

Intellectual Norms and Foundations of Mind

Author(s): Tyler Burge

Source: *The Journal of Philosophy*, Vol. 83, No. 12 (Dec., 1986), pp. 697-720

Published by: Journal of Philosophy, Inc.

Stable URL: <http://www.jstor.org/stable/2026694>

Accessed: 11-04-2017 02:21 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://about.jstor.org/terms>



Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*

THE JOURNAL OF PHILOSOPHY

VOLUME LXXXIII, NO. 12, DECEMBER 1986

INTELLECTUAL NORMS AND FOUNDATIONS OF MIND*

TWO paramount tasks for philosophy are articulating the nature of propositional attitudes and articulating the nature of the intellectual norms that govern thinking about objective matters. The first task is fundamental to explicating the notion of mind. The second is fundamental to saying what is significant and distinctive about being human.

With respect to the first task, I have previously argued that propositional mental-state and event kinds are *nonindividualistically individuated*. The mental natures of many of an individual's mental states and events are dependent for their individuation on the individual's physical and social environments.¹ The arguments I have given indicate that a person's mental-state and event kinds might, in principle, vary with variations in the environment, even as the person's physical history and constitution, described nonintentionally and individualistically (without relation to the physical and social environments), remain constant.

A deep source of interest in these arguments lies in the help they provide with the second of the two tasks I described. They sharpen

* I have benefited from conversation with Rogers Albritton and Keith Donnellan at early stages of this project and with Willi Vossenkuhl at a later stage. I have also learned from discussion at Berkeley, Konstanz, Munich, Stanford, and Tucson, where I gave talks based on the paper.

¹ One such argument occurs in "Individualism and the Mental" *Midwest Studies*, iv (1979): 73–121. A second occurs in note 2 of the same article and in "Other Bodies," in Andrew Woodfield, ed., *Thought and Object* (New York: Oxford, 1982). Aspects of both arguments are further discussed in "Two Thought Experiments Reviewed," *Notre Dame Journal of Formal Logic*, xxiii, 3 (July 1982): 284–293. A third argument occurs in "Individualism and Psychology," *Philosophical Review*, xlv, 1 (January 1986): 3–45, and in "Cartesian Error and the Objectivity of Perception," forthcoming in a volume edited by John MacDowell and Philip Pettit, Oxford University Press. As is usually the case in philosophy, the important thing is less the conclusion itself than the routes that lead to it.

our conception of intellectual responsibility. This notion undergirds the proprietary concepts of dialectic, rationality, understanding, spirit, and rule-following that Plato, Aristotle, Kant, Hegel, and Wittgenstein, respectively, tried to explicate. The conditions under which people become responsible to intellectual norms are complicated. They involve considerations regarding etiology, minimum competence, and intention. I will not be concerned here with these conditions or with implications of the notions of agency and responsibility. I shall discuss rather the nature and source of the norms' authority and the role of those norms in constituting the notion of mind.

In this paper I propose a new argument for the view that mental-state and event kinds are nonindividualistically individuated. The argument uses the basic strategy of its predecessors, but provides a different perspective on the factors that underlie the conclusion.

The argument proceeds on two levels. At the more general level the key idea is that some necessarily true thoughts or statements can be doubted.² A few philosophers have rejected this idea. Perhaps some rationalists and certainly some positivists and possible-worlds semanticists have maintained that one cannot believe an impossibility. Most of these would add that one cannot doubt a necessary truth. I find these positions antecedently implausible, and I do not find their motivations compelling. (In some cases the motivations involve implicit commitment to individualism.) Moreover, these positions account poorly, in my view, for the epistemology of logic, mathematics, analyses of meaning, and a posteriori necessary truths. They do raise interesting issues. But I shall not confront them in detail here.³

I take it that most thoughts that are necessarily true are dubitable. The anti-individualistic argument I propose in section II is perhaps most clearly applicable to empirical necessary truths (e.g., that water is H₂O). I shall say little about these truths or about their role in our argument because they fit so straightforwardly into our scheme.⁴

² Actually, the argument need not rest on an assumption about necessity. What we need are general thoughts or statements so central to the correct identification of a type of thing, property, or event, that, under ordinary conditions, if the thought failed to apply to some given entity *x*, we would correctly and almost automatically refuse to count *x* an instance of the type. For expositional convenience, I shall write in terms of necessity, without further qualification. I shall use the term 'doubt' in the weak sense of 'withhold belief' (analogously, for cognates like 'dubitable')—except where the context indicates some different usage.

³ For a motivated development of the opposing semantical view, see Robert Stalnaker, *Inquiry* (Cambridge, Mass.: MIT Press, 1984). Positivist motivations will be discussed briefly in section I.

⁴ For a different reason, I shall not discuss logical and mathematical truths. Nearly all these are both necessary and dubitable. But the issues involved in relating them to my over-all view are complex and require special attention.

Thus the most general level of argument will largely stay backstage, with occasional bit appearances.

Most of the paper will proceed at a more specific level. At this level the key idea is that a particular subclass of necessary truths is dubitable. The subclass is made up of those necessary truths which are intuitively central to giving the meaning of an empirically applicable term or to providing an explicit general understanding of such a concept. I concentrate on this subclass because the ramifications of the more general argument against individualism are richest and least obvious for this case.

In section I I provide a rationale for the fact that this subclass of necessary truths can be doubted. In section II I lay out a thought experiment which anchors the argument against individualism. Section III is devoted to a brief discussion of alternative construals of the thought experiment. Section IV introduces a distinction, implicit in the argument, between cognitive value and conventional or idiolectic linguistic meaning. In section V I use this distinction to sketch an argument to the effect that even nonindividualistic "use-based" theories of mind (in senses to be specified) cannot be correct.

I

My aim in this section is to explain why necessary truths that intuitively give the meaning of an empirically applicable term (or purport to provide general understanding of such a concept) are dubitable.

As a preliminary, let us review some recent history. The most significant source of appeals to indubitability in this century has been a set of assumptions about meaning, understanding, and belief which emerged in positivist theories of knowledge. Some positivists counted logical and mathematical truths "true by convention" or "true purely in virtue of meaning" in order to protect empiricism against the charge that it could not account for knowledge in the deductive sciences. Propositions of logic and mathematics and truths that have criterial status (such as 'Sofas are pieces of furniture made or meant for sitting') were held to express not knowledge, but convention. In effect they were counted indubitable. This conclusion could be drawn in either of two ways. One could reason that doubting something requires understanding it; but criterial truths are true in virtue of meaning, so understanding them entails realizing that they are true; so it is impossible to doubt them. Alternatively, one could emphasize the allegedly degenerate nature of criterial truths: since such truths express only convention or fiat, anyone who understands them must realize that doubting them is pointless or meaningless; and doubt requires understanding; so doubt is impossible.

Neither argument would be given today. But their assumptions are

worth making explicit. They assume that doubt requires complete understanding of what is doubted; that criterial truths are true by fiat or convention, or meaning alone; and that understanding these truths requires realizing this (alleged) connection between their meaning and their truth.

All these assumptions are unacceptable. The argument that I shall develop in section II does not attribute incomplete understanding. But I shall develop reasons for doubting the first assumption shortly. (Cf. also the first *op. cit.*, note I.) W. V. Quine's work long ago threw the second into deserved disrepute.⁵ There is no explanatory use for claiming that the relevant truths are true purely in virtue of meaning, as opposed to partly in virtue of meaning and partly in virtue of (perhaps "obvious") ways the world is. That is, there is no reason to see them as degenerately true. The third assumption is clearly vulnerable to the same Quinean considerations.

Taken as self-conscious academic doctrine, positivist theories have not survived the century's sixth decade. But the general picture they offer of the relation between the meaning, understanding, and truth of "criterial" statements still informs much philosophical common sense. The first and third premises just discussed come closest to painting the general picture: the idea is that *in the cases at issue, understanding sufficient for carrying on responsible ratiocination—which is identified with understanding ordinary linguistic meaning—necessarily requires recognition of truth.*

Substantially the same picture has been painted by some self-styled Wittgensteinians. It has been assumed that belief entails full understanding, and claimed that a "criterion" of understanding is recognition of the truth of the relevant statements. The attempt of the positivists to provide a theoretical ground for this connection between understanding and belief (second assumption above) is dismissed as quixotic. Of course, we often do use apparent disbelief as

⁵ I do not back all elements in Quine's attack on the view that there are truths that are true purely in virtue of meaning (in particular, his empiricism and some of his attacks on the notion of meaning itself). But I think that his earliest criticisms of analyticity are sound: see "Carnap and Logical Truth" secs. I–V, X; "Truth by Convention," both in *Ways of Paradox* (New York: Random House, 1966). I take it that Quine's challenge to justify a disjoint distinction between "truths of fact" (possibly including, so far as this argument is concerned, necessary facts and facts known by reason) and *mere* "truths purely by virtue of meaning"—has gone unmet. This section supplements Quine's view by showing in some detail why certain "truths of meaning" are simultaneously "truths of fact", or better—why they are not degenerately true. Cf. also section IV below. Nothing I argue in this paper depends on acceptance of any particular theoretical conception of meaning. I merely make use of an untendentious intuitive notion of meaning to isolate a rough class of statements or thoughts.

ground for attributing failure of understanding, and even sometimes for reconstruing a person's words to avoid attributing disbelief. But these criteria are defeasible, as the cases we shall later discuss indicate. Sometimes the criterion is mechanically invoked in philosophical discussion as if it mirrored some absolute practice that we all recognize. Such invocation often serves to undergird the picture inherited from positivism.

There is a common enthymematic argument which also serves this picture. It goes as follows: The meaning of (say) 'sofa' is the same as the meaning of 'piece of furniture [of such and such a construction] meant or made for sitting'; but, if one understands both phrases, one cannot doubt that sofas are pieces of furniture . . . meant or made for sitting, since the expressions, meaning the same, are interchangeable; such a doubt would be a doubt that sofas are sofas; but *that* is not a doubt that anyone could be thought to have. So anyone who *appears* to be doubting that sofas are pieces of furniture . . . meant or made for sitting, cannot really be doing so.⁶ This reasoning again assumes that doubt requires full understanding; and it again assumes that understanding the relevant meaning entails recognizing truth. The reasoning, however, will provide a basis for exploring why these two assumptions are mistaken in cases relevant to the issues at hand, and thereby explain why necessarily true, "meaning-giving" statements involving many empirically applicable terms are dubitable.

My strategy will be to consider in some detail how synonymies are grounded and how we come to affirm them. I shall develop the view that the role of examples in the cases of the relevant synonymies guarantees that the relevant object-level statements or thoughts are dubitable.

What does *having the same meaning* involve in the cases of empirically applicable terms? Synonymies are grounded in practice: *the most competent speakers would use the two relevant expressions interchangeably*.

The notions of interchangeability and competence, which are central to this characterization, require discussion. First, interchange-

⁶ The argument from the substitutivity of synonyms, at least in embedded cases, received considerable discussion in the 1950s. The discussion was started by a different argument in Benson Mates, "Synonymy," in Leonard Linsky, ed., *Semantics and the Philosophy of Language* (Urbana: Illinois UP, 1952), p. 215. The points at issue in Mates's discussion are developed in some detail in my "Belief and Synonymy," this JOURNAL, LXXV, 3 (March 1978): 119–138. These issues have antecedents in earlier discussions of the "paradox of analysis." It should be emphasized that we are concerned with whether substitutivity of all ordinary synonyms is guaranteed necessarily or logically. In *most contexts*, of course, exchanges preserve the point and certainly the truth value of the report.

ability. Competent speakers are normally willing to exchange synonymous expressions in extensional, counterfactual, and many other contexts *salva veritate*; they use one expression as a means of identifying entities to which the other applies and as the primary and sufficient explanation of the other, apart from any particular perceptual context. This is a rough, minimal characterization; but it will serve.

One should not interpret “interchangeability” more strictly than our practices warrant. Abbreviative definition in formal systems is in most cases a misleading model. It is simply not the case that the only hold competent speakers have on one of a pair of synonymous expressions is always through the other. As any book on synonyms will show, there are many reasons for this. One example will suffice here. Sometimes there are syntactical differences between the relevant expressions which occasion differences in the ties they have to perceptual experiences. Thus ‘sofa’ is commonly applied, defeasibly but directly, to objects that look a certain way. The components of ‘pieces of furniture of such and such construction meant or made for sitting’ have independent perceptual and conceptual ties. The application of the longer expression to sofas often seems to involve a kind of computation.

Let us turn to “competence.” The language does not present a standard of competence independent of individuals’ activity. Minimal competence consists in conformity to the practice of others. “Greatest competence” consists in abilities to draw distinctions, to produce precisifications, to use numerous linguistic resources, to offer counterexamples to proposed equivalences—that elicit the reflective agreement of other competent speakers. We may imagine a vast, ragged network of interdependence, established by patterns of deference which lead back to people who would elicit the assent of others. (Of course, we idealize from this network; a person’s degree of competence may vary over time and with the case at hand, and may develop or regress.) To put it crudely, a person counts as among the *most competent* if he or she would be *persuasive* to other competent speakers in the use and explication of the language. The point about persuasion is fundamental. I shall develop it by considering the dialectic that commonly leads to statements that explicate meaning.⁷

⁷ My characterization of dialectic will overschematize enormously. I shall ignore the variety of types of synonyms (some of which would require different treatment) in order to concentrate on certain types that will further my purposes most directly. And I ignore many cases of explaining meaning that do not aim at synonymy in any sense. Such cases are rife, and important to our view. But I must omit discussing them here.

It is obvious that few, even among the competent, can quickly and accurately hit upon meaning-giving characterizations for ordinary terms. Usually they need thought or discussion to improve their initial attempts. Beliefs about *what Xs are* typically contain minor errors, often because of over- or undergeneralization from common examples of Xs. This fact already motivates some distinction between the sort of competence necessary to engage in ratiocination involving use of the term or notion, and the sort necessary to have an accurate understanding of the conventional meaning of the term. For a person may, in the course of the dialectic, have thoughts involving the concept X , and yet have beliefs about what Xs are which are incompatible with the conventional normative characterization or with the characterization he or she would regard on reflection as correct. Thus I may have numerous ordinary chair beliefs (that that is a chair), yet believe incorrectly that chairs must have legs. Such a belief betrays an incomplete understanding of the conventional meaning-giving characterization of chairs, but does nothing to exempt me from responsibility to communal norms of evaluation when I have or express my chair beliefs.

The dialectic attempts to arrive at what might be called *normative characterizations*. These are statements about *what Xs are* that purport to give basic, “essential,” and necessarily true information about Xs. They are used as guides to certifying the identity of entities: something that is cited as an X but does not fulfill the condition laid down by the normative characterization will not normally be counted an X. A subclass of normative characterizations not only purport to state facts that set norms for identification; they also provide linguistic meaning—set a norm for conventional linguistic understanding (‘A knife is an instrument consisting of a thin blade with an edge for cutting, fastened to a handle’; ‘To walk is to move on foot at a natural unhurried gait’; ‘A baby is a very young child or very young higher animal’). Not all normative characterizations provide meaning in this sense (‘Water is H₂O’). But my primary interest is in cases where they do.

Meaning-giving characterizations, for ordinary terms, are usually arrived at through reflection on *archetypical applications*. These are perceptually backed, indexically mediated applications (or imagined projections from these) to “normal” or “good” examples (‘That’s a knife’; ‘That woman is walking’). The dialectic consists of an attempt to find a fit between such examples and the characterizations that are dominant in selecting them. To provide meaning, a proposed normative characterization must accord with archetypical applications and must treat the characterizations that competent

users actually give, as at least approximations to the norm. The conventional linguistic meaning of a term has been correctly specified when, under these restrictions, the most competent speakers have reached equilibrium on a characterization.

I have pointed up the central role of persuasion in the dialectic. To understand the dialectic one must understand the nature of the persuasion. Sometimes we see the most competent speakers as creating new uses through impressive employment of verbal resources. In such cases, persuasion in the strict sense need not be involved. Imitation may be based primarily on the attractiveness of the style of speech, the power or status of the speaker, or the impressionability of the hearer. But often when someone is seen as more competent it is because he or she is persuasive on matters about which there are objective rights and wrongs and on which substantive reasons have a bearing. The agreement reached is not to a decree, nor is it merely the result of practical reasoning about how to adjust our usage for smoothest communication. The agreement concerns the proper ordering of applications of a term which we have already made or are disposed to make and, ultimately, the correct characterization of the examples that those applications pick out. In the course of the dialectic, we *stand corrected*: we recognize ourselves as convicted of *mistakes*, not merely infelicitous strategies for communication. We come to know something that characterizes empirical entities and sets standards for characterizations to which we regard *ourselves* as antecedently committed. Thus the most competent speakers are pre-eminent not merely because they impress the impressionable. Their influence is based on persuasion that is subject to dispute and cognitive checks.⁸

These disputes usually concern two matters at once. One is how

⁸ Of course, there is nothing about the signs or about the most competent speakers *per se* that makes it right for others to use the signs "interchangeably." This point has encouraged the claim that some people's using expressions interchangeably *does not ever* signal that it is cognitively correct for other people to do so. Cases of linguistic correction are assimilated to cases of strategic accommodation for achieving a common communicative purpose. This position is part of Donald Davidson's theory of linguistic interpretation. Cf. his *Inquiries into Truth and Interpretation* (New York: Oxford, 1984), e.g. "Radical Interpretation." It is also implicit in Quine's views about translation. Cf., for example, *Theories and Things* (Cambridge, Mass: Harvard, 1981), pp. 49/50, and the role of stimulus meaning in the translation theory of *Word and Object* (Cambridge, Mass.: MIT Press, 1960). It seems to me that this position is mistaken. It gives no plausible account of the fact that people frequently do stand corrected for not using (or for using) expressions interchangeably, for making nonmetalinguistic, cognitive mistakes expressed in their own terms. And it has certainly not explained why they ought not stand corrected, or why they are immune from that sort of error. For further discussion, cf. "Individualism and the Mental," *op. cit.*, pp. 89–103. (See also note 10 below.)

correctly to characterize the relevant entities: whether all chairs have legs or must have legs. The other is how to state the meaning of the term as such: whether according to (by) definition chairs have legs, whether according to the standard accepted meaning of 'chair', chairs have legs. The second question is higher-order: either it is explicitly about words, or it otherwise contains an extra intensional context ("according to"). The first question is formulated more directly about extra-linguistic entities (chairs). Later I will indicate that these questions are not equivalent. Their answers can come apart. But for now, what is important is that, even where there is no practical point in distinguishing the questions, the second is typically pursued by trying to answer the first. Questions of meaning are pursued by attempting to arrive at factually correct characterizations of empirically accessible entities, the examples.

This point rests on two features of the dialectic which I have already noted. One is the central role that *examples* play in arriving at meaning-giving characterizations for ordinary terms ('Here is a chair and it lacks legs', 'That man is pulling'). The dialectic aims at capturing archetypical applications in a way that sharpens explications that competent speakers naturally give. In being responsible for correctly characterizing core examples, explications of meaning are sensitive to empirically available facts. Examples are ineliminable from our procedures of meaning-giving. For where expressions are regularly and nondelusively applied to perceivable extralinguistic reality, examples necessarily tend to be created and legitimated. (Cf. note 16 below.) If some terms were not so applied, no terms would have objective, empirical meaning.⁹

The second notable feature of the dialectic is that, as the participants work toward an expression of communal meaning, they typically do not discuss the matter as outsiders. Usually, all participants begin the discussion without being able to give a precisely correct normative characterization; all or most would make minor errors in attempting to do so. But this does not entail that any lack object-level thoughts expressible, by the rest of us, with the term whose meaning is in question. If it did, most people would have few if any object-level thoughts so expressible—an absurdity. Participants commonly regard their object-level thoughts (thoughts about, say, chairs) as undergoing correction in the course of the inquiry. They stand corrected on substantive matters. This is exactly what one would expect in view of the role of examples in the discussion.¹⁰

⁹ Cf. my "Belief *De Re*" this JOURNAL, LXXIV, 6 (June 1977): 338–362, sec. II.

¹⁰ I have conducted the discussion on what seems to me to be the natural and correct view that an individual's terms' meanings sometimes are dependent on

We may now return to the argument for indubitability from the substitutivity of synonyms (seventh paragraph of this section, page 701). Suppose that a given normative characterization is both true and expressive of communal meaning. What would make it indubitable? The answer is alleged to derive from the synonymy of the relevant expressions. The expressions' synonymy is grounded—in-*deed* in a sense consists in—their being treated as interchangeable by the most competent speakers.

Now it is true, as noted, that in ordinary cases synonymy provides a rough norm for understanding and belief: apparent rejection of the relevant statements derived from synonymies is usually a sign of mistake resulting from incomplete understanding. But the nature of authority developed and manifested in the course of the dialectic indicates that criterial statements derived from synonymies are (typically) not indubitable. Authority derives in part from an ability to answer doubts—to persuade oneself and others that erstwhile beliefs do not accord with facts about core examples or with usage to which the participants hold themselves responsible.

The norms for *understanding* provided by synonymies take their force from actual or potential agreement partly stemming from persuasion. But the *truth* of normative characterizations does not rest primarily on agreement. In the first place, agreement is elicited through a process that involves reasoning and citations of extralinguistic fact. In the second place, even when equilibrium is reached—even when one arrives at a complete expression of and agreement on relevant norms, at complete understanding of linguistic meaning—there is no transcendental guarantee that people cannot agree in making mistakes. The authority of the “most competent” rests on an ability to turn back challenges. But usually there is no method for demonstrating that every possible relevant challenge has been answered. The role of examples in the dialectic makes such assurance

commitment to standards for understanding that others can hold us to. (Cf. note 8.) But the argument for the dubitability of “meaning characterizations” that follows still works if one assumes that only the individual can criticize his or her own characterizations. For the dialectic of persuasion can be seen as carried out purely by oneself. Each characterization one gives oneself could be corrected or come to be doubted because of examples to which one is antecedently committed. The rest of the paper can be preserved, under fairly simple qualifications and transpositions, if one maintains this antisocial picture of meaning. But it seems to me that there is no good reason to maintain it. In fact, if an individual can correct him- or herself by reference to publicly available examples, it is hard to see why the individual should not stand corrected by others using similar considerations. Although each person participating in the dialectic is reflecting on his or her *own language*, there is a willingness to take criticism from others.

impossible. I conclude that our reflective conception of synonymy not only does not support the indubitability of meaning-giving normative characterizations for many empirically applicable terms; it actually underwrites and makes inevitable their dubitability.

II

The dubitability of meaning-giving normative characterizations can be converted into a demonstration that social practices are not the only or ultimate nonindividualistic factor in individuating mental states and events. I have elsewhere argued for this view on other grounds. Some mental states (for example, some perceptual states) depend for their identity on the nature of the physical environment, in complete independence of social practices. What I want to show here is that, even where social practices are deeply involved in individuating mental states, they are often not the final arbiter. This is because the sort of agreement that fixes a communal meaning and norms for understanding is itself, in principle, open to challenge. The argument that follows articulates this fact.

We begin by imagining a person *A* in our community who has a normal mastery of English. *A*'s early instruction in the use of 'sofa' is mostly ostensive, though he picks up the normal truisms. *A* can use the term reliably. At some point, however, *A* doubts the truisms and hypothesizes that sofas function not as furnishings to be sat on, but as works of art or religious artifacts. He believes that the usual remarks about the function of sofas conceal, or represent a delusion about, an entirely different practice. *A* admits that some sofas have been sat upon, but thinks that most sofas would collapse under any considerable weight and denies that sitting is what sofas are pre-eminently *for*. *A* may attack the veridicality of many of our memories of sofas being sat upon, on the grounds that the memories are products of the delusion.

A is willing to test his hypothesis empirically, and the sociological tests he proposes are reasonable. *A* also offers to demonstrate by experiment how the delusive memories are produced. He is sophisticated, and the tests would require elaborate controls. We can even imagine that the theory is developed so as to be compatible with all past experience that might be thought to have falsified his theory. Thus a normal but sophisticated conception of confirmation accompanies *A*'s unusual theory. We may imagine that if we were to carry out his proposed experiments, *A* would come to admit that his theory is mistaken.

As a second step, imagine a person *B* (or *A* in nonactual circumstances) who is, for all intents and purposes, physically identical to *A*. He has the same physical dispositions, receives substantially the same

physical stimulations, produces the same motions, utters the same sounds. Like *A*, *B* hears, though seldom, word forms that are counterparts to the truisms that *A* hears. But in *B*'s situation, these word forms are not taken as truisms; they are contextually appropriate remarks that do not purport to convey a general meaning. (They could be lies or jokes, but it is more natural to take them as not-completely-general contingent truths.) The objects that *B* is confronted with are objects that look like sofas, but are, and are widely known to be, works of art or religious artifacts sold in showrooms and displayed in people's houses. Many of these objects would collapse under a person's weight. There are no sofas in *B*'s situation, and the word form 'sofa' does not mean *sofa*. Call the relevant objects "safos." *B* assumes that most people would take these objects to function primarily as seats and that the remarks he hears are communally accepted truisms. But, like *A*, *B* develops doubts. At least by the time *B* expresses his skepticism and his theory, he is correctly doubting that safos function as furniture to be sat upon. Thus *B*'s thoughts differ from *A*'s.

The self-proclaimed skeptics face different responses when they express their views. *A*'s view is resisted. Although *B* thinks he is opposing received opinion, his claim about what safos are is taken as a matter of course. *A* mistakenly thinks that sofas do not function primarily to be sat upon. *B*'s counterpart thoughts do not involve the notion of sofa and could not correctly be ascribed with 'sofa' in oblique position. He correctly thinks that safos are works of art or religious artifacts. There are numerous correlative differences in the thoughts of the two people which I shall not spell out. *A* is right about the sociolinguistic practice of his comrades, but wrong in the relevant sofa thoughts. *B* is right about safos and wrong about the surrounding sociolinguistic practice (though his usage and beliefs may accord with it perfectly).

The conclusion is that *A* and *B* are physically identical until the time when they express their views. But they have different mental states and events. *A* has numerous mental events involving the notion of sofa. *B*'s skepticism does not involve thinking of anything as a sofa. I do not assume that what I have said about the case in non-propositional-attitude terms *entails* that *A* and *B* have different thoughts. I think, however, that it is overwhelmingly plausible that they do and, more importantly, that there is no general objection to the natural view that they do.

The arguments of "Individualism and the Mental" and "Other Bodies" (*op. cit.*) ascribe incomplete linguistic understanding and ignorance of expert knowledge (respectively) to the relevant protag-

onists. By contrast, *A* may be a sophisticate. He need not lack linguistic understanding or be unapprised of expert or common opinion. The present argument features not incomplete understanding or ignorance of specialized knowledge, but nonstandard theory.

This orientation makes the argument extremely comprehensive in its application and very resistant to objection. Nearly anything can be the topic of nonstandard theorizing. Similar thought experiments apply to knives, clothing, rope, pottery, wheels, boats, tables, watches, houses. Both technical and everyday natural-kind notions clearly fall within the domain of the argument. (Science has generated nonstandard theory with respect to most such notions.) Concepts of other ordinary objects and stuffs, which are not natural kinds, are equally good examples: earth, air, fire, mountains, rivers, bread, food, dung. Notions associated with common verbs are also subject to our argument. What it is to eat, to talk, to sing, to own, to walk, to sleep, to fight, to hunt, to have intercourse can each be subjected to strange theory. The last already has been.

I claim that the argument can be adapted to any substantive notion that applies to physical objects, events, stuffs, properties. The point I wish to press is that the notions may be as ordinary and as observational as one likes. They may be the basic tools of common sense and child rearing.

The thought experiments instantiate three general points. The first is that propositional mental-state and event kinds are individuated by reference to intentional notions. The second is that there are certain relations between an individual and the environment that are necessary to the individual's having certain intentional notions. If one conceives the environment—or the individual's relations to the environment—as varying in certain ways, one must conceive of the individual's intentional notions as varying. The third point is that there is possible slack between the relevant environmental facts and relations to the environment, on one hand, and what the individual knows and can discriminate, on the other. The sort of slack relevant to the present argument is the dubitability of our beliefs about individuation. Clearly this third point applies not only to the special case of meaning-giving normative characterizations, but also at the more general level mentioned in the introduction.

These points schematize the thought experiments. The third point functions in the argument in two ways. It allows ascription of error, ignorance, incomplete understanding, or nonstandard theory in the first stage of the thought experiment. And it allows us to conceive of the relevant environmental facts or the individual's environmental relations as varying without varying the individual's individualistically

described nonintentional discriminatory powers. The second point indicates that these variations prevent the ascription of error or ignorance which was possible in the first stage of the thought experiment. Under certain conditions one cannot have acquired certain thoughts. And the first point indicates that mental-state and event kinds have varied between the two stages because intentional notions have varied.

The iconoclastic theory will often require ingenuity. But we can draw upon traditional skepticism to elaborate cases. To discourage attempts to show by transcendental arguments that the protagonist's doubt is incoherent or impossible, I have made the skepticism testable. By traditional standards, the sort we require is quite modest. The protagonist need not believe it; he or she may simply consider it. It need not be possibly true. It is enough that it be epistemically possible in the weak sense that it is thinkable.

III

I now want to defend the argument briefly against some possible rejoinders.¹¹ Since we shall mainly discuss the first step, let us look briefly at the second in order to lay it aside. Consider a parallel case. Suppose that someone, who can recognize sofas as well as anyone, told us, "These things, sofas, are not art objects, contrary to what you think; they are pieces of furniture commonly made, intended, and used to be sat upon." We might regard the person as having delusions about our views. But, barring further evidence, we would regard the opinions about sofas as unproblematic and would concentrate on the sociological perversity. I believe that the point is fairly obvious. Rather than belabor it, I turn to the first step.

That step presupposes that widely accepted thoughts about what Xs are are not indubitable: that it is possible to be uncertain about them. I believe that a claim that A literally cannot doubt what he appears to doubt is obviously implausible. But I will consider some ways of making the claim.

It would be unacceptably superficial to gloss the case by claiming that A is refusing to speak our language. We have no difficulty understanding that he is raising questions about what sofas really are. The proposals for deciding the question are exactly what we would expect, given a literal interpretation of what he says.

One could claim that, whatever he says, A has no thoughts literally attributable to him by using 'sofa' in oblique position, but only metalinguistic thoughts about the word 'sofa'. The metalinguistic ma-

¹¹ I discuss related objections in a different context in "Individualism and the Mental," *op. cit.*, sec. III.

never is even less plausible as applied to the present argument than it is as applied to the argument of “Individualism and the Mental.” *A*’s doubts focus on empirical facts and are to be tested by empirical methods. The metalinguistic maneuver is appropriate when the speaker has no minimal competence. But there need not be any failure or incompleteness of understanding on *A*’s part, much less a lack of minimal competence.

A related objection would hold that *A* thinks only that *what most people think of as sofas* are works of art or religious artifacts. The word ‘sofa’ should, on this view, be reinterpreted in terms of the italicized phrase. Of course, we must ask how to interpret ‘sofa’ in the italicized phrase. Answering this question will lead to complications that favor our view. But there is a simple reply to the suggestion. We may assume that *A* would say that what most people think of as sofas are sofas. If he makes the distinction between meta- and object-level, we cannot collapse it in interpreting him. What he questions, quite explicitly we may imagine, is not whether these things—or what people think of as sofas—are sofas, but *whether sofas are what people think they are*.

One might hold that the protagonist is to be attributed a “reduced” notion of sofa, like one an anthropologist might employ on coming into a society that uses a term for objects that he or she can recognize, but whose use he or she has not yet determined. The key to understanding this objection lies in being scrupulous with the term ‘reduced notion’. *A* is not an outsider; he has fully learned communal usage. *A* does use ‘sofa’ in a “reduced” way in that he prescind from assuming as true the “truisms” that explicate the term’s communal meaning. But it does not follow that ‘sofa’ should be re- or de-interpreted in literal attributions of *A*’s propositional attitudes. When one tries to be more specific about what the “reduced” notion is, difficulties emerge.

Suppose, for example, that the “reduced” notion were tied to perceptual aspects of sofas. Sofas in the “reduced” sense are just things that look like those things (where sofas are indicated). Clearly, this notion need not be what *A* utilizes. He may be unwilling to commit himself to how all sofas look or to there being no counterfeits. He may rely on others in determining whether some samples are sofas. In the interests of brevity, I shall not pursue alternative “reduced” interpretations. I think it is always possible to show, compatibly with the thought experiments, that revised construals of ‘sofa’ need not capture the protagonist’s way of thinking.¹²

¹² This sort of dialectic can be pursued at considerable length. Another possibility for a “reduced” sense of ‘sofa’ in attributions to *A* is *thing of a kind relevant to*

One might claim that two ways of describing the dispute must be “equivalent.” *A* could be said to doubt whether sofas are really such and such, agreeing that *these* are sofas. Or *A* could be said to agree that sofas must be such and such, and doubt whether these are sofas. I think that any claim of this sort is quite mistaken. There may be such cases of descriptive equivalence, but I see no reason or plausibility for the claim that our attributions to *A* are necessarily descriptively equivalent with different ones. In view of the way *A* expresses his reasoning—agreeing that these are sofas, acknowledging the communal agreement on the defining condition, but questioning that condition by attacking its empirical presuppositions—it is difficult to imagine an equally good redescription of the case that would have him accepting the defining condition but doubting alleged exemplifications of it. The reasons that *A* gives support one description rather than the other (cf. note 16).

I think that it must be admitted that the various maneuvers for blocking the first step in the thought experiments, taken as claims about our common conceptions, are not very plausible. There is simply nothing in our ordinary practices that precludes our taking *A* as literally entertaining the doubts I have ascribed. Whatever impetus there is behind the objections derives not from antecedent practice, but from the feeling that there must be some way of resisting the ascriptions. This feeling derives from habits that stem from background philosophical doctrine. The main sources of philosophical opposition to the thought experiment’s first step are the claims about indubitability discussed in section I. These views provide no

understanding what those things are (where some sofas are indicated). This sort of suggestion again (at best) confuses reference fixing and way of thinking. If one takes the proposal literally, the italicized phrase expresses a way of thinking of things which is indexical; it will shift its application or referent from occasion to occasion. ‘Sofa’ in our attributions to *A* is not indexical. (Cf. “Other Bodies,” *op. cit.*, sec. II.) We can build into the thought experiments behavioral tests to help establish the point. For example, it could emerge under questioning that *A* would regard himself as subject to correction if he were to allow the referent of ‘sofa’ to shift with context. One could perhaps modify the italicized phrase to *thing of a kind relevant to understanding what I (or we) usually refer to when I have this sort of experience (or, alternatively, when I use ‘sofa’)*. But now it is clear that no one but a philosopher would think of sofas in that complex, meta-level way. A complex analysis or theory of reference should not be conflated with the way *A* thinks of sofas. The counterfeit is recognized by Frege’s test. Such analyses, even if true, are informative. *A* might have to reason to decide whether they are true. So the analyses need not yield the way *A* thinks of sofas. *A* might even lack notions that occur in the analyses without lacking his notion of sofa. No version of the reduced-sense line of objection that I know of is acceptable.

good reason for regarding our common cognitive practices as in need of revision.¹³

IV

The story inherited from positivism holds that understanding sufficient for responsible ratiocination—which is identified with understanding ordinary linguistic meaning—necessarily requires recognizing the truth of normative statements derived from synonymies. There are two things wrong with the story: the two sorts of understanding should not be identified; and neither sort necessarily requires belief in the relevant statements.

Consideration of the dialectic for arriving at meaning-giving characterizations (section 1) indicates different types or levels of understanding. There is, first, understanding sufficient to engage in responsible ratiocination with a notion expressed by a term.¹⁴ We presuppose that a person has this sort of understanding when we use the term in oblique position literally to specify attitudes—when we treat the person as responsible to cognitive criticism by reference to the relevant normative characterizations. Such responsibility is usually incurred before normal competence is achieved. But even normal competence—the sort of understanding that one can reasonably expect of a person who is a member of the community—typically throws up mistaken explications that meaning-giving normative characterizations must correct. Finally, there is full understanding of the meaning-giving normative characterization—the ideal, articulable mastery achieved by the most competent speakers or by an individual who fully understands his or her own idiolectic usage. The critical point is that neither understanding sufficient for ratiocination nor normal competence requires full understanding.

¹³ There are, of course, the indeterminacy theses of Quine and Davidson and other views that challenge the objectivity and cognitive status of mentalistic attributions. For the indeterminacy theses, see Quine, *Word and Object*, and Davidson, *Inquiries into Truth and Interpretation*, *op. cit.* These views are somewhat orthogonal to our primary theses here, since they bear more directly on the status of our conceptions than on their character. (It is doubtful, incidentally, that Davidson's view is individualistic.) However, I find unpersuasive all current doctrines that hold that mentalistic discourse is noncognitive or even (as a matter of principle) less cognitive or "factual" than physicalistic discourse. Some of what I say here and in "Individualism and Psychology," *op. cit.*, is material for supporting this view. In this latter article I argue that nonindividualistic, mentalistic methods of individuation are legitimate for and present in scientific explanation.

¹⁴ Animal ratiocination is ontogenetically prior to any use or understanding of linguistic symbols. In omitting animal thought from the present discussion I am not overlooking or denigrating it. I think, however, that it does not involve certain sorts of intellectual responsibility, the norms for which are our present concern.

Moreover, understanding the meaning-giving characterization for a term does not necessitate acknowledging it as true. It is sufficient to be able to apply the term correctly and to give approximations to standard characterizations that may be adjusted under reflection, criticism, or new information. There is a potential gap between even the best understanding of accepted usage and belief. The consensus of the most competent speakers can be challenged. Usually such challenges stem from incomplete understanding. But, as our arguments in sections I and II indicate, they need not. One may always ask whether the most competent speakers' characterizations of examples (or one's own best characterizations) are *correct*, and whether all examples usually counted archetypical *should* be so counted.

The impetus to accept these conclusions, and their anti-individualistic corollaries, already lies implicit, I think, in Quine's first point against positivism. I have in mind his point that there is no separating truths of meaning or truths of logic (or criterial truths, or truths of reason, or necessary truths) from truths of fact. (Cf. note 5.) In stating a truth of meaning (however one construes the notion), one is not stating a degenerate truth. To put this crudely: in explicating one's "meanings," one is equally stating nondegenerate truths—"facts." So giving a true explication is not separable from getting the facts right. It is a short step from this point to the observation that truths of meaning are dubitable.

Of course, the key idea is much older. Our cases develop a theme from the Socratic dialogues: Thought can correct meaning. Although dialectic typically serves to correct wayward belief and incomplete understanding, there remains the possibility that the discussion will overflow the boundaries of the established norms. A conventional norm is nearly always subject to the most general evaluative question: "Is it true?" If new empirical facts or new insights are imported into the discussion, the background assumptions of normative characterizations may be undermined, and the characterizations themselves may be shown to be mistaken.

The fact that thought can correct meaning has significant consequences for a philosophy of mind. A consequence I shall develop is a distinction between cognitive value and conventional (or idiolectic) linguistic meaning. There are two closely related arguments for this distinction. One derives from the possibility of doubt. The other issues from reflection on the epistemic possibility of radical theoretical change.

We must be able to specify a person's thoughts in situations where meaning-giving normative characterizations come under fire. Since it is possible to doubt such characterizations (and their logical con-

sequences), it is possible to find them informative. Doubts of such characterizations ('Sofas are pieces of furniture . . . meant for sitting') are supported by publicly recognized means different from those which would support doubt of the corresponding identity judgment (say, the judgment that sofas are sofas). In short, such characterizations have different cognitive values from those of the corresponding identity judgments. So in interpreting a specification of a belief that sofas are pieces of furniture . . . meant for sitting, one must assign different cognitive values or units of potential information—to the conventionally synonymous phrases ('sofa' and 'piece of furniture . . . meant for sitting') as they occur in such specifications. Thus cognitive value and conventional meaning should be distinguished.¹⁵

The same conclusion can be derived from considering possible linguistic or conceptual change pursuant on disagreement over meaning-giving normative characterizations. For example, it is possible to doubt that sofas are all and only pieces of furniture of a certain construction meant or made for sitting. If the doubt were to prove well founded, the conventional meaning of 'sofa' would be forced to change. But despite the change, it might remain appropriate, before and after the change, to attribute propositional attitudes involving the notion of sofa. Both before and after, *A* and his opponents would agree that *these are sofas*. Before and after, they would be characterized as having disagreed over whether all and only *sofas* are furnishings of a certain structure made or meant for sitting.

¹⁵ The argument, of course, is a variant of Frege's for distinguishing senses from one another and from denotation. The argument should not be construed in a narrowly linguistic manner. Its aim is to illumine the cognitive phenomena underlying the language about belief. Cf. "On Sense and Reference" in P. T. Geach and Max Black, eds., *Translations from the Philosophical Writings of Gottlob Frege* (Oxford: Basil Blackwell, 1966). Although there are differences between my view of "cognitive value" and Frege's, the conclusion of my argument as so far stated is also Fregean in spirit. In Frege's view, what sense an expression expressed was fixed by a deep rationale underlying the expression's use and understanding—a rationale that might not have been understood by anyone. Cf. my "Frege on Extensions of Concepts: From 1884 to 1903," *Philosophical Review*, xciii, 1 (January 1984): 3–34, esp. pp. 3–12, 30–34, and "Frege on Sense and Conventional Meaning," forthcoming in David Bell, ed., *The Analytical Tradition in Contemporary Philosophy* (New York: Cambridge, 1986). Other, lesser reasons for distinguishing meaning and cognitive value appear in my "Sinning against Frege," *Philosophical Review*, lxxxviii, 3 (July 1979): 398–432. I might add here that, although I have been speaking of cognitive value and conventional linguistic meaning as if they were different *entities*, nothing so far said commits me to regarding them in this way. One might take the expressions as labels for different methods of interpretation. The pressure to speak of entities, which is I think more insistent in the case of cognitive value, derives from providing a systematic semantical interpretation of the relevant discourse.

This situation bears some analogy to cases of theoretical change in science. Dalton and his predecessors *defined* 'atom' (and its translations) in terms of indivisibility. Major theoretical changes intervened. The definition was discarded. Despite the change, we want to say, Dalton wrongly thought that *atoms* were indivisible: despite his erroneous definition, he had the "concept" of atom (not merely the referent of 'atom'). I think that this sort of attribution is defensible in a wide variety of central examples of scientific developments, even in the light of Kuhnian insights into scientific revolutions. Although I cannot develop the point here, the distinction between theoretical meaning and cognitive value is in many ways analogous to that between conventional meaning and cognitive value.

It would be a mistake, however, to assimilate common-sense notions to a theoretical paradigm. Although meaning-giving characterizations from ordinary terms or notions are vulnerable to theoretical change, they differ from theoretical definitions of terms whose original home is a systematic theory, not only in that they are more stable and in that sense less vulnerable to theoretical criticism. They also differ in the means by which they are known and checked and in the ways in which they are vulnerable. (This is one durable element in the problematic observational/theoretical distinction.) The dialectic that we have been exploring seeks to derive characterizations from reflection on perceived examples picked out by common indexical usage. By contrast, the natural sciences, whose methodology we best understand, do not expect to reach their normative characterizations through simple reflection on usage or common perceptual experiences. Theoretical terms are not indexically applied to perceived objects. ('This is a quark; see the quark'—not kosher.) Ordinary words (or notions) for ordinary entities are given their life and meaning through such applications. Not only meaning-giving characterizations for these words but also doubts about these characterizations must accord with the bulk of ordinary, indexically aided, perception-based uses of these words—uses which yield the archetypical examples. This fact is part of what underlies the relative security of common-sense notions for observables.¹⁶

¹⁶ It is, for example, appropriate (as distinguished from merely possible) to doubt the existence of entities categorized by an ordinary notion, like *sofa*, *only* by appealing to doubts about the veridicality (referentiality) of typical perceptual experiences of entities so categorized. "Ordinary" notions are those, unlike *phlogiston* or *witch*, which in an ordinary sense and in ordinary circumstances are thought to apply to entities that are perceived and are not taught or applied by reference to a theory that depends heavily on assumptions about events or powers that are not

On the other hand, the fact that ordinary conceptions are more familiar and less embedded in a systematic theory does not indicate that normative characterizations involving them are indubitable or immune from theoretical intrusion. Criticism can emerge unforeseen from almost any quarter. This open-endedness is a matter of principle. It motivates the distinction between cognitive value and conventional linguistic meaning.

There was never any obvious reason why the two notions should coincide. They are responsible to different paradigms, and their explanatory purposes are distinct. Cognitive value is fitted to explicating possible differences in attitudes—cognitive perspective. Such differences are invoked to explain action and epistemic inquiry. Frege's test for differential dubitability, when accompanied with requirements that doubt be supportable by publicly recognized methods, is a defeasible but profoundly valuable tool in individuating cognitive values. Attributions of continuity through changes of conventional or theoretical meaning provide another touchstone for cognitive value. Conventional meaning, by contrast, is fitted to describing reflective agreement on the means of conveying, in short order, accepted usage. Dialectic ending in reflective equilibrium among competent speakers is a defeasible but reliable tool for individuating conventional meanings.

The two notions differ not only in their explanatory or descriptive purposes but in the ways they provide a basis for applying norms in our linguistic and cognitive practice. Conventional meaning provides norms for ideal competence by reference to which a person's usage and beliefs may be corrected. The normative function of cognitive value is more complex. On one hand, the notion marks the minimum competence necessary for ratiocination specifiable with the term. To be attributed the "concept" of sofa—and to have one's attitudes literally and correctly specified with 'sofa' in oblique occurrence—, one must be able to use the term 'sofa' sufficiently well to be attributed mistakes and true beliefs about what sofas are. On the other hand, the notion of cognitive value marks the openness of established normative characterizations to correction by reference to nonconventional theory. One may use one's concept of sofa to criticize conventional meaning.

commonly perceived. As long as archetypical perceptual applications ('That's a sofa') are not criticized wholesale, doubts about normative characterizations of such entities categorized by ordinary notions can be quite radical without undermining the basic usefulness of the categorizing notion.

Earlier in this section we distinguished different levels of understanding, culminating in ideal, articulable mastery of a conventional normative characterization. Is there a further sort of ideal understanding associated with cognitive value? There is no simple answer. In order to think with a given cognitive value (or concept), one need only attain the minimal level of understanding mentioned earlier. Usually, the best understanding one can achieve of a cognitive value is that offered by accepted normative characterizations and whatever background information accompanies them. Thus full understanding of cognitive value is normally not distinct from ideal understanding of ordinary usage and meaning. (In such cases, the cognitive value expressed by a term is, however, still individuated differently from its linguistic meaning.) When thought does correct meaning, one may achieve a revised understanding of cognitive value based on theoretical realization that goes beyond ordinary usage and meaning. Understanding of this sort bears comparison with the sort of ultimate insight, championed by the rationalist tradition, that was regarded as concomitant with deep foundational knowledge. Perhaps foundations and ultimacy are not to be expected. But our cognitive commitments and potential go beyond the boundaries set by conventional (or idiolectic) linguistic meaning.

v

A central aim, some have said the central aim, of the analytic tradition has been to give an account of thought and meaning by reference to some pattern of activity by language-users or thinkers. There are various positions on the relative priority of thought and meaning. One approach, distinctive of this century, is to explicate thought in terms of linguistic meaning and thinking in terms of using symbols with such meaning. A more traditional view is to explicate linguistic meaning in terms of thought. A third view, which I regard as correct, is that the two notions are interwoven in complex ways which render it impossible fully to analyze one in terms of the other. But with all these approaches, a syndrome of philosophizing characteristic of the analytic tradition has been to explicate both mind and meaning in terms of some pattern of activity: behavior, use, functional role, verification or confirmation procedure, computational role, subjective probability metric, teaching or interpretation procedure, illocutionary or perlocutionary act potential—and so on.

Surely, some pattern of actual physical, linguistic, or mental activity is essential to investing any given act with linguistic meaning or cognitive value. Such activity is even, I think, necessary to making meaning and cognitive value possible. Several of the use-based accounts have illumined and deepened this point. But, when these

accounts have posed as reductive explanations, they have been notably unilluminating. Most have been extremely programmatic. And where they have become more specific, they have become clearly implausible. Ironically, these implausibilities have often been made the basis of skeptical attack on intentional notions. I think the proper object of blame is the assumption that a reduction must be possible.

One significant obstacle to reduction is the complexity forced on reductive accounts by the failure of individualism. Both the conventional linguistic meaning and the cognitive value expressed by an individual may vary with the individual's environment, even as the individual's activities, individualistically and nonintentionally specified, are held constant.

Despite these similarities, conventional linguistic meaning and cognitive value relate to use, or nearly any actual pattern of activity, in significantly different ways. I shall first sketch an application of the point to linguistic use, and then generalize it.

Schematically, our thesis is that in certain senses, "use," socially conceived, necessarily "fixes" conventional linguistic meaning; but "use" (individualistically or socially conceived) does not necessarily "fix" cognitive value. This general schema comprehends a variety of specific instantiations. I shall try to articulate two.

The most straightforward specification of the thesis proceeds from an intentional conception of use and an epistemic conception of fixing. Ordinary, thorough, unimpeded, rational reflection by members of a community on the totality of their actual communal practices—all patterns of assertion, deference, teaching, persuasion—would necessarily suffice to yield a characterization that in fact gives the term's conventional meaning. This fact reflects the common observation that there can be nothing hidden or esoteric about the public conventional meaning of an expression. Such meaning derives from ideal projections from what people do and what people believe about what they do.

But this limitative point does not apply to the relation between "use" and cognitive value. Reflection on and analysis of linguistic use at any given time need not suffice to explicate cognitive value correctly. For the question of whether communally (or individually) acceptable normative characterizations *should be accepted*—whether the totality of communal (or individual) practices *should be revised*—can always be raised. Such questions may survive reflection on actual usage and may lead, by means of new theory, to changed practice. Cognitive value partly functions to make such questions possible and to interpret cases where their answer forces revisions in conventional meaning and actual use.

A second version of my thesis employs a noncognitive conception of fixing. Hold constant all discriminations that members of a person's community make, or are disposed to make, with a given term at a given time. (One may *include* among these discriminations all beliefs and assertions expressed with help of the term, with the proviso that one leaves unspecified the meaning and cognitive value of the term.) Then the linguistic meaning of the term at the given time could not possibly vary; but its cognitive value could—if the nature of the environment, or the subjects' relations to it, were in certain ways varied.

These points about the distinction between conventional meaning and cognitive value generalize to nearly all other conceptions of meaning that are similarly explicated or "fixed" in terms of conformity to some actual pattern of activity or dispositions to activity. For insofar as meaning is representational or purports to convey information, the question of whether the information has been correctly identified or characterized can be raised. One can thereby question whether the relevant underlying pattern of activity *should be as it is*. Cognitive values are individuated in such a way as to allow such radical questions to be raised.

Intentional mental states and events are individuated in terms of cognitive value. We have no other systematic, cognitively informative way of individuating them. Since communally accepted characterizations as well as expert opinion can be doubted, the ultimate authority regarding the application, explication, and individuation of a subject's intentional mental events does not derive solely from the actual motions, behavior, actions, usage, practices, understanding, or even (except trivially) thoughts of any person or social group. Our conception of mind is responsive to intellectual norms which provide the permanent possibility of challenge to any actual practices of individuals or communities that we could envisage.

TYLER BURGE

The University of California/Los Angeles