

SOME REFLECTIONS ON THE SPORT OF LANGUAGE

Mark Norris Lance
Georgetown University

To understand the relation between language and belief we must revise the image of the language game. Language, I claim, is essentially a sport and mental states are among the equipment with which that sport is played. The primary point of this picture is to allow us to account for one of the central motivating puzzles about the intuitive notion of belief, namely, that beliefs are both essentially causally engaged physical entities and at the same time essentially individuated according to public semantic norms. Unlike elements of games which are purely formal, or normative, entities, a football is essentially characterized both in terms of its physical makeup and the normative proprieties governing it. Similarly, mental states are essentially physical states of brains, and the like, with specific causal roles, and also individuated according to their normative status.

The resulting picture thus takes seriously the embodiment of belief states, and linguistic dispositions generally, while retaining a conception of them as essentially playing roles in the public linguistic practice. It stands, therefore, in stark contrast both to the normative formalism of Brandom, who takes semantic content to be a matter of inferential role and understands the latter purely in terms of public normative status, and also to the reductive naturalistic accounts of causal theorists, functionalists, and others.

Given constraints on length, the view is presented in a way which requires patience on the part of the reader. Although well known arguments are gestured at, many other approaches to these issues are dismissed in a cursory manner. I thus make no claim to have shown that one must follow the sort of account given here, only to have made it clear that there is such an account, and why one might be inclined toward it.

1. The Fundamental Problem of Belief

Twentieth century philosophical reflection on belief is distinguished most clearly from that of earlier periods by a rejection of the view that beliefs are immediately perceivable. In the third person case, it has always gone without saying that beliefs are posits, but under the pressure of challenges from such disparate sources as Freudianism, reductive materialism, pragmatism, and the interpretationalist approaches to mind and language of such philosophers as Quine,

Davidson, Dennett, and Lewis, the transparency of the doxastic has been largely abandoned in the first person case as well.

If, however, the category of belief in general, and typical specific beliefs are to be understood as posits, two questions must arise. First, what sort of posit is belief (or a belief) and within the context of what broader project does it find its use? Second, given this context, is it a posit that we ought to continue positing? There are at least two broad traditions regarding the first question, and it is becoming more and more tempting, within each of them, to embrace a negative answer to the second.

Both the temptation to reject belief as a useful theoretical posit, and also the dispute regarding the general context it would have to function within, arise out of the curious dual nature of belief. On the one hand, the category of belief is distinguished from other mental categories in causal terms. The difference between beliefs, desires, emotions, dispositions to devil's advocate, entertainings, etc., lies in the way in which beliefs as a body produce behavior, especially non-linguistic behavior. On the other hand, beliefs are individuated semantically. We identify particular beliefs as "the belief that P," where 'that-P' is an intersubjectively available content in public language. Many have argued that this latter aspect implies that the individuation of a belief state involves far more than the physical state of the believer, encompassing (at least) facts about that individual's society and immediate physical environment.

One strand in the philosophy of mind focuses on the first of these aspects, the causal. Since the folk category of belief is defined by its role in producing behavior, it is quite sensibly assumed that beliefs function globally as posits within a folk theory designed to explain and predict behavior. Our second, evaluative, question then concerns whether such posits survive as a significant element of those scientific accounts of behavior which must necessarily supplant folk explanations. Though they differ in their final attitude toward belief, functionalists, eliminative materialists, and Dennettian "stance-theorists" all fall into this broad theoretical perspective in virtue of the assumption that the category of belief exists as a potential predictor and explainer of behavior.

The other strand in twentieth century philosophy of mind has its historical roots in Wittgenstein, Heidegger, and some work of the pragmatists, while reaching its most sophisticated form in Brandom's development of Sellars's work. On this conception, mental state attribution is parasitic on public language and is thoroughly tied to epistemic evaluative projects. To attribute a belief is, roughly, to attribute an internal state which disposes one to make a public move of asserting a propositional content within the social game of giving and asking for reasons. On this view, we begin with an account of public asserting which, on Brandom's conception, is the playing of a public chit which has the normative significance of committing the player to defending against all appropriate challenges, the proposition asserted. Asserting is, in a slogan, an undertaking of justificatory responsibility.¹

Though the Sellarsian account must begin with publicly available acts which constitute assertings—since the responsibilities undertaken are to respond to chal-

lenges reasonable within the context of a dialogical practice—this does not preclude the possibility of internal states individuated “on analogy with” (in Sellars’s formulation) external language. One can have an internal disposition to play an external token within the game of giving and asking for reasons, and one can be in an internal state which commits one already to such an act of public normative undertaking. In either case, however, beliefs are seen not as causal explanatory posits in a theory of behavior, but as normative posits which round out the epistemic characterization of an agent in terms of commitment, entitlement, rationality, and the like. The second question as it arises for this tradition, is whether we ought to fill out—in terms of *belief-like entities*—the normative characterization of people which begins with their explicit assertions.

1.1. Two routes to eliminativism

Surely the most plausible and well-developed attempt to carry out a reduction of beliefs to causal-explanatory states of organisms is functionalism and it seemed, for a time, to make room for a genuine notion of belief, since computational roles—causal state transitions within a mind—mirror inferential contents in an obvious way. Reflection on this “inferential” dimension of internal computational states promised an account of propositional content, while the functionalist account of behavioral control allowed for distinctions between the various propositional attitudes.

Focusing on attitude rather than content, *belief* that P can be distinguished from entertaining the possibility that P, lying that P, hypothesizing that P, or defending P for the sake of argument. Belief that P is a state which is defined in terms of both intra-linguistic transitions and language-behavior transitions, that is, in terms of causal-computational connections to other belief states and also by connections had by it—perhaps in union with other states—to behavior. So one who believes that P will believe things which follow from it and also perform actions on the basis of the belief. Devil’s advocating that P, for example, is a state which simply preserves the intra-mental connections of belief that P, but severs the connection to action. Someone who is defending P for sake of argument will make the same computational moves and be disposed to the same public assertional undertakings—but will not base non-linguistic behavior on the claim. While one might well deny that this is a complete account of the difference between the two states, urging also, perhaps, an affective difference between them, this sort of fact is certainly an essential constituent of our grounds for a refusal to posit genuine belief.

On this line, there is a certain priority that belief enjoys in such a systematic account of the causal engagement of computational contents. Beliefs are the states which, we might say, are *fully* engaging within the individual’s behavior-generating psychology. That is, the functional architecture of the individual will be defined so as to include input states—connections between states of the sensory organs and cognitive states—internal transition states—connections between beliefs and other beliefs, for example—and output states—connections between internal states

and behavior. The overall *telos* of this system will be to mediate behavior. The overall evolutionary *telos* of the system will be to mediate behavior in the light of input, evolving so as to produce behavioral variability with situation which is suited to keeping the organism from being eaten by tigers and the like. All the processing which goes on in between functions in the first instance to facilitate such transitions.²

Also, of course there will be a good deal of interdependence among the internal states. Typical computational transitions will depend not on a particular belief state, but upon that state and a range of background states. Similarly, it is typically a belief, a desire and a large body of background states which fully determine behavior, but however the details of the mechanism are understood, beliefs are those states which engage with the entirety of this integrated and interdependent mechanism. Hypothesizings, entertainings, devil's advocatings, etc. all bracket to a greater or lesser extent some range of this integrated functional structure. They are to be understood as states which result when one addresses a proposition while withholding some aspect of the usual functional totality which is the belief state, thus becoming modifications, by limitation, of the believing state.

So the functionalist project seems well suited to making the distinctions between these various attitudes, which is not surprising since the project was one of causal explanation and these distinctions all have to do with the sort of causal engagement had between content and individual. When it comes to the other aspect of belief, however—its individuation in terms of semantic content—things do not go so smoothly. Semantic content, as many have argued, does not supervene on the physical state of the individual. I may not know that Fermat's theorem follows logically from various simple mathematical beliefs I hold—and that implication may not have any significance for my action or other belief—but the semantic content of the simple beliefs are such as to entail it nonetheless. Or I may, to use Kripke's example, assign exactly similar psychological roles, descriptive contents, etc. to 'Feynmann' and 'Gel-Man', but if I believe that Feynmann is going to be at the party, this belief is true or false according to whether *Feynmann* is. It is a belief that Feynmann will be at the party, then, in virtue of things outside my head. The reader will also recall the cases developed by Tyler Burge concerning such concepts as 'arthritis' which again show that the semantic content of belief varies with societal factors beyond the individual.³ Finally, John O'Leary-Hawthorne and I have argued that semantic content outruns, conceptually, even the totality of dispositions to behavior across society.

Thus, it would seem that beliefs are ill-suited to serve as causal-explanatory posits in individualistic psychology and, so, it has become common to reject, in one way or another, the dual life of belief in favor of an entity more naturally suited to individual causal explanation such as narrow content—essentially idiosyncratic computational role—or something else even less belief-like.⁴

Similar eliminativist tendencies can arise within the Wittgenstein-Sellars tradition as well. Though the semantic individuation of beliefs is as well motivated on this account as it is for public language, what is apparently irrelevant is the

causal differentiation of belief as a category. If the purpose of doxastic attribution is merely to record an agent's epistemic commitments, and dispositions to undertake others, then there is no clear reason why we should care how the disposition connects to non-linguistic behavior. Imagine two people who assert that junk food is bad for you. Suppose one believes the claim—as we would say intuitively—while the other merely asserts it to play devil's advocate. The difference between these two lies in the fact that the former state will incline the person to avoid eating junk food⁵, while the latter will have no such causal force. But this difference does not obviously impact upon the agent's epistemic status. Either person is a legitimate target of justificatory challenge, and both must meet these challenges with the same arguments.

For this sort of reason, just as many in California advocate the elimination of beliefs on the grounds that the notion seems gerrymandered from the point of view of mature psychological explanatory theory, so many in Pittsburgh would do away with the concept in the game of giving and asking for reasons. In Brandom's official vocabulary, there are only the normative statuses of commitment and entitlement. Belief is conspicuous by its absence from *Making It Explicit*.

It is my desire to vindicate the notion of belief from within the Sellarsian project and, thus, I focus on whether there is a genuine need for causally individuated belief as a posit within the philosophy of language and epistemology. That is, if we take language to be a Sellarsian game of giving and asking for reasons, see asserting to be an undertaking of justificatory responsibility within that game, and postulate only those internal states necessary to round out the normative description of the individual required for linguistic and epistemic evaluation, do we have a reason to commit ourselves to the concept of belief?

2. On Games and Sports

The difference between a sport and a game lies in the degree of abstraction from embodiment. The normatively individuated elements of a sport may also essentially be instantiated in particular physical sorts, while tokens in a game are, to a large degree, abstractions from the physical, defined purely by normative roles. For something to be a rook is for it to be the element of a game of chess which is required to begin the game on a1, a8, h1, or h8, is allowed to move from 1 to 8 squares vertically or horizontally so long as it does not pass through an occupied square or beyond the edge of the board, etc. A huge variety of things have been caught up in such a web of normative constraint: pieces of plastic, wood, marble, and metal in a variety of shapes and sizes, humans, and clusters of electrical impulses functioning as digital codes within computers to name just a few.

We *could* define a 'snook' to be something which is both a rook and between 1.1 and 1.2 kilograms of metal, but there would be no motivation for doing so. Though appropriately sized hunks of metal might be interesting from the point of view of certain technological or scientific projects, the line between those currently subject to the norms of rook-hood, and those which aren't, is not a marker

of a significant distinction within these projects. Similarly, weight and composition is of no concern to the game of chess.⁶

Football is similar to rooks in that for something to be one essentially involves its being caught up in the constellation of rules definitive of football. Something must be the thing the movement of which by a center constitutes the start of a play, and hence licenses defensive players in coming across the line of scrimmage, in order to be a football. However, in this case, there are also clear and specific physical constraints on the type of entity which can legitimately instantiate the normative role of football. Footballs must be made of specific materials, and fall within narrow size and weight tolerances.

This embodiment requirement is intimately connected to the fact that sports are normatively structured systems of *actions*. A track race is one in which we try *to run* according to certain rules and in the pursuit of a particular goal. The norms of baseball, in the first instance, constrain such actions as pitching, batting, fielding, etc. Norms governing the physical instantiation of the elements of the game⁷ are motivated by and flow from the initial starting point of the physical instantiation of actions in human bodies.

In the case of a game, however, it is the events resulting from the game which are primarily constrained. Here, for example, it is the fact that the knight has gone from b1 to c3 which is at issue, not the act of moving it. Such acts take place, of course, but merely as means to the end. This is not to say that there aren't constraints on the acts of playing a chess game. There are rules about moving the piece one touches, about how much time can be taken, etc., but as I emphasized the distinction between games and sports is really a matter of degree. And there is, in fact, a huge gap in the level of essential embodiment between something like football and something like chess. To see this clearly, note how much more is conveyed by a score of a chess game—something which tells you only the normatively defined moves which, for all the score records, could have been instantiated in just about any sort of physical events following from just about any other actions seeing to them—than would be by a similar log of a football game which recorded ball position, score, possession, and down for each successive play while leaving out what all the players *did*. Televised chess and newspaper box scores of football are both boring.

3. Language as Game and as Sport

The dominant metaphor within the tradition we are considering has been that of the language game. Wittgenstein and Sellars could fairly be said to have built their conceptions of the linguistic upon the image of language as a game, or set of games. What our discussions of the dual nature of belief and of the distinction between games and sports suggests is that we consider whether this image needs to be modified to that of a “language sport”. Wittgenstein and Sellars both would have agreed that language cannot be understood except in relation to mind. In particular, assertions must be seen to be the products of beliefs and assessed

against the context of a background of beliefs. Sellars, in particular, is quite explicit about this. Though he denies that we can explain external linguistic content in terms of a logically prior notion of internal mental content—no internal state, for Sellars, can count as contentful unless it is normatively embedded within the context of a public language—he insists nonetheless that there cannot be language without mind.

Beliefs, though, are commonly understood to be entities which are causally efficacious, which collaborate with desires to produce behavior. Thus, among the “chits” in a language game are entities—beliefs—which must be so instantiated as to produce (non-linguistic) behavior in the relevant ways. In this much, the tokens of the language game are essentially embodied, and in this much language is a sport.

Or so it would seem intuitively. Brandom’s development of Sellars precisely calls into question the status of language as a sport. On Brandom’s view we certainly need internal states. A fully articulated epistemic normative agent must embody a wide range of dispositions to undertake publicly commitments in the face of various linguistic or non-linguistic promptings, but the physical instantiation of these dispositions—even such abstract properties of them as how they connect with non-intentional behavior—is beside the point. Brandom, we might say, is the first philosopher to try consistently to understand language as a game rather than as a sport.

4. Beliefs and Brandomian Epistemic Appraisal

In bridge, as in language, the immediate status of a player in the game is a function of what cards she has played. This determines, in part, her score, what sorts of other cards she can play next, and what cards others can play. However, the cards played so far are not the only determinants of normative status. It is also crucial what *unplayed* cards one holds. Not only do these determine the range of possible plays in the future—one can play only those cards one holds—but also the *significance* and *permissibility* of various plays depends upon one’s “internal” holdings as well. Thus, it is permissible to play a 2 of hearts after someone leads a 3 of clubs, only if one holds no clubs. Indeed, for a physical instantiation to be a hand just is for it to be that which is taken in the game to matter to all these normative issues, and features of hands have a function within the game just in case they make a difference to some such normative issue.

Similarly in the game of giving and asking for reasons an assertion may be rational or irrational depending upon whether one has various other dispositions to assert in one’s “hand”. If asserting is understood as the placing of normatively contentful cards on the public table, cognitive states are whatever features of the player make a difference to the significance and appropriateness of such playings. Players within the game of giving and asking for reasons will attribute internal states with an eye to their utility in codifying epistemic normative statuses of persons.

Assertings license others in responding in particular ways. For Jones to say that the economy is going to be expanding in the next quarter, is—among other things—for it to become appropriate to ask *him* to explain what will compensate for rising interest rates. Obviously such proprieties depend crucially upon a huge range of attributed internal dispositions to undertake additional commitments in various circumstances. Every propriety we explicitly assign to the discourse will presuppose a vast array of other attributions, the totality of which could never be made explicit.

Ultimately the goal of the practice is to determine those claims to which an assertor is epistemically entitled. Such determination allows for communication—the passing on of entitlement from one agent to another—and inferential inheritance of the entitlement—theoretical and practical inference to other claims and non-linguistic acts which are taken as justified by virtue of following from the justified claim. Thus, we all come to recognize that it is warranted to drop our work we are otherwise obliged to perform and to scurry up the tree in virtue of one agent's justified assertion that there is a tiger nearby.

In this practice of placing a person in the game of giving and asking for reasons—in this process of assigning commitments on the basis of assertions and entitlements on the basis of people's ability to defend the public propositional contents of those sentences asserted—we do take account of internal states which lead to the making of these assertions. Further, our folk practice seems also to make the distinction between an agent who is merely disposed to assert that P and one for whom the state responsible for this disposition engages with the agent's broader psychology so as to lead to predictable patterns of behavior. In ordinary language this is precisely the distinction between what is merely said and what is believed. One may say something—indeed, even defend it with solid evidence—without believing it, and that distinction seems precisely to have to do with whether the commitment in question arises from a performance which is causally engaged with the agent's psychology so as to influence behavior.

So there certainly *is* a notion of causally significant belief within the manifest image of commitment attribution, but *should* there be, and if so, why? That is, just as in the case of a theoretical project of causal explanation, we can ask whether the project of assigning epistemic and semantic status has any *use for* the distinction between mere commitment and causally significant commitment, i.e. belief. It could be that as mature science replaces folk causal explanation we will lose the need for semantically individuated psychological states, while at the same time folk epistemology will be replaced by a mature normative system which sees these causal distinctions as ultimately irrelevant.

This line of argument is instantiated, as we said, in the recent work of Robert Brandom.⁸ Brandom, develops his philosophy of language purely in terms of the normative statuses of commitment and entitlement. It may be, on Brandom's account, that there is some probabilistic connection between action and commitment—or even causal connections which hold in a large portion of cases—since there may be some true rough generalizations concerning connec-

tions between commitments people undertake and their actions, but there is no necessary causal component to commitment. It may even be conceptually necessary that people are *generally* disposed to act in accord with their epistemic commitments, but it is obvious that they are not always so. What matters for one's epistemic appraisal—for the assessment of what claims one should be taken as committed to defending and of what claims one is entitled to use in theoretical or practical inference—is simply the content of those claims one has undertaken to defend. Whether those undertakings engage with a broader action-guiding psychology is really beside the point of Brandom's project. Whereas psychology had good use to make of the distinction between belief and the various other propositional attitudes, but found no point for semantic individuation, the epistemologist obviously wants to individuate assertions in terms of semantic content, but has no clear reason for sorting the psychological causal antecedents of assertion between the various attitudes which might give rise to them.

5. Toward a Vindication of Belief

Sellars saw the attribution of semantically contentful, causally significant states to be an essential element of the manifest image, and claimed the latter to be itself ineliminable from our thought. (It is to be joined with, rather than reduced to, the scientific image.)⁹ However, the reasons for this are a bit obscure in Sellars's own work. Note that he was quite willing to allow a theoretical scientific image to supersede the manifest for purposes of causal explanation, so one wonders why a similar process could not take place in the epistemic project as well.

On Brandom's account, we place an agent normatively without committing ourselves to the causal structure of their psychology, aside from the bare presumption that some internal states dispose the agent toward acts which constitute the normative status in question. On this development, the manifest image simply splits into two autonomous theoretical realms, the scientific image and the normative image.

In order to avoid this course, we must find a systematic way in which distinctions concerning the broad causal-functional engagement of the internal state responsible for an assertion makes a difference to the rational epistemic evaluation of that assertion. That is, we can vindicate belief if we can explain why, in addition to the public defenses one makes of an assertion and the implicit commitments and entitlements characteristic of one's place in the game, the extent of the causal engagement of these states matters to one's normative position.

5.1. *Output-based responses—discourse as guiding and coordinating action*

In order to assess the virtues and limitations of such explanations, we need to return to the general *telos* of linguistic practice. In our earlier discussion of the causal function of organisms we noted the familiar point that the evolutionary

role of computational systems within the individual is to facilitate functions from sensory input to behavioral output. While the *nature* of various internal states depends upon the details of its functional relations to other states, the *upshot* of the whole system—that upon which it is ultimately tested in the court of biological viability—is the way they relate input to output across various circumstances. The point of noting this is not to suggest that there is any sort of reduction of computational role either to input-output relations, nor to evolutionary function. But the evolutionary *telos* of our psychological architecture nonetheless provides a necessary conceptual background for any understanding of psychological function.

A crucial insight found in the work of Heidegger, Wittgenstein, and Sellars is that despite deep connections between the linguistic and the mental, one does not correctly situate the problem of explaining semantic content within the broad *telos* of the biological evolution of the individual, but rather within the more complex social evolution of practice. It is only by looking at the development of persons *qua* members of social practices—Being-with-others, forms of life, participants in the game of giving and asking for reasons—that we see what is at stake in the evolution of semantic norms.

What is at issue socially, and that which provides the context for an adequate understanding of semantic content, is the coordination of social action. Social norms aggregate within communities, and the norms which govern discourse survive or perish, in the first instance, on the basis of their ability to facilitate social coordination, to produce structures of mutually reinforcing action patterns across the community which allow for the continuation of the social practices and institutions within which they play a role.

This point immediately suggests a line of response to our worry about the causal function of semantically contentful states. Suppose that the ultimate point of linguistic discourse is the coordination of behavior. Grant, further, the central assumption of our problem: that semantic content is a notion which arises only within the public space of linguistic discourse. Then there appears to be no mystery, since for social coordination to result from the attribution of normative statuses, these statuses must causally connect with behavior.

In the end, suitably understood, there is much to this line of argument, but several additional components must be in place for it to be successful. For example, it might seem that since coordination requires knowing what others will do, a clarification of the response can be found in the idea that semantic state attribution is a mechanism for the prediction of complex behavior. This is the central starting point of Dennett's analysis of the "intentional stance,"¹⁰

No doubt this idea is generally congenial to the Sellarsian tradition. It is undeniable that we gain a predictive foothold on others' behavior by way of mental state attribution, but it is crucial neither to overestimate this role nor to misunderstand the way such prediction as we do achieve arises. While the present context allows for neither an elaboration of the evidence, nor a detailed examination of Dennett's view, it seems clear to me that our predictive abilities *vis a vis*

other people are quite limited. Imagine trying to predict what someone would do in a typical situation of human interaction. Will the next action of the person cooking dinner be to reach for the oregano, to protest the treatment of Kurds in Turkey, to stir the pasta sauce, to brush back his hair, or to flirt with a nearby philosopher? Obviously if the goal is anything like the specific prediction of behavior, the intentional stance is of little use.

The point is not to deny that there are predictive gains inherent in the attribution of contentful states, but to emphasize three features of such prediction. The first point is that insofar as we predict what others will do in the ordinary course of things, our predictions are basically negative. We expect someone not to leap into the pasta sauce, or to begin twirling the cat around like a baton. Of course in no case are these expectations all explicit. We would be inclined to say that these unexpected actions are unlikely, *if asked*, but we do not formulate a list of such things. Not only do we not do so in the ordinary course of things, but it is clear that no complete list could be formulated, for the simple reason that there is an open-ended set of such actions, each of which is equally unexpected.

The second, and more important point, is that what negative predictions are embodied in our practice are deeply tied to our own involvement in that practice. Human interaction is largely a skill, but it is a social skill. Our practices are complex dances we learn together, not by each one learning alone a particular sequence of steps which they then predict will avoid stepping on the toes of the other dancers—or even less by learning a choreography of others' moves first, and then formulating a suitable way of joining them—but by coming together, through mutual trial and error, to trade off sequences of improvisations in ways which generally cohere according to a sort of emergent social logic.

Thus, what we expect from others is that they will continue the dance in a recognizable manner, do something which will strike us as neither bizarre nor incomprehensible, perform in a manner we will spontaneously know how to react to. Anything which leaves us at a loss for a way to respond, even which requires us to consider explicitly what is going on or how to respond to it, is one of the actions we can be said to have predicted not to occur. It is, I submit, precisely our disposition to so react which constitutes the primary sense in which we can be said to “predict” actions.¹¹

This brings us to the third point, which has specifically to do with mental state attribution. In the ordinary course of interacting with others we do not explicitly attribute mental states at all. As we walk through a crowded room at a reception—an amazingly complex activity which requires a deep tacit understanding of patterns of human interaction—one need not posit any particular mental states in order to carry off the task. We don't attribute a range of mental states to the woman with the martini so as to guess how to avoid bumping into her because we don't consciously assess how to avoid her at all. It is only when something goes wrong that postulation arises as an issue. If someone suddenly starts gesticulating wildly, we may attribute to the agent anger, or a belief that there is a stinging insect in the vicinity. We do so in order to return ourselves to a

situation from which we can skillfully manage our interaction with the other person.

Such an ability to turn a context of confusion back into one in which one knows how to behave depends upon our ability to empathize with others. We know that one who is gesticulating in an effort to ward off a bee is not someone who is going to attack us, and that the gesticulating is likely to cease shortly. We know this, in the usual case, not as a sort of theoretical knowledge, but because it is how we would behave. Thus, we predict via psychological attribution—insofar as we do—by empathically integrating the attributed state into our own psychology, and seeing what range of actions seem reasonable and unreasonable for ourselves.

Having gone this far in spelling out the folk practice of behavioral prediction via psychological attribution, we can see that it is not going to provide an answer to our worry about belief. The suggested solution was that we attribute beliefs—causally integrated internal commitments—so as to achieve this kind of predictive advantage, so as to be able to re-orient those with unusual behavior within the space of actions we are capable of coping with in the ongoing social practice. However, while it is true that this reintegrative process of empathic interpretation requires that we relate the other's internal states to our own, there is no explanation here of why these states should be individuated according to publicly available linguistic contents.

This is to say that the need to predict in this way is fully consistent with the radical Brandomian separation of the causal and the normative. Prediction requires only that one attribute a computational state to the agent in question. While it is true that the state attributed must causally engage with the agent's behavior producing psychology, there is no need for the individuation to be governed by public epistemic and inferential norms. Indeed, the process of empathic understanding we have been sketching will only work if the attributed state functions computationally in roughly the same way for the object of attribution as for the attributer. So it would seem that the very best way to go about it would be to attribute precisely those computational contents one is inclined to enter into in various circumstances.

There is, however, another aspect of social behavioral coordination which seems to bring in the linguistic more directly. So far we have considered only the interpretive activity which goes on when we attempt to cope with unusual behavior. I have urged that such reflection on folk prediction offers no explanation of why the attribution of causally efficacious states must be part and parcel of the normative state attributions which form the practice of linguistic discourse, and out of the structure of which semantic contents arise. However, a moment's reflection on this linguistic practice makes clear that there is a deeper connection. We not only attempt to *predict* the behavior of those who act oddly; we *argue* with them as well.

We might try to convince the person in our earlier story that waving his arms will actually increase the probability of his being stung by the insect we take him

to be aware of. In so entering into the discursive practice, we hope not only to redeem entitlement to our claim, not only to place an onus upon the other fellow to agree with us, but most significantly we expect his doing so—his very state of undertaking a commitment to the truth of our claim—to cause him to stop waving his arms. Plausibly here is the connection between linguistic content and full causal engagement with behavior guiding psychology: we argue not merely for the purpose of generating normative statuses, but to influence the behavior of others. Indeed, this role of language would seem to be the fundamental fact underlying what we called its *telos* of social coordination. Social coordination is not merely a matter of figuring out what others *do* believe so as to guess what they *will* do. Rather, it is more centrally a matter of showing them what they *should* believe so as to influence how they *ought* to behave.

Though it is true that the connection between dialogue and action, properly understood, can be seen to require attributions of states with both public normative criteria of individuation and full causal engagement with agentive psychology, more is needed to make out this connection. The above observation is that it is essential to the practice of linguistic discourse that people alter their behavior in the face of argument.

The proposed solution to our problem is that the public nature of language, and the public normative criteria of individuation of its commitment-types, is transferred to those states which cause behavior via the language-exit transitions built in at the very core of the language game. When we argue, the assertions we make are individuated according to public normative criteria precisely because they are put forward as counters in a public game, as generic reassertion licenses. We have just now noted, however, that one of the fundamental purposes of this practice of public argumentation is to influence others to behave in various ways. We argue largely so as to come to agreement, and this practice of coming to agreement in propositional contents develops against the background of a social need to come to agreement—better, coordination—in non-linguistic behavior.

This observation, however, does not yet solve our problem of vindicating a dual-nature posit of belief because the fact that normative posits place rational constraints on an agent's non-linguistic actions does not imply that we postulate an essential causal engagement of those posits. Nothing in the above discussion suggests that the members of the linguistic community *will always* act as is rationally required by the commitments they are obliged to accept. Just as in the case of theoretical inference, practical inferences can be missed, ignored, etc. That is to say that the requirements of the language game, *vis a vis* non-linguistic action, is *normative*.

To be sure, since the guiding *telos* of the practice is to provide *actual* coordination of behavior, discourse as a practice would break down were people generally to fail in their rational commitments to action. This is just a species of the more general Davidsonian point that we cannot interpret a community as undertaking linguistic commitments unless we take them to be generally rational. But there is no reason to require that the very state which is a disposition to assert

always be a state which guides behavior in the way made rational by the content of that assertion. Further, and this is the crucial point, we have not yet explained even why the distinction between assertional dispositional states which are appropriately causally engaged—i.e. genuine beliefs—and those which are not matters to the practice of linguistic discourse. We have shown only that it is a part of the practice of discourse that normative proprieties attach to action in virtue of the normative status of assertions, and that this must holistically connect to actual behavior. That is, people must in general live up to their epistemic duties to act. Both of these are necessary if linguistic practice is to serve its coordination function, but they stop short of requiring a distinct category of belief. To put the point another way, the fact that we want to keep track of which of those non-linguistic actions made obligatory by discursive practice are actually engaged in and which are not does not imply that we need to keep track of those implicit commitments which are causally engaging and those which are not.

To see this, let us imagine a character who habitually enters into linguistic discourse in a way we would intuitively characterize as not involving genuine belief. Suppose this person frequently utters claims he is not antecedently prepared to act on, perhaps—as we might say intuitively—to see how they are received in the public epistemic arena. Such an agent—we might call him ‘Clinton’—could have a range of dispositions to assert which remain largely disengaged from his behavior guiding psychology. There will, of course, be *computational* states which do engage with this psychology, but if these do not form the source of his assertional behavior, we have no apparent need to individuate them in terms of public content.

Let us suppose further, that when Clinton does utter assertions in the public realm, he argues for them in fairly normal ways. Further, if one of his claims is refuted, he may well systematically see to it that he does not act in ways incompatible with whatever claims have been publicly demonstrated. It is, perhaps, overwhelmingly important to him to avoid criticism. It appears that he could live up to the public standards for the epistemic regulation of behavior, while regularly failing to produce assertions as a result of dispositions which are fully causally engaged.

Intuitively, this is a strange person indeed. We would characterize him “as not really believing what he says,” “as making assertions insincerely,” “as adjusting his public commitments and his non-linguistic actions in a way which is non-genuine and lacking in true commitment.”¹² The intuitive folk attributional practice, in short, takes the psychology of such a person to be radically different from that of a normal participant in linguistic practice, but our current reflections on the normative structure of that practice have shown us no way in which Clinton will fail to live up to our standards of rationality and, so, no reason why those standards should take account of the difference between Clinton and people who really believe what they say.

Now it is worth remarking that there certainly *is* a difference between Clinton and normal folk, one which is apparent from a purely descriptive stance.

There is a far more complex causal link between whatever features of his internal architecture lead to his production of language and those features which lead to the rest of his behavior. So the descriptive psychologist will see Clinton as an odd and interestingly different sort of person. However, *the descriptive psychologist* has no reason to individuate anything on the inside of Clinton in terms of public semantic content. What we desire is an explanatory project which *both* individuates internal states in terms of public conceptual role, *and* sees a relevant distinction between Clinton and people with integrity.

5.2. An input-based response: the rational significance of causal engagement

This last way of putting the distinction—that it is one of “integrity”—may point us in another direction. Clearly the hypothetical assertor we have been considering lacks integrity in a literal sense; his action guiding psychology and his assertional dispositional states are not integrated, at least not in the usual way. But to say that someone lacks integrity also implies a substantial evaluation. Integrity, we seem to think, is a virtue, and my characterization of Clinton clearly implied an evaluation.

It is not a new idea to suggest that there is normative significance to being the sort of entity whose behavior is causally responsive to reason in the usual way. Many moral theorists would even seek to identify freedom, or moral responsibility, with one’s being such as to bring together the space of reasons and the space of causes.¹³ The current claim is more radical, however. We are trying to show not just that moral responsibility requires that one be causally responsive to the normative force of assertion, but that the entire sport of language requires such an embodiment of the conceptual. We are attempting to derive a category of belief from the generic features of rational discourse.

In the previous section we considered connections between the propositional contents of public discourse and action which reside on the “output” side of discourse. That is, we focused on the *upshot* of commitment attribution, or of arguments for assertional commitments. Our new emphasis on the value judgments which attach to people whose internal states lack causal-conceptual integrity suggests, on the contrary, that we attend to the “input” side of linguistic practice. If we disvalue the very psychology of one for whom dispositions to assert are not fully causally engaged, then perhaps we can find a role that this disvaluing has in the rational normative structure of linguistic discourse. Perhaps then we see integrity as a *rational* virtue.

When considering the role of language in coordinating social behavior, we spoke of the need for argument to carry practical normative import, but substantive argument is obviously not the most common case of language use. Rather, the usual case is that of one person saying something and others simply accepting the claim.¹⁴ Any epistemology which is to avoid skepticism must imply that one enjoys *prima facie* entitlement to most of one’s current commitments. Of course such entitlement is defeasible, but unless one can rely on one’s current commit-

ments barring substantive criticism of them—unless the onus of proof, however weak it may be, is upon claims incompatible with currently held commitments—there can be no justification at all, since there are always a myriad of alternative views incompatible with all one believes and yet consistent with any data one might come up with.¹⁵

Similarly an account of epistemically rational discourse will have to employ what I have elsewhere referred to as a principle of trust.¹⁶ We must also assign a *prima facie* entitlement to *others'* assertions, making acceptance of what others tell us the default position and accepting an onus of proof when we choose to disagree. To deny this would be to allow as rational the simple dismissal of typical claims by one's interlocutors, without basing this on arguments from other claims which they would have any tendency to accept. Such a state of general distrust would eliminate the possibility of language serving any sort of coordinational role. Thus, it would seem to be *a priori* necessary that linguistic communities involve a norm of trust. Thus, we don't even so much as look for evidence when others tell us their name, what kind of food they like, or what time it is, when we look things up in history books, or—most significantly and generally—when we are taught the meanings of words at our parents' knees.¹⁷

Indeed, social status is largely a matter of the range of claims one is taken to have such a *prima facie* authority regarding. To be taken to be a member of the linguistic community—a person—requires that trust be extended for the most part, and in things run of the mill. To play the role of epistemic expert within a community—be it in physics, in chess, or in morality—is to be accorded trust on matters regarding which people in general are not.¹⁸

However, it is not quite correct to say that one's expertise is measured by the extent to which one's *assertions* are worthy of being accorded authority, nor that one's personhood is recognized by the community insofar as they apply epistemic trust to one's *assertions*. The reason is that there is a large caveat which seems to apply to all such phenomena. It is precisely those assertions which one *believes* which are accorded *prima facie* authority.

This is not to say that we consciously distinguish, in the usual case, between those claims an interlocutor actually believes and, for example, those she is merely putting forward for the sake of argument. Rather, the default assumption is that others believe what they say, that an agent's commitments arise from dispositions which are fully causally engaged with her action-guiding psychology, and thus to accord rational warrant to their claims. However, anytime this default is overridden—anytime we attribute a source for the undertaking which is not itself a belief—we override as well the *prima facie* authority granted to the assertion.

Thus, if Jones tells us that she was born on a particular date, the rational course is not even to question this. People are experts in these matters, and it is not in order to demand evidence unless one is prepared to offer positive arguments that the circumstance is unusual. If we were told, however, that Jones doesn't actually believe her claim, but is rather trying it out “for the sake of

argument,” or some such, things are quite different. Then the claim loses all *prima facie* authority and must confront the court of rational argument from a neutral position.

The same goes for more specific expertise. We may well trust the judgment of the chess master when she says that white is better in a given position, or the judgment of the moral expert who says that it would be wrong to pry into someone's reasons for action, even if these experts can articulate little in support of their judgment. The mere fact that they so judge the matter is often good enough. However, this tacitly presupposes that they believe the claim in question. Were they to be expressing mere dispositions to assert which did not play a full role in their action guiding psychology—if the chess master was disposed to bet on black, if the moral expert had no inclination to refrain from prying—their expertise would be moot and we would demand substantive justification before accepting the claim.¹⁹

This is all to say that epistemic trust of claims derives from trust of people. We take agents to be worth trusting, and for this reason we trust the claims they produce. It is perfectly correct to put this as the claim that it is a fundamental feature of linguistic practice that we take other members of the practice to be—by and large, and barring argument—reliable reporters across normal circumstances, so long as we don't think of this as providing an independent justification of trust. We do not have some independent means of checking on the general reliability of the other members of our linguistic community. So dependent is our knowledge upon others that such a process of testing could never get off the ground without a prior posit of warranted trust. Whether characterized as a socially instituted onus of proof on challengers, or as the general reliability of the mechanisms which produce non-inferential beliefs in people, trust is a basic methodological assumption which undergirds the very possibility of rational discourse.

The paradigm case of such rationality in virtue of reliable production is perceptual knowledge. We take people to be justified in making immediate perceptual reports not because of any argument they can give, but because of the fact that we take the claim to be a result of a belief which was produced by a perceptual mechanism which is reliable. The distinction between justified and unjustified observation reports is most typically a result of a differential assessment of the reliability of the perceptual faculties of the reporters, whether this differential is a result of differing hardware—one person is color-blind, for example—different training, or differing circumstances of observation.

The relevance to our current topic comes in the fact that the process to which reliability is reputed crucially involves the production of a specifically belief-like attitude. It makes a crucial difference to our assessment of the reliability of processes of perceptual judgment formation whether they result in a belief attitude. Contrast a typical situation in which a person, after looking, comes to have a belief that there is a table in front of her, with one in which a similar irradiation of the eyes leads to a state in which she is inclined to play devil's advocate on behalf of the claim.

This explanation of the epistemic relevance of causal-psychological engagement generalizes from the perceptual case to the extent that one thinks of epistemic authority as arising from faculties of discernment not reducible to any sort of explicit reasons available to the assertor. Views of epistemology which emphasize such tacit skills have recently arisen in two contexts. In epistemology and the philosophy of science skill based models of intelligence have been pursued by Dreyfuss, Hacking, Rouse, and many feminist epistemologists. These approaches emphasize ways in which experts in various fields come to their judgments on the basis of perceptual and judgmental faculties which remain largely implicit. In all cases, it is clear that their status as experts depends on the assumption that it is their *beliefs* which are being expressed.

The claim that discernment is a skillful exercise of a faculty of expertise which must remain largely inarticulate also arises within the context of recent work in virtue ethics.²⁰ McDowell, Nussbaum, Murdoch and others all emphasize in one or another way that moral knowledge often has the character of simply *seeing* that the situation calls for a given action or is characterized by a particular moral quality. Again, whether such a claim is to be taken seriously will depend upon whether it is one which flows from the expert faculty. Moral experts don't always make claims via their expert perception, after all. Sometimes they conjecture, play devil's advocate, defend a claim so as to see where it goes etc.

This suggests that belief is indeed a viable posit, which can be seen to have an indispensable role within this sort of skill-based epistemology, an approach which takes epistemic virtue quite seriously. Now, I cannot claim to have shown that this is the only approach to epistemology which can rescue belief as a useful posit. It seems to me that an epistemic system will need the distinction between beliefs and other dispositions to assert insofar as it places importance on trust. Presumably most anyone would grant that it is generally reasonable to accord *prima facie* trust to others' beliefs and to withhold this trust to non-belief-like inclinations to assert. But for traditional foundationalists, certain sorts of "deontologically" inclined epistemologists, or generally anyone who thinks that justification could in principle all be made explicit, the role of trust will be derivative. Though as a kind of shorthand we might postulate beliefs, we could ultimately put our epistemic house in order by recording all the evidence and argument for our view. And presumably, when in the philosophical or scientific mode, this is what we ought to do. So for such a theory, it would appear that the category of belief is of no importance for properly epistemically rigorous agents. Insofar as justification is a codifiable matter, having to do solely with the justifications one does or could produce in favor of a claim, one can and should abstract from considerations of embodiment.

Interestingly, Brandom himself offers one sort of argument that this can never be the right course. On his preferred model of understanding there is, of course, such a thing as explicit theoretical knowledge. But this is understood—in roughly Heideggerian fashion—to be itself a complex sort of epistemic skill, an inferen-

tial skill governing the manipulation of linguistic entities.²¹ What is crucial for our purposes, is that these Heideggerian skills could never be made fully explicit.

Are there other epistemic conceptions which assign an essential role to trust? Perhaps. My own inclination is to think that there are no other viable ones, but this is not the place for such an argument. Nor is it the place to develop a comprehensive account of socially embedded rationality, something which would be a rather large job. Still, the outlines of an account of the significance of various causally founded psychological distinctions is coming into focus and it should be reasonably clear why these would matter to one who hoped to develop this sort of epistemology of social skills.

6. Conclusions

Ordinary language contains a much richer array of distinctions regarding the mode of integration into one's psychology of a disposition to put a claim forward as fodder for giving and asking for reasons than is often noticed by philosophers. Typical philosophical theory usually considers only belief and desire, with perhaps a nod to various imperatival and interrogative attitudes. Such a narrow field easily obscures the issues under consideration here. Bayesians expand the belief attitude to degrees of belief and this is helpful, but ordinary language speaks of seeing that P, suspecting that P, (educatedly) guessing that P, feeling compelled to conclude that P, finding no other choice than to believe that P, persuading oneself that P, recognizing that P, (suddenly) noticing that P, a calm assurance that P, a profound realization that P, reconciling oneself to P, disingenuously putting forward P, "mouthing off" that P, and many more. Each of these attitudes suggests subtly different engagements with functional-computational systems and each invites a different evaluative attitude toward the agent to whom such a state is attributed.

A developed account of epistemic virtue would provide insight into the sorts of authority inherited by claims resulting from each of these attitudes and the related cognitive faculty. It would raise questions about the kinds of objections and argumentative strategies rationally appropriate with regard to a person who is led to assert by each of these states. It should be a pressing concern of epistemologists to develop such a theory of the epistemic virtues which pays due attention to the richness of epistemically relevant psychology.

In §5.1 we considered the ability we have to predict each other's behavior. This ability, it was observed, depends on the fact that we all share a common sort of skill at navigating the socially and environmentally defined worlds we inhabit. Roughly, we anticipate each others' actions because we all behave largely similarly in the sorts of situations we come across and construct together.

This account will be part and parcel of an understanding of knowledge in terms of social skill. Undertakings and attributions of commitment arise against the background of shared practices and shared practical skill-repertoires. Coming

to hold a belief is a way of revising one's repertoire of behavioral skill in specific ways from the usual cultural norm. Thus, when one comes to believe that there is a dangerous hole in the sidewalk, one walks *just as one does*, except that one avoids that spot. We do not build up a theory of how to behave from scratch on the basis of our various cognitive states. Rather, we start with an elaborate and flexibly revisable behavioral repertoire, and introduce cognitive states so as to emend this in ways which are themselves part of this repertoire of behavioral skill.

Attributions work similarly. We attribute states to others so as to understand the ways they are likely to diverge from the cultural norm, attributing states to them and then expecting of them behavior that we ourselves would engage in were we to commit ourselves to the state in question. Why do we not typically attribute non-belief states to others, perhaps seeing them as lacking integrity in the way our earlier discussion of Clinton did? One could do this, but in general *should not*, for to do so is to assign them a lesser epistemic status. The more we read others' moves in the sport of language as having a non-integrated psychological origin, the less we can accord them epistemic trust. This not only excludes them from full membership in the discursive community, but threatens the very fabric of that community.

It is, then, quite correct to say that the usual practices of belief attribution and behavioral prediction from the intentional stance, require the attribution of beliefs, and the significance of the distinction between beliefs and other assertion-generating states. They require this, however, only because they are part and parcel of a practice of normative assessment which arises out of a communal practice in which we are all presumed competent until shown irrational. We attribute beliefs, because we have a linguistic practice which assigns normative statuses beginning with epistemic trust. We have such a linguistic practice because it grows out of a non-linguistic practice of shared behavior and the communal need to coordinate that behavior across changing circumstances.

Language is not a game, but a sport. It is essentially embodied in agents who share an implicit, and largely non-articulate, form of life. Linguistic authority grows out of a sort of mutual respect for those who share that form of life, and it does so precisely because the embodied sources of moves in the public linguistic arena are systematically integrated into the causal architecture of those agents, who are themselves integrated into a coherent social practice. Only by accepting most of what others are inclined to assert, can we come to share enough by way of background that debate over the remaining details can be productive. This is true whether the details are so tedious as the number of buffalo in yonder herd, or as profound as the nature of fundamental forces. Asserting requires the sport of giving and asking for reasons, a sport which is founded, in part, on epistemic trust. The practice of epistemic trust requires the postulation of belief.

Or so the situation seems to me.

I have presented a picture of the relation between language and belief, and gestured at much which goes along with it. I have done so in an admittedly cursory and preliminary manner. Although well known arguments were gestured at,

many other approaches to these issues have been rejected far too quickly. To have done otherwise would have required, just to pick one example, a defense and elaboration of the view of Heidegger, Wittgenstein, Dreyfuss, and Brandom that knowledge is fundamentally and irreducibly a matter of skill. I thus make no claim to have shown that one *must* follow the sort of account given here, only to have made it clear that there is such an account, and to have shown why one might be inclined toward it. This was not, however, a mere exercise in philosophical argumentation. In hopes that it is relevant, I promise that I do believe in it.

Notes

1. See Brandom's Brandom, Robert. "Asserting," [*Nous* 17, pp. 637-650, 1983], and the development of this article in *Making It Explicit*, Cambridge, Mass: Harvard University Press, 1994.
2. That is not to say that there will not come to evolve standards of correct inference which do not directly bear upon survival, or even on behavior at all. Any number of interesting inferential networks can develop once the internal structure is up and running. But the mechanism as a whole is one which will be seen to be quite essentially tied into a role as facilitator between perceptual input and behavioral output.
3. See Tyler Burge, "Individualism and the Mental" in *Midwest Studies in Philosophy*, Vol. IV. Eds. P. French, T. Uehling and H. Wettstein, Minneapolis: University of Minnesota Press, 1979, Saul Kripke, *Naming and Necessity*, Cambridge, Mass: Harvard University Press, 1980, Hillary Putnam, "The Nature of Mental States," in *Readings in the Philosophy of Psychology*, vol. 1, ed. Ned Block, Harvard University Press, Cambridge, Mass, 1980, and *The Grammar of Meaning: Normativity and Semantic Content*, Mark Norris Lance and John O'Leary-Hawthorne, Cambridge University Press, 1997.
4. Of course there remain those who hold out the hope of finding a home for genuine belief within a causal-explanatory account of behavior, and I do not take myself to have argued against them here. My goal is to point to those familiar arguments which motivate the course pursued in what follows.
5. Of course there may be other states which incline her otherwise. The connections to behavior are complex and holistic as we discuss below.
6. I said at the outset that something is a sport insofar as it is essentially embodied, for though there are clear gaps in practice, the distinction is a matter of degree in principle. Thus, there are actual physical constraints on chess pieces in particular competitions. Given human ability, one could not legally insist on playing a tournament game with 400 pound pieces, etc. But these limitations, we want to say, have little to do with the game per se, which can be played with a vast array of differing piece-instantiators.
7. This includes the human elements. So we may restrict participation to women, or to people who haven't used various performance enhancing drugs.
8. Most systematically in his *Making It Explicit*, Harvard University Press, 1995.
9. See "Philosophy and the Scientific Image of Man," in *Science, Perception, and Reality*, Routledge and Kegan Paul, 1963.
10. See, for example, *The Intentional Stance*, Cambridge, Mass: MIT Press, 1987.

11. Of course there are non-primary senses as well. Sometimes, we do make explicit predictions, but the claim is that these are not the usual case, do not account for the general utility of social understanding of others, and that these cases are understandable only in terms of the more fundamental sort.
12. It is suggestive that 'commitment' in ordinary usage is ambiguous between those duties one has regardless of whether one is aware of them, and duties one has psychologically embraced as one's own. When we speak of 'true commitments' it would seem to be the latter, causally engaging, sort we have in mind.
13. For a recent argument along these lines, see R. Jay Wallace, *Responsibility and the Moral Sentiments*, Harvard University Press, 1994.
14. This is not to claim that there could be a language which lacked dialogical norms. Without being embedded within a structure of discursive norms—hence, of inferential norms—no amount of passing back and forth noises, or even of altering behavior predictably as a result of this, could count as a language. The point is exactly analogous to Sellars's insistence in "Epistemology and the Philosophy of Mind," that while no observation can have epistemic or conceptual significance without being embedded in a linguistic practice not all the elements of which are non-inferential reports, there still are genuine non-inferential reports which form a part of the more complex inferential practice. The present point is a dialogical correlate. Acceptance of the claims of others without argument presupposes practices of argument, but there are genuine instances of immediate acceptance, and further, these instances remain the most common sort of interaction.
15. For one version of the argument for this structural conclusion, see Lycan's "Conservatism and the Data-base," in *Reason and Rationality in Science*, N. Rescher (ed.), 1984. A specific version of the claim, more closely tied to the current discussion, can be found in *The Grammar of Meaning*.
16. *The Grammar of Meaning*.
17. For a development of this point and an argument that the default authority of trust cannot be derived from any sort of prior data, see Mark Webb's, "Why I know about as much as you do," *Journal of Philosophy*, 1994.
18. Indeed, other sorts of authority and social power, as Foucault and others have argued, can be understood in terms of a societal granting of trust regarding one's utterances of imperatives. There is much to be said here, but the account is intended to be uniform and to highlight the fact that social power is primarily in all cases a matter of others' accepting one's claims as binding.
19. One obvious consequence of this point is that charges of hypocrisy can be relevant to the reasonableness of accepting someone's moral claims, for the hypocrite forfeits any authority which his unargued assessment might otherwise have.
20. This analogy and some aspects of its development below was suggested to me by Michael Kremer. I am also most grateful to Maggie Little for extended discussions of the connections between particularism in ethics and skill based models of epistemology. Much of my early thinking on this topic took place during a joint seminar we taught on moral wisdom at Georgetown University in 1996.
21. See *Making It Explicit* and "Heidegger's Categories in *Being and Time*," *Monist* 66, no. 3, July 1983, pp. 387–409.