practical commitments. For, first, though Davidson does not think of them that way, in the context of the model presented here, taking intentions to be a kind of judgment *is* taking them to be a kind of commitment. Davidson uses 'judgment' as the genus and allows cognitive and conative, or descriptive and evaluative, species corresponding to beliefs and intentions. The idiom presented here uses 'commitment' as the genus and allows doxastic and practical species corresponding to beliefs and intentions.

For Davidson, there are two sorts of "evaluative" judgment: those that express pro-attitudes suited to be elements of primary reasons for action (that is, those that express merely prima facie or ceteris paribus evaluations), and those that express intentions and are directly responded to by the production of suitable performances. In the case of practical reasoning whose conclusion is the formation of an intention, the agent is noninferentially disposed to respond reliably by producing suitable performances—which are qualified as actions by having such a provenance; these two sorts of evaluations appear in Davidson's account in the role of *premises* and of *conclusions*, respectively. The account offered here denies that what is expressed by the prima facie evaluative judgments Davidson understands as codifying pro-attitudes need appear as explicit premises in such reasoning; they correspond to the endorsement of a pattern of practical inference as entitlement-preserving. Such practical inferential commitments may be made explicit in the form of doxastic commitments with assertible contents (and so be available for duty as explicit premises) if suitable explicitating vocabulary is available—just as theoretical inferential commitments may, but in general need not, be codified explicitly by the use of conditionals. In this use, then, normative expressions exemplified by 'should' as it appears in rules of conduct play an inference-explicitating role on the practical side that is analogous to that played by 'if . . . then . . .' on the doxastic side;[60] in neither case does the omission of a premise that codifies a material propriety of inference, whether practical or doxastic, result in an enthymeme. Besides this *permissive* use of normative locutions such as 'should', which corresponds to Davidson's prima facie evaluative judgments, there is also a *committive* use, which corresponds to the "all-out" evaluative judgments that serve for him as intentions.

Recall the discussion of intentional explanation in Chapter 1. One of the ideas advanced there in connection with the suggestion that intentional states be understood in terms of deontic statuses and (propositional) attitudes toward them is that the conclusions of intentional explanations in the strict sense are *normative*, rather than *descriptive* claims. One attributes beliefs and desires (or other evaluations or pro-attitudes) and concludes from those attributions, not that the agent *will* perform an action of a certain kind, but that the agent is *committed* by those beliefs and desires to do so, that in the light of those other attitudes the agent *ought* (rationally) to do so. Intentional explanation illuminates what was done by showing why the intentional agent was *committed* to acting in that way. Under various circumstances it is possible to continue the inquiry and to ask why the agent acted in accord with that commitment. The response to such a question is not an intentional explanation, however, but a different sort of account—one showing why it is useful to offer intentional explanations of this individual, why treating the individual as a rational agent is a useful predictive and explanatory strategy. Explanations of this supplementary sort may appeal to how the organism is wired up and how it was trained so as to be able to respond reliably to the acknowledgment of a practical commitment by producing a performance of the sort specified in the content of that commitment. Such considerations are offstage from the point of view of intentional explanations proper, for these go only as far as showing what an agent should (rationally) do, what the agent is committed to do by the doxastic and inferential commitments that agent acknowledges.

Intentional explanations display sample pieces of practical reasoning, attributing theoretical and practical deontic statuses as premises and attributing a practical commitment as a conclusion. To serve as an intentional explanation of something the agent did, or to draw a conclusion about what the agent should do on which a prediction might be based about what the agent will do, these must be treated as commitment-preserving inferences. For the conclusion is that the intentional agent was or is *committed* to act in a certain manner. When the deontic scorekeeper attributes various commitments and concludes that therefore the agent *should* perform an action satisfying a particular description, the evaluative judgment expressed is of Davidson's second, all-things-considered kind.

So on the deontic scorekeeping approach there are two sorts of 'should', corresponding to the two sorts of evaluative judgment that Davidson considers: one involving prima facie evaluations suitable to serve as premises in practical reasoning, and one involving all-in evaluations suitable to serve as conclusions in practical reasoning. The first sort of 'should' is used to make explicit the endorsement, undertaken or attributed, of a pattern of practical reasoning, as in "Bank employees should wear neckties." The second sort of 'should' is used to make explicit commitments to act, which are attributed as the conclusions of committive practical inferences attributed in the course

of intentional explanation. To say this is to say that in the latter sort of use, 'should' expresses in the third person what is expressed in the first person by 'shall'. (Indeed, 'shall' and 'should' are etymologically linked in just the way suggested by this doctrine.)

Translated into the official scorekeeping terminology of deontic attitudes, this is the claim that while 'shall' is used to make explicit the *acknowledgment* (and therefore the *undertaking*) of a practical commitment to make some claim true, 'should' is used to make explicit the *attribution* of such practical commitments. The same piece of practical reasoning can be presented from either social perspective. In first-person, deliberative terms, the agent may acquire a practical commitment that would be made explicit (if the idiom encompasses sufficient expressive resources) by an overt utterance of "I shall wear a necktie," as the result of an inference from acknowledged commitments that would be made explicit (perhaps in response to a challenge to demonstrate entitlement to the practical conclusion) by an overt utterance of "I am a bank employee." In third-person, scorekeeping terms, the scorekeeper may attribute a practical commitment, adopting an attitude that would be made explicit (if the idiom encompasses sufficient expressive resources) by an overt utterance of "He should wear a necktie," as a result of an inference from attributed commitments that would be made explicit by an overt utterance of "He is a bank employee." The same piece of practical reasoning can be exhibited either by the one undertaking a practical commitment or by the one attributing it—significant either in deliberation regarding action or in assessment of such action. Indeed, deliberation—my considering various practical inferences in order to decide what I *shall* do—is just the internalization of assessment, the consideration of what anyone, given the relevant collateral commitments and circumstances (as they are taken by the assessor to be), *should* do.

6. *Weakness of the Will*

Of course 'should' has first-person uses as well. Some of these express only endorsement of patterns of permissive inference and so correspond to Davidson's merely prima facie evaluative judgments: "I should wear a necktie" (since I am a bank employee, but only if there is no better reason not to do so). But some also are self-attributions of commitments, in which one takes up a third-person perspective toward oneself, drawing conclusions about what one's reasons commit one to do: "I should (all things considered) drive to the airport." With practical commitments, as with doxastic ones, although acknowledging a commitment entails attributing it to oneself, the converse is not the case; though attributing a commitment to oneself is one way to *undertake* that commitment, this can be a consequential undertaking, rather than an acknowledgment. In particular, a self-attribution of the 'I should . . .' variety need not trigger the reliable differential dispositions to

respond to an acknowledgment of a practical commitment by producing a suitable performance. The noninferential significance of the deontic attitude that is made explicit by 'I should . . .' can be different from that of the deontic attitude that is made explicit by 'I shall . . .' Thus one can self-attribute a practical commitment without acknowledging it in the sense that matters for eliciting action.[61]

This possibility is one of the phenomena philosophers have discussed under the heading of weakness of the will, or *akrasia*—knowing the better and doing the worse. The attitude expressed by 'I should . . .' in its all-in sense does indeed commitment-entail that expressed by the corresponding 'I shall . . .' statement. But the difference between acknowledging a commitment and consequentially undertaking it depends on the fact that one does not always acknowledge the consequences of commitments that one acknowledges. In scorekeeping terms one can nonetheless be said to undertake those consequential commitments because the initial acknowledgment licenses others to attribute them. This distinction remains even when one takes up a third-person point of view toward oneself, as in deliberation about various possible courses of action when the agent traces out what commitments would be undertaken consequentially were certain others acknowledged. For in deliberating, an agent considers what commitments would be attributed by scorekeepers, under various circumstances. It is for this reason that one must be able to assess the conduct of others in order to deliberate about one's own.

The akratic's deliberations and intentions are out of step; the commitments acknowledged in the third-person theoretical way are incompatible with those acknowledged in the first-person practical way. The akratic agent is accordingly the analog on the practical side of the believer who undertakes incompatible doxastic commitments. It is one of the cardinal strengths of the deontic scorekeeping approach to intentional states in terms of normative statuses that there is nothing conceptually mysterious about the possibility of such incompatible commitments. Difficulties in coherently understanding akratic action and endorsement of incompatible beliefs arise from exclusive emphasis on a causal-functional model of intentional states.

The account of action presented here is a thoroughly Kantian one. For Kant, *will* is just the capacity for practical reasoning—that is, the capacity to derive performances from a conception of laws.[62] In the terminology introduced here, this is just the capacity to respond reliably to acknowledgments of commitments (the pragmatic version of "deriving from conceptions of laws") by producing suitable performances—suitable in terms of the way their descriptions line up with the contents of the practical commitments they either acknowledge or by the acknowledgment of which they are responsively elicited. For Kant the expressive role of 'ought' or 'should' *(Sollen)* is to make norms explicit in the form of imperatives. Specifically, such normative vocabulary "indicates the relation of an objective law of reason to a will

which is not in its subjective constitution necessarily determined by the law."[63] As construed here, normative vocabulary (of which 'ought' and 'should' are paradigmatic) has the logical expressive function of making explicit in the form of something that can be *said* (put in the form of a claim) an attitude that otherwise could be implicit only in what is *done*—namely, the endorsement of a pattern of practical reasoning. The propriety of a form of reasoning is the practical correlate of Kant's "objective (= valid, binding) law of reason," and its acknowledgment as constraining actual practical reasoning is its "subjective" relation to the will.

The rational will as described here is not a particularly puzzling phenomenon. Its normative dimension is explained by extending the account of discursive commitments to encompass not only doxastic but practical deontic statuses. Its causal dimension is explained by exploiting the analogy between discursive entries and exits, between action and perception. The relation between the normative and the causal aspects of rational willing or practical reasoning is explained by appealing to the causal efficacy of the deontic attitude of *acknowledging* commitments: acknowledgments of doxastic commitments can be reliably differentially elicited as responses to environing situations in perception, and acknowledgments of practical commitments can reliably differentially elicit performances as responses in action. Reasons can be causes because deontic scorekeeping attitudes can play both normative and causal roles. There is much still to be learned about the empirical details of the differential responsive dispositions that make possible these discursive entry-and-exit practices, but it is not hard to understand in principle how there can be such things. We are rational creatures exactly insofar as our acknowledgment of discursive commitments makes a difference to what we go on to do—on the side of action, insofar as we incorporate a connection between what is expressed by 'should' and what is expressed by 'shall'.