



The Transcendental Significance of Phenomenology

Stephen L. White
Department of Philosophy
Tufts University
Medford, MA 02551
stephen.white@tufts.edu
© S. White 2007

PSYCHE 13/1, April 2007

KEYWORDS: agency, demonstratives, experience, meaning, phenomenology, skepticism, transcendental arguments

Abstract: There is a well-known line of thought, associated with Donald Davidson, that connects the notion of a perceptual given—of non-linguistic or non-conceptual experience of the world—with skepticism. Against this, I argue that the notion of what is given in perception leads to skepticism only on certain interpretations. I argue, in fact, that there must be perceptual experience such that there is “something it is like” to have it, or that would provide the subject of a phenomenological analysis, if we are to block skepticism in its most radical forms. In particular, I claim that there is a distinctive phenomenology of the experience of agency. These phenomenological claims are conclusions of a transcendental argument according to which our having such experience is a condition of our having a meaningful language. Moreover, the same transcendental argument is sufficient to show the incoherence of radical skepticism about the external world. And I argue that the proper understanding of perceptual experience—as object involving—renders the standard objections to transcendental arguments ineffective.

1. Introduction

Transcendental arguments in their modern guise consist of premises describing conditions of our having a meaningful language (having intelligible thoughts, being subjects of experience) and an anti-skeptical conclusion. In this paper I argue that there is a distinctive phenomenology of agency, and the argument will be transcendental in this sense. But what exactly is the skepticism at issue? Most immediately, of course, it is skepticism about the meaningfulness and truth of our talk about the phenomenology of experience in general and agency in particular. Thus it entails skepticism about there being something it is like to experience the world and to experience ourselves as agents.

The argument, then, will be that it is a condition of our having a meaningful language that we actually have such subjective experience of the world. And the claim that there is a distinctive phenomenology of agency is central to the argument.

The argument for a phenomenology of agency, however, stems from the requirement that a meaningful language be grounded in an appropriate relation to the world. And this requirement connects the conclusion about agency to another conclusion of a more traditional sort: that the conditions for our having a meaningful language are intimately tied to the conditions for its being possible to justify empirical beliefs about the external world. To put the point in different terms, what we might call *meaning skepticism*—skepticism about the satisfaction of the conditions of our having a meaningful language—goes hand in hand with traditional epistemological skepticism about (the possibility of justifying our beliefs about) the external world. Thus the argument that a certain kind of phenomenology is necessary for our having a meaningful language is also an argument that it is necessary for our having justified empirical beliefs. Moreover, I shall argue that the phenomenology in question is *sufficient* to counter the fundamental skeptical argument against the possibility of such justification.

This may seem surprising—even paradoxical—to those familiar with the claims of Donald Davidson and others that it is the postulation of “experiential intermediaries” between our empirical beliefs and their objects that leads to skepticism about the external world. As Davidson puts it,

. . . Quine and Dummett agree on a basic principle, which is that whatever there is to meaning must be traced back somehow to experience, the given, or patterns of sensory stimulation, something intermediate between belief and the usual objects our beliefs are about. Once we take this step, we open the door to skepticism, for we must then allow that a very great many—perhaps most—of the sentences we hold to be true may in fact be false. It is ironical. Trying to make meaning accessible has made truth inaccessible. (Davidson 2001, pp. 144-145)

In what follows I shall criticize some of the assumptions behind Davidson’s position. What is ironical, given this fundamental disagreement, is that Davidson and I agree on many of the basic theses in this area. Unsurprisingly, we understand them very differently. The moral, I think, is obvious. The concepts currently invoked in these discussions are seriously underspecified. Any useful consideration of these issues requires that they be unpacked and made more explicit and more precise in ways that have not yet been attempted.

2. The Transcendental Argument in Outline

What are the most fundamental conditions for our having a meaningful language? Certainly among them is the condition that language be grounded in a connection to the world. We can see what this means by reflecting that there could not be an infinite backward regress of lexical definitions, of words defined in terms of other words. That is, there could not be an infinite backward regress of merely word-to-word connections. If we are to have a meaningful expression, such a regress must terminate in a connection to the world—an ostensive definition, a demonstration, a word-to-world connection. In the absence of such word-to-world connections, we would have nothing over and above a

formal calculus. That is to say, we would have nothing more meaningful than an uninterpreted formal language, formal rules of inference, and a subset of the set of well-formed formulas as axioms. Certainly we could, in an obvious sense, specify inferential roles in such a “language.” But this simply shows that such a formal inferential role is not sufficient for what we think of pre-theoretically as meaning.

What, then, constitutes the connection necessary for a meaningful language? Is it anything more than the causal connection that virtually everyone acknowledges between our referring expressions and their referents in the world? Intuitively, if a speaker’s grasp of meaning is concerned, a causal connection—one to which the speaker need have no access—does not seem an appropriate substitute for an ostensive definition or an act of demonstration. But why? And what is the connection between such intuitions about meaning and the kinds of transcendental arguments that concern us?

The answer is that the idea of a meaningful language is complicated by the wide range of things that might count as meaning or content—broad content, narrow content, language meaning, speaker’s meaning, meaning within an idiolect, meaning in virtue of a division of linguistic labor, and so forth. The notion of meaning most directly relevant to transcendental arguments and to skepticism, however, is that in terms of which we can rationalize a subject or ourselves. That is, it is the notion in terms of which we can do justice to the rationality and coherence of the subject’s beliefs and desires, and to the practical and theoretical rationality of the subject’s actions and speech. And this is because it is generally (and I believe correctly) assumed that such rationality and such coherence are constitutive of (are internally connected to) the very possibility of ascribing intentional states. Certainly Davidson is committed to such an assumption. (Davidson, 1973, Lewis 1974) Were it impossible by the skeptic’s own lights to make sense of, or ascribe such content to, his or her own thoughts or utterances, we would clearly have a successful transcendental argument. Skepticism of the sort in question would have proved self-defeating in the strongest possible sense.

The kind of meaning, then, that is relevant for our purposes is, to use the Fregean term, *cognitive significance*. This gives us a second necessary condition on meaning. Not only must a subject’s words be grounded in something non-linguistic. In addition this grounding must be in accordance with our commitment to doing justice to the rationality of the subject, and hence to the cognitive significance of the subject’s terms. At a minimum this means that ascriptions must satisfy Frege’s constraint in the various versions in which it is currently formulated. (Stephen Schiffer’s version: If x believes y to be F and also believes y not to be F , then there are distinct modes of presentation m and m' such that x believes y to be F under m and disbelieves y to be F under m' . Schiffer, 1978, p. 180.) As I have argued elsewhere, the satisfaction of these conditions is necessary but not sufficient for doing full justice to the rationality of subjects and to the cognitive significance of their terms (White 2007). But these are not the issues that concern us here. I shall say, as a kind of shorthand, that the condition for a language’s being meaningful with which we are concerned is that it be grounded subject to Frege’s constraint. But it is more accurate, if more cumbersome, to say that the condition requires us to ground the language in such a way as to do full justice to the rationality of speakers. And this requirement commits us to grounding language in such a way that we can solve all the relevant Frege problems as they arise.

We have, then, the relevant conditions for meaningfulness. And we have, as well, our answer to the question whether language can be grounded in a merely causal connection. The answer clearly is no. As is well known, two expressions, such as ‘Hesperus’ and ‘Phosphorus’, may be causally connected in the relevantly appropriate way to the same object, even though they differ in cognitive significance. And if it is objected that the causal chains themselves will be different, then it must be remembered that these are differences to which a normal subject need not (and in general will not) have access. Thus whatever role causal relations play, they are not themselves sufficient to ground meaning in the sense of cognitive significance. We need, it seems, some intermediaries, accessible to the speaker, to serve as the appropriate modes of presentation that the satisfaction of Frege’s constraint requires.

In the case of ‘Hesperus’ and ‘Phosphorus’, of course, there is evidently no problem. ‘Hesperus’ and ‘Phosphorus’, we might suppose, are associated with different descriptions, both of which are fully accessible to the subject. Such descriptions might be, for example, ‘the first heavenly body’ visible in the evening’ and ‘the last heavenly body visible in the morning’, respectively. But notice: The association of a description with a proper name just postpones the real problem. This association is just one more word-to-word connection, and what we required was a connection between words and the world. It seems, indeed, that we face a dilemma: If we try to establish a direct link between words and the world—one unmediated by linguistic-descriptive content—by appealing to a causal connection, we violate Frege’s constraint. And if in response we postulate descriptive modes of presentation, we reintroduce the backward regress of word-to-word connections.

We can see the difficulty more clearly if we consider Frege’s problem in its demonstrative guise. In Gareth Evans’ example, one may be prepared to accept, pointing to the bow of a ship out one window, “That ship was built in Japan.” Equally, one may be unprepared to accept, “That ship was built in Japan,” pointing to the stern of a ship out another window. And this is possible even though one has, without realizing it, pointed to the bow and stern of the very same ship (Evans 1982, p. 84). The temptation again is to introduce descriptions, but this move, as we have seen, is self-defeating. We need some direct or demonstrative connections to ground language. And if in every case we go on to postulate new descriptions in order to satisfy Frege’s constraint, we will have failed to explain what makes the language meaningful.

Frege’s constraint, then, imposes a very strong requirement on any account that purports to do justice to the meaningfulness of a speaker’s utterances and to the contents of his or her beliefs. And it seems we have no choice but to see the difference in cognitive significance of the two references to the same ship as a matter of the two different sets of sensory experiences that accompany them.

3. The Humean Argument for Skepticism

In spite of the arguments, however, the imposition of such a constraint and the willingness to ascribe such a role to perception will seem to some misguided. Davidson, as we have seen, has argued that the postulation of epistemic “intermediaries”—“the testimony of the senses: sensation, perception, the given, experience, sense data, the passing show”—leads to scepticism (Davidson 2001, pp. 141 and 144-145). Davidson

doesn't say what argument for skepticism he has in mind, nor what that skepticism consists in. It isn't clear, for example, whether the skepticism in question follows from the thesis that many or most of the sentences we hold true may be false or is identical with it. And if the latter, some might object that the so-called skepticism is merely common sense and, in any case, nothing with which we cannot comfortably live. (I shall have more to say about this later.)

There is, though, a form of skepticism so strong it precludes either of these reactions. Indeed one of the sub-themes of this discussion is that Davidson and others have seriously *underestimated* the significance of the skepticism that follows if we accept a sensory given *as traditionally conceived*. What I shall argue, however, is that Davidson misdiagnoses the deep source of this skepticism. A correct diagnosis, I shall claim, allows us to see an adequate phenomenology of experience not as the source of skepticism, but as the only possible defense we have against it.

The argument for a stronger form of epistemological skepticism than the claim that many or most of our beliefs could be false is Hume's (Hume 2001, Section XII, Part I, p. 184; cf. Hume 1968, Book I, Part IV, Section II, p. 212). Suppose one considers one's present conscious perceptual experiences (experiences, say, as of a room, furniture, etc.). Such experiences, we normally assume, are caused by what they purport to be experiences of—a room and so forth. But certainly, it seems, the same experiences (or experiences indistinguishable from these) could be caused in completely different ways—by a dream, an evil demon, the Matrix, and the like. Thus there is a gap between the character of our experiences and their causes. And bridging this gap requires an inference (since closing the gap between perception and its causes couldn't itself be a matter of perception). Because the connection is causal, however, the inference could not be *a priori*. Neither, though, could it be *a posteriori*. For such an inference could only be justified by a principle connecting the character of perceptual experiences with external causes. And the legitimacy of appealing to any such principle is exactly what the Humean skeptic questions. The conclusion is that one couldn't be rationally justified in preferring the commonsense hypothesis (that one's experiences are caused by a real world much as we take it to be) over any of the skeptical alternatives.

This conclusion is obviously much stronger than the proposition that many or most of our beliefs could in fact be false. But what has not been sufficiently seriously considered is that there is an even stronger conclusion in the offing. Suppose we accept, for the sake of argument, Hume's conception of experience as interpreted by contemporary sense-datum theorists. What we are given, then, in visual experience (to take the most important example) are mental entities—sense-data—that actually have the visual properties—shapes, colors, relative sizes—that external objects *seem* to have. The sense-datum normally caused by a round table viewed at an angle, for example, will be elliptical, and that normally caused by a white cup in red light will be red. On such a conception, we never experience external objects directly—indeed, we could not do so in principle. Doubts about the connection between sense-data and the world, then, could never in principle be resolved, because all we could ever get are more sense-data.

Such a conception of experience clearly supports the Humean argument for epistemological skepticism, since exactly the same sense-data could be caused in radically different ways—including all those suggested in the usual skeptical scenarios.

And such scenarios provide the doubts that could never be resolved in principle. But now put aside, for the time being, the epistemological question: whether we could ever know, or be rationally justified in believing, that ordinary external objects really exist. Consider a more basic question. Can we even attach any *sense* to the terms in our external object vocabulary? Grant, for the moment, that we understand the idea of a pattern among all of our sense-data—past, present, and future, actual and possible. (Of course even this assumption would be problematic for a consistent Humean. I shall discuss this below.) For example, what one thinks of as the table involves many predictable sense-datum patterns—including the one in virtue of which one would take oneself to be walking around it. What justifies the assumption that what we mean by ‘table’ is any more than the totality of such patterns?

Suppose we stipulate that ‘table*’ is to be understood in just this way. Then we can ask what justifies the assumption that ‘table’ means anything different from ‘table*’? Obviously the two terms are experientially equivalent on the Humean or sense-datum view of experience. Thus it seems that anything one could point to in characterizing the meaning of ‘table’ is something one could point to in characterizing the meaning of ‘table*’ and vice versa. What about a difference in inferential roles? It is tempting to think that we could define a physical object as a *cause* of sense-data and hence as distinct from any collection of them or pattern among them. But this just pushes the issue one step back, presupposing as it does that we have a concept of causation as something more than just another sense-datum regularity. (Hume was, of course, famously skeptical on this point.) The same point applies if we appeal to counterfactuals to talk about the back of the table that exists at the same time we are looking at the front (and not just if and when we walk around to see it). And indeed, the point is a perfectly general one about any appeal to a difference between the inferential roles of ‘table’ and ‘table*’.

So far, the conclusion is merely phenomenalistic. And those who thought we could learn to live with Humean epistemological skepticism¹ might suppose that we can live with this as well. But recall: We have been assuming unproblematic access to patterns among past, present, and future sense-data and among sense-data that are possible as well as actual. And the same considerations apply here. We cannot supply any content to the notion of a past or future sense-datum over and above what we understand in the notion of the relevant pattern of sense-data in the present. (These sense-data include, of course, those we class as memories and anticipations). Nor can we assign meaning to the notion of a possible sense-datum over and above the relevant pattern among those we actually experience. The upshot is what we might call “analytical solipsism of the present moment”—we can attach no meaning to terms that refer to anything other than our actual, present sense-data or their actual, present properties and relations. And this, surely, is a form of skepticism too extreme for anyone to live with.

4. The Source of Humean Skepticism

It seems, then, that there is no problem with Davidson’s claim that for *a certain conception* of perceptual experience—one in which it plays the role of an epistemic “intermediary” (in one sense of the term)—such experience leads to skepticism. Indeed, as we have just seen, Davidson grossly underestimates the seriousness of the skeptical consequences. But the question is how we should generalize the result. What, in other

words, is the source of the skeptical upshot? We can put the question another way by asking what the skeptic needs for a Humean skeptical argument to go through.

Notice first that the Humean skeptic is clearly not trying to call everything into doubt simultaneously. There is a well-defined domain and a corresponding set of propositions—*a posteriori* propositions about the objective, external world—with regard to which it is claimed that we have no rational basis for our beliefs. And there is a counter-domain and its corresponding set of propositions with regard to which no such challenge is raised. (For simplicity, I shall drop the distinction between the domains and their corresponding sets of propositions in what follows.)

What relation, then, between the domain and the counter-domain, is required to generate Humean skepticism regarding the domain? Certainly a necessary condition is that there should be no conditions regarding the counter-domain (where ‘condition regarding a domain’ is used to include its existence) that are sufficient for any condition regarding the domain. (This would rule out, for example, phenomenalism.) We can put this point by saying that Humean skepticism about a domain requires a counter-domain that is *domain neutral*. In the case of Humean skepticism about the external world, there must be an external-world-neutral counter-domain—that is, a counter-domain that can be completely characterized in a way that is logically and conceptually neutral with regard to the existence and character of the external world. Clearly Hume’s domain of conscious experience, understood in terms of sense-data, was intended to (and arguably does) have this feature. And equally clearly, Davidson’s radically externalist conception of content, according to which the objects of a belief are the causes of that belief, does not (Davidson 2001, p. 151). This is not the end of the story, however, since, Davidson’s radical externalism cannot satisfy Frege’s constraint. Thus it cannot do justice to the rationality of the subject to whom the mentality is ascribed—an extremely ironic result, given the constitutive role that rationality plays for Davidson himself in the ascription of intentional states to others.

But, it seems equally clear that this necessary condition (call it the *neutrality condition*) is not sufficient. Consider an example I have used elsewhere (White 2004a, p. 276). Imagine someone confined to a windowless room (perhaps under house arrest—a latter-day Galileo). Suppose that although he cannot leave the room, he has access to the outside world via all the usual modern media, as well as unlimited opportunities to communicate with others outside the room and to receive them as visitors. Clearly this example satisfies the independence condition. *A posteriori* propositions regarding the world outside the room are logically independent of those regarding things inside.² And, since the connection between events inside the room and events outside is merely causal, the external causes of events inside—the causes of the images on TV, the words in the email, the speech of visitors—might be radically different from what would normally be supposed. Thus with regard to the causal sources outside of events inside, one could entertain obvious analogues of most of the usual skeptical scenarios. In fact, the room and the events within it could exist even though *nothing* existed outside. But does this example really generate “outside-the-room skepticism,” analogous to Humean skepticism about the external world? (From now on I shall use ‘Humean skepticism’ loosely to refer to Humean epistemological skepticism as generated by the Humean argument we saw above.)

Intuitively the answer is no. And seeing why will help us see what has gone wrong with Davidson's argument that the postulation of anything that would ground or justify a phenomenological description leads to skepticism. First, as we saw, the Humean argument that there must be an inference from the character of our perceptual experience to our beliefs about the external world presupposes that we cannot perceive external objects directly—i.e., without there being an inference in any relevant sense. (Of course, the question of what counts as a relevant sense will have to be examined, and I shall do so below.)

With regard to outside-a-room skepticism, however, this presupposition is, arguably, false. Imagine, for example, that the person who never leaves the windowless room does so because he is the younger and more resourceful (but equally phobic) brother of Howard Hughes. Such a person would have, in addition to the resources already described, the money and access to scientific talent necessary to create prosthetic perceptual devices (some of which exist, some of which are currently in development) literally to look through the walls. Alternatively, the instruments already at his disposal might become virtually transparent with use. In such a case it might be completely appropriate to say that he is seeing not, say, images on a screen or any other features of an instrument, but external objects themselves. (And at the limit, it isn't clear that this alternative is really distinct from the first one. I shall say more about this below.) And, of course, combinations of these possibilities might give someone with almost limitless resources a view of the world unlike, and far superior to, that of any normal subject. Finally, of course, there is nothing in principle to prevent someone in such a situation from cutting a hole in the wall, or—the most radical solution of all—facing down his phobias and walking out the door.

These possibilities block any application of the Humean skeptical argument to generate "outside-a-room skepticism." The Humean argument presupposes that any empirical belief about the external domain *must* (logically, conceptually) be based on an inference from what is given in perceptual experience. And if that presupposition is false, then the argument that any such inference would be unjustified is irrelevant. But there is another way in which the Humean argument as applied to the case of house arrest fails. For the Humean argues that no *a posteriori* inference could get us from our sense experience as given "from the inside" to a conclusion about the external source of that experience. But why not? The answer is that such an inference would have to be justified by a principle connecting the subjective character of the perceptual experience with external objects, and any such principle would be among the things taken already called into question by the skeptic. Hence the appeal to such a principle would beg the question. But now notice that the subject confined to a room can do experiments to establish *inside the room* generalizations about causal connections not only between objects inside the room, but between objects inside and outside and objects all of which are outside. Can the Humean skeptic try to disallow this? Can the skeptic claim that we have evidence only for causal principles as they apply within the room? It seems not. What could it be claimed is special about the distinction between what goes on inside the room and what happens outside? Suppose it is alleged that the spatial distinction matters. Then experiments could be performed inside the room that would support the claim that some causal generalizations established inside are spatially invariant over at least reasonable distances. And for some events, the distance between events inside and outside the room

will be less than the distance between events inside. If it is alleged that it is a matter of the materials with which the room is constructed, we can bring those materials inside the room and experiment with transmissions through them—and so forth.

Finally, there is one more way in which the Humean argument goes wrong—and it goes wrong in the original context as well as in this one. It is an assumption of the Humean that the inference from the character of one's perceptual experience to conclusions about the external world could not be *a priori*, since we are talking about empirical knowledge. But this is open to question in either context, since the fact that the issue is empirical knowledge does not obviously rule out the possibility of a transcendental argument for anti-skepticism about the external world. Such an argument might deliver either the conclusion that some of our beliefs about the external world *must* be justified or that some are self-justifying or in no need of the justification that the Humean assumes their rationality requires.

Could it be objected that this treatment of outside-a-room skepticism is question-begging? In the first reply we took it as given, apparently without justification, that because we have unproblematic access to things inside the room, we can suppose that our access to things outside is sometimes, or at least can be, direct. In the second we seemed to take it as given that because we have unproblematic access to causal relations within the room, we can suppose we have such access to causal relations connecting things inside and outside. And in the third reply we apparently supposed that transcendental arguments could justify the truth of a realist conception of the external world, or that an anti-realist conception is more than just another version of skepticism. (Barry Stroud has questioned both these assumptions.³ I shall return to the issue of transcendental arguments later.)

What this possible objection points up is that the Humean skeptic's argument must take place against the background of certain theories. I shall argue in particular that it must take place against a background that includes a philosophical theory of perception and a philosophical theory of meaning. But how is this claim justified? After all, it flies in the face of another apparent assumption of Stroud's—that we cannot defeat the skeptic by appeal to a theory of meaning. Stroud's objection to such an appeal to demonstrate the incoherence of the skeptic's claims seems to be that we are as confident in our understanding of what the skeptic is saying as we could be in any theory that entails its unintelligibility. (Stroud is explicit about this where verificationism is concerned, but the point seems to be a general one. Stroud 1984, p. 205). Clearly there is something to this objection. And in any case, why is it *these* background theories that are required?

The answer to the last question is easy. The Humean argument says that there must be an *inference* from the character of our perceptual experience to the character of the external world and that neither an *a priori* nor an *a posteriori* inference could ever be justified. And since we can take it (because the Humean does) that a legitimate *a priori* inference will be grounded in meaning broadly conceived, the Humean clearly owes us an account of it if the need arises. (And it will in the context of the possibility of a transcendental argument for anti-skepticism.) Similarly, since it is claimed that there must be an *inference* to the nature of things in the external world, the Humean is committed to the claim that we cannot *in principle* perceive objects in the external world directly (i.e., in a way that needs no inferential support). Thus the Humean owes us an account of

perception from which this follows if, as it will be, such an account is necessary to clarify the nature of Humean skepticism itself—and in particular to clarify its nature by distinguishing it from outside-a-room skepticism. The choice of the two background theories, then, is dictated by the structure of the Humean argument itself. And, of course, arguments and theses from other areas, such as metaphysics or philosophy of mind, will be relevant to the extent they are necessary to make good on the deliverances of the background theories to which the Humean is committed directly.

Thus the appeal to such background theories, the theory of meaning in particular, is nothing like the appeal, to which Stroud objects, to such controversial theories as verificationism. Rather, it is an appeal to the theory of meaning (and perception) not in the first instance to refute the Humean skeptic, but to clarify the presuppositions of the Humean's argument. More specifically, it is intended to elicit the necessary and sufficient conditions for such an argument to go through. Indeed, this inquiry into the presuppositions of the Humean argument is so natural, it is something of a mystery why it has not been pursued more systematically. I can only speculate that theorists have been inhibited, either by Stroud's argument for the irrelevance of theories of meaning, the fact that the issues cut across the major divisions within contemporary philosophy, the fact that it requires serious consideration of such unfashionable topics as sense-data and phenomenology, or, more likely, some combination of the three.

Suppose that, on the assumption that we have such background theories on the table, we have specified the domain and counter-domain, and specified our access to the counter-domain as unproblematic. Then necessary and sufficient conditions for the Humean skeptical argument to go through regarding the elements of the domain can be stated as follows.

Regarding the elements of the domain (objects, facts, events, etc.), the Humean argument for Humean skepticism applies, if and only if, it is a consequence of the theory of meaning that there are no conditions regarding the counter-domain that are sufficient (logically, conceptually) for any condition regarding the domain, and it is a consequence of the theory of perception that none of the elements of the domain can be perceived or observed directly.

Consider skepticism about the external world against the background of the following assumptions. The theory of perception says that all we can in principle (as a matter of conceptual necessity) perceive or be given directly are sense-data. The theory of meaning says that there is no *a priori* inference from sense-data to elements of the external world (as there would be, for example, if phenomenalism were true). Given these assumptions, the Humean argument applies, and arguably it goes through. (Again we are putting aside transcendental arguments for the time being).

With regard to the question whether there is a Humean argument for “outside-a-room skepticism,” the answer is more complicated. This potential form of skepticism was defined by saying that inside the room we have unproblematic access to ordinary objects. But, as we have seen, it is not immediately apparent what this entails regarding the question whether objects outside the room are, in principle, directly perceivable. As we saw above, the initial argument that there is no such thing as “outside-a-room skepticism”—that the Humean skeptical argument cannot be made to apply—was

inconclusive. For it seemed a possible objection that to suppose objects outside the room could be perceived directly, even in principle, begged the question against the would-be skeptic.

This, however, is where the more abstract and general characterization of the necessary and sufficient conditions for Humean skepticism pays off. For what philosophical theory of perception could the would-be skeptic appeal to? If the background theory of perception said that all we could perceive directly were sense-data, then the Humean argument we have been considering would show (arguably) that we could not have unproblematic access to ordinary objects *inside* the room. And this violates one of the conditions to which the would-be skeptic is committed. For the would-be skeptic wants to claim that we *can* make sense of “outside-a-room skepticism” on the analogue of Humean skepticism about the external world. Thus the proponent of “outside-a-room skepticism” is committed to a background theory of perception that gives us *unproblematic* access to ordinary things and events within the room, and the Humean argument shows that this means direct perceptual access to some such objects. However, a perceptual theory that gives us direct perceptual access to ordinary objects within the room will give us, at least potentially, perceptually direct access to objects outside. And, as we have seen, this is all we need to defeat the Humean argument in the context of house arrest.

What is more relevant in the broader context than the commitments of the would-be skeptic about the world outside the room is that we think that (sometimes) *we* have access to ordinary external objects unmediated by anything like sense-data. And the point of the comparison between Humean skepticism and a range of (possible) analogues is to see what is essential to Humean (or as I shall call it later “philosophical” or “transcendental”) skepticism, and whether in fact it arises naturally out of our ordinary epistemic practices. If a Humean skeptical argument required a sense-datum theory of perception among its theoretical background conditions, as some seem to assume, it would be easily blocked (Searle 1983, pp. 71-76). (As we shall see shortly, however, it doesn’t.) If, as Stroud and the other new Humeans seem to hold, there are *no* controversial theoretical presuppositions that could conceivably be rejected, then it seems that such skepticism is indeed a fact of life. Settling this issue will be crucial later to distinguishing between philosophical and what I shall call “real” skepticism.

External-world skepticism based on a sense-datum theory is genuine philosophical skepticism because for any hypothesis regarding the external world, there is, in principle, no way of obtaining evidence for or against it. Thus, as Hume claims, there can be no rational justification for preferring one hypothesis about the external world over another. Outside-a-room skepticism isn’t Humean or philosophical skepticism because there are perfectly reasonable ways of gathering evidence for and against the various hypotheses that might be entertained about the character of the world outside.

But, are there other cases of genuine philosophical or Humean skepticism that don’t involve the sense-datum theory? The answer is yes, and by now the reason should be obvious. Consider any form of skepticism, and suppose that it draws a *principled* distinction between elements in the domain and elements in the counter-domain. Suppose also that it proceeds on the basis of a theory of perception according to which that distinction entails that we have unproblematic access to (including direct perception of)

some of the elements of the counter-domain and no possibility of direct perception, even in principle, to elements of the domain. Then assuming that there are no conceptually sufficient conditions regarding the counter-domain for conditions regarding the domain, we have an example that yields philosophical skepticism. And this is true regardless of what the theory of perception is like in respects other than those specified. Moreover, as we shall see, there are many such theories of perception, of varying degrees of plausibility.

The sense-datum theory, then, is sufficient for generating philosophical skepticism, but not necessary. And we can see immediately another advantage of the general characterization of Humean or philosophical skepticism. This characterization explains why differences in the assumptions about what counts as a directly perceivable or sense-datum property are irrelevant where philosophical skepticism is concerned. For it is a characteristic of all the variations of the sense-datum theory (that don't define external objects in terms of sense-data) that what are perceived (given etc.) directly are the sense-data, however rich they may be conceived as being. Thus it is a characteristic of all such theories that we perceive external objects only in virtue of our perception of (access to) sense-data—hence indirectly. Of course, some theories of perception are more plausible than others. And there is something perennially appealing about a theory like Berkeley's according to which the properties given directly in (visual) experience are identical (or very tightly connected) to the properties of the image on the subject's retina. (This appeal is, in my view, completely spurious. I shall have more to say about this below.)

For an example of Humean skepticism that does not involve or presuppose the sense-datum theory, consider our access to the past. It is natural to think that what we are given or perceive directly is in the present and that things in the past are unperceivable (directly) in principle, or as a matter of conceptual necessity. Now consider a theory of perception according to which we have direct perceptual access to ordinary objects and to many ordinary properties (properties of external objects, not sense-data, etc.). In fact, as we saw above in the analogous case of "outside-a-room scepticism," we have to make this assumption if we are to have unproblematic access (in principle) to things in the counter-domain (the present). In such a case, and in contrast to the case of house arrest, the Humean argument does apply and does, apparently, go through. There is a gap between what we can in principle be given in perception and any belief about the past. And because in principle we could never be given anything about the past directly, the gap is one that no justified *a posteriori* inference could bridge. For anything to which we could appeal in support of a hypothesis about the connection between the past and present will simply give us more information about the present. Thus, unless our conception of the past is a form of idealism analogous to phenomenalism regarding the external world, skepticism follows. And, as we have seen, pushed to the limit, the skepticism in question is meaning skepticism.

For another example, consider the problem of other minds. Imagine that we have unproblematic access to all of the kinds of external objects that we normally think we do (objects about which we think we can, therefore, at least in principle, form justified beliefs), except other minds. The assumption that one cannot, on conceptual grounds, experience in the relevant sense another's pain—that anything that could be experienced

from the inside as pain would be one's own—is very widely accepted. And the idea that one could make an analogical inference to a justified belief about what another person is feeling on the basis of one's knowledge of what one feels oneself in similar circumstances is not generally considered convincing. It is, at least, unclear what could ground such an inference since one case seems to provide a weak basis for an inductive argument.

If this were the strongest version of the problem of other minds, of course, it would not support Humean skepticism. Even very weak inductive support for a proposition is *some* relevant evidence. And what characterizes Humean skepticism is that nothing could *count* as evidence for or against any hypothesis about conditions in the relevant domain. But this way of putting the problem with the argument from analogy is inadequate. For our so-called knowledge of our own mental states is not knowledge of something that could in principle be had by anyone else. This is because of the assumption, made by proponents of the argument, that our knowledge of our own mental states is nothing more than what we have in virtue of experiencing them. Alternatively we can say, for example, that all that we can mean in using 'pain' is what we can mean in virtue of knowing when *we* have it. But if this is the case, then it is a matter of conceptual necessity that there is nothing we could *mean* in ascribing a pain to someone else. We have no conception of pain as something independent of what it is like to have it, hence nothing independent of its occurring in us. If this is the case, the inductive inference is not simply weak, but meaningless, as is any analogical inference or inference to the best explanation. Thus we have Humean skepticism about other minds. And if so, we have another example of the kind we wanted—a Humean skeptical argument that has no connection with any commitment to a sense-datum theory. Similar arguments could be made for a wide range of other domains, including causation, the future, modalities, meaning, and, on some conceptions of what is perceivable, theoretical entities.

In each case what generates Humean skepticism is not the sense-datum theory, but the unobservability in principle of the elements in the domain (assuming the neutrality condition holds). Thus there is a sense in which philosophical skepticism *is* the result of epistemic intermediaries if these are understood broadly—so broadly that some type of thing counts as an epistemic intermediary between a subject and a domain just in case the following holds: Elements of the domain are in principle not directly perceivable by the subject, and that subject perceives those elements indirectly in virtue of perceiving (being given, having access to) things of the type in question directly. But, as we have seen, epistemic intermediaries understood so broadly come in many different forms. Understood in this sense, arguably (and not implausibly), present events (some of them) have this status for us relative to events in the past and future; actual states of affairs have it relative to possibilities, conjunctive relations have it relative to causal relations, and bodies have it relative to other minds. Thus abolishing sense-data, the given, perceptual experience, and the like is virtually no use in helping us understand and deal with philosophical skepticism.

Of course it might be suggested that if we deal with philosophical skepticism about the external world, we can simply learn to live with everything else. But this is a mistake. For we don't really get the external world simply by eschewing the given. What we get is the present moment shorn of all its modal characteristics, causal relations and

potentialities, meaning of any kind, and other minds. And recall that where philosophical (Humean) skepticism is concerned, the epistemic deficit is not merely a matter of certainty. Rather it is the lack of any rational justification for preferring any hypothesis about the problematic domain over any other. Recall, moreover, that philosophical or Humean skepticism has its immediate counterpart in meaning skepticism regarding the terms in which we purport to describe the elements of the domain in question. Even, then, if we completely repudiate any notion of a sensory given, we are left with a form of skepticism with whose truth or unanswerability we could not even pretend to be reconciled.

Thus far I have sketched the outlines of a transcendental argument that our having a meaningful language (even one of such ordinary objects as tables and chairs) requires our having perceptual experience of the kind Davidson evidently wants to deny. For we need, it seems, perceptual experience as opposed to linguistic-descriptive content in order to deal with the demonstrative versions of Frege's problem. And, as we have seen, Davidson's attempt to associate such experience with philosophical skepticism fails. What we have instead is a clear association between such skepticism and epistemic intermediaries in the broad sense defined above, including, as a special case, sense-data. That is, there is a clear association between skepticism and things perceived (given) directly in virtue of which the elements of a problematic domain are, as a matter of conceptual necessity, perceived indirectly at best. What Davidson never provides is an argument for the assumption that *all* sensory experience must have this character. In the next section I shall provide a more detailed version of the transcendental argument, in the course of which I shall argue that this assumption is false.

5. The Transcendental Argument

As we saw in the outline of the transcendental argument (Section 2), linguistic-descriptive content is not sufficient, either for the characterization of a meaningful language, or for the specification of the contents (in terms of cognitive significance) of a subject's beliefs. There must, in addition, be content in virtue of which the subject's language is *grounded* subject to Frege's constraint. More explicitly, there must be connections between words and the world that are direct in the sense that they are unmediated by linguistic-descriptive content—and which, nonetheless, provide the modes of presentation that the satisfaction of Frege's constraint requires. And this result, as we have seen, stems from the impossibility of an infinite regress of linguistic descriptions, the need to distinguish a meaningful language from a merely formal calculus, and the requirement that we deal with the Frege problems in their demonstrative versions.

But to say that our access to external objects is (sometimes) direct in this sense (i.e., unmediated by linguistic description) is still not enough. Such a requirement is fully compatible with a theory like Hume's or a sense-datum theory according to which our words get their meanings through association with collections of, and patterns among, sense-data. What rules out such theories is not the regress argument, the inadequacy of formal inferential roles in grounding meaning, or the need for non-linguistic modes of presentation in cases of demonstrative access to external objects. Such theories are ruled out, as we have seen, by a different argument—that because they preclude the direct

perception of external objects in principle, they make the Humean skeptical argument and its meaning-skeptical counterpart unanswerable.

It follows, then, that our access to external objects must, in at least some cases, be direct in another sense. It must, as we have seen, be non-inferential. But, again, our access must be direct in this sense while being perspectival. Equivalently, our access must be direct, while doing full justice to all the relevant distinctions in content where our intentional states are concerned and where content is understood as cognitive significance. This means it must be direct while allowing us to do full justice to the rationality of the subjects of intentional states and thereby satisfying Frege's constraint. And with the sense-datum solution ruled out, it is unclear what an account of such access could look like.

To see what such access might amount to, consider two more definitions of directness that, while ostensibly different, are, when properly understood, arguably equivalent: (1) Non-descriptive access to some element of a domain in virtue of an intentional state (i.e., the availability to the subject at the personal level of that element in virtue of an intentional state) is *direct* just in case the subject has such access and does not have it in virtue of having access to anything else. (2) Non-descriptive access to some element of a domain in virtue of an intentional state is *direct* just in case if the object doesn't exist, there is no complete content. We can see that these come to the same thing, since if our access (to an external object, for example) is direct in sense (1), then there is nothing (at the intentional and personal level) to which we have access and in virtue of which we have access to the object. The object itself, we might say, is an ineliminable part of the intentional content that provides or constitutes the subject's access to the object (See Evans 1982, pp. 64-73). Thus, if the object doesn't exist, the intentional content will necessarily be incomplete. And if the access is direct in sense (2) so that if the object doesn't exist there is no complete sense, then there couldn't be an intentional item to which we have access and which is such that in virtue of this access we have access to the external object in question. For such an intentional item would complete an intentional content, assuming it was complete in other respects, even if the external object itself didn't exist.

Not only does it seem that there is a natural understanding according to which (1) and (2) are equivalent, it is also natural to suppose that they are both equivalent to the requirement that our access to the objects in question be potentially (or, in the case of external objects at least sometimes actually) non-inferential. For if our access is direct in the sense of being non-inferential, then there is nothing given to us in virtue of which the object is given—if there were, then there would be room for an inference, and if so, one would be required. And if there is nothing in virtue of which the object is given, then there is no room for an inference and, arguably, no need for one.

But what does it mean for access to be direct but perspectival? Or, to put the point in a different and, I think, more useful way, what kind of account could we give of such access? We can see immediately, it seems, that it will not be a reductive account in any ordinary sense. For if we try to characterize such access in causal terms (e.g., causal terms involving physical or functional states, or even states of a mental substance) we fail to capture one of its essential features given its defining role in satisfying Frege's constraint—its accessibility to the subject to whose rationality we are committed to doing

justice. The account, then, will be in terms of states that are both subjective and intentional. But, again, what could such an account look like?

To begin to answer this question, I want to make a provisional assumption. The access we want—access that is direct but perspectival—is a matter of our perceptual states. This assumption is natural given what is apparently indisputable: that language is grounded in our demonstrative access to objects. For in our example of the *demonstrative* version of Frege’s problem, the Evans’ ship example, the difference in the two modes of presentation of the same ship seemed clearly perceptual in nature. But to say that the assumption is provisional is not to say simply that it is one for which, as yet, no argument has been given. Rather it is to say that the assumption is, in this context, a partial truth, and that in a broader context, and in something closer to the final analysis, it will appear in a different light.

Of course, to say that our access to the ship in the Evans case is a matter of our two different perceptual experiences (one of the bow, one of the stern) just pushes the problem one step back. What account can we give of the nature of the perceptual experiences themselves? The two (apparently) different perceptual contents cannot be understood solely in terms of linguistic-descriptive content, since this would not take us outside the circle of word-to-word connections and hence couldn’t *ground language*. Nor could the difference between the perceptions be a matter of different causal chains, since (given that these are not necessarily, and not generally, available to the subject) this wouldn’t serve to ground language *subject to Frege’s constraint*. This seems to leave as the only other alternative a difference in the apparent visual properties available to the subject in the two experiences. But we have, as yet, no way of understanding this, compatible with the satisfaction of Frege’s constraint, except in terms of sense-data. And, as we have seen, far from grounding language, the sense-datum theory of perception leads to the most radical form of meaning skepticism.

To see the existence of an alternative, consider the things that the two different perceptual experiences allow us to *do*—in particular the basic actions they allow us to perform. The visual perception of the bow, for example, allows one to point to it, trace its outline, move toward or away from it, move in relation to it so as to get a better view—to produce a better line of sight, or to obscure the line of sight (perhaps so one cannot be seen by a subject located on the bow), etc. We can think, then, of a very large package of basic action possibilities associated with the view of the bow. And this will be a different package or set from the one associated with the view of the stern. Where the project of grounding language is concerned, we could express the underlying intuition in terms of the familiar slogan (though one which I believe has elicited more lip service than genuine appreciation) “Meaning is use.”

In spite of the intuitive plausibility of the move to basic-action possibilities in the context of grounding language, however, it seems anything but intuitive where perceptual experience is concerned. (There is no slogan “Perception is use,” nor does it answer to any currently felt need.) And to reduce perception to agency, or (worse) the possibility of agency, seems to leave out precisely what is essential to perception—that it is experience, that there is something it is like to have it. What I am suggesting, however, is neither a reduction of the perceptual to the agential, nor an analysis or an explanation of the former in terms of the latter. As will become clear, this analysis of direct access could just as

easily be called a perceptual analysis of agency as an agential analysis of perception. That is to say, the analysis treats neither perception nor agency as more primitive than the other. At the moment, however, I want to focus on perception and make three points. The first is a point about the intimacy of the connection between perception and action. It involves an argument designed to show that, contrary to what we might suppose, there is no logical or conceptual gap between our conscious perceptual awareness and our corresponding agential capacities. Needless to say, the argument is designed to make the agential approach to direct perception less unintuitive. The second point is made with the same aim, but proceeds from the opposite side. That is, I argue briefly for the inadequacy of the sense-datum conception of visual perception (the source of virtually everyone's "pre-theoretical" intuitions, regardless of their theoretical commitments). And I do so on phenomenological grounds. The intention is to counter the claim that the treatment of perception I'm putting forward is motivated by exclusively theoretical considerations and ignores our ordinary intuitions. The third point is a more explicit explanation of the theoretical payoff of this treatment of direct perception.

On the first point, consider an argument I have used elsewhere—that the notion of perfect blindsight is incoherent (White 2004b, pp. 215-216). As is well known, subjects with certain kinds of brain lesions report blind areas in their visual fields. When pressed to answer questions about visual features of external images falling in such areas, however, their rates of success are considerably better than chance. Such subjects, then, will have many of the agential capacities cited in connection with Evans' ship. There is no reason to suppose that they are limited to describing the various kinds of shapes and motions they detect. They could (almost certainly) point in the appropriate direction, trace a rough trajectory of a movement, trace the rough shape of a stationary image, and so forth. In spite of this, taking their testimony at face value, they have no visual experience. Thus it might be thought that this is exactly the example to show the conceptual *independence* of perceptual experience and agential capacities. But now consider a thought experiment. Imagine that the agential capacities of a blindsight subject get better and better until they match those of a normal subject. Imagine, that is, that this so-called blindsight subject can do all the same things, answer all the same questions, and do so with the same confidence and spontaneity as a normal subject. Suppose in particular that whereas the guesses of actual blindsight subjects have to be elicited by others, beliefs about the part of the world "hidden" in the "blind" area occur spontaneously and in the same profusion as they do for the normal subject. And suppose that far from seeming alien, or as though they "come from nowhere," these beliefs are as well integrated with the subject's practical capacities as those of any subject with normal sight. Suppose, in fact, that the occurrence of such beliefs is not different in any relevant way for the parts of the world inside and outside the "blind" area. Can we still suppose there is any sense to the idea that regarding that area the subject is blind?

Imagine looking at your computer keyboard with your eyes closed. Certainly you can answer some of the same questions about its visual appearance that you could answer with them open, but many of them you could not. For example, where are the brightest highlights? Where are the darkest shadows? Where is the Intel logo in relation to the Windows logo? Where is it in relation to the 'P' key? Now imagine that you get better and better at answering such questions. And bear in mind: This is not simply a matter of certain words popping into your head. You have the same demonstrative visual abilities,

and so the same demonstrative access, that you have with your eyes open. You can see complicated visual patterns at a glance that you would be hard pressed to describe in words, but they are ones you could match with confidence to congruent patterns flashed on a screen. You can trace the outlines of such patterns, you know how to move your head to get a better view, and you know when you have moved it too far to see the whole keyboard. You can obscure any feature of the keyboard by holding your finger at an appropriate location between it and your eye. And instead of describing the relation between the Intel logo and the 'P' key, you can use your finger to trace a connecting path. Moreover, when performed, these are full-blooded actions. They are given in experience as doable before being performed, they are performed with confidence, and they are given as successful when the performance is complete. Can you really imagine all this and not imagine that your experience is phenomenologically equivalent to seeing?

My second point starts from the side of visual experience and is offered to help defuse the idea that the proposed treatment leaves something out. The point is a critique of the sense-datum view, but on phenomenological, not theoretical, grounds. In fact, I shall argue for what I have called elsewhere an inflationary/deflationary phenomenology of experience (White 2004b; White 2004a pp. 281-282). By 'inflationary' and 'deflationary' I mean inflationary and deflationary relative to a strict sense-datum theory. And by this I mean one according to which, for visual perception, the subjective visual field is modeled on the retinal image. On such a theory, the only sense-datum properties will be properties that have obvious retinal counterparts—shape, color, and relative size. An inflationary/deflationary phenomenology, then, is one according to which what we are given directly in experience (and I shall be concerned primarily with visual experience) is both richer (in some respects) and more impoverished (in others) than such a sense-datum theory would allow.

Consider deflation. Here there is a wide range of examples suggesting that our most direct access is not to sense-datum properties because, the evidence suggests, we have no such access at all. Most people, for example, cannot approximate the apparent shape of their car windshield (as viewed from the driver's seat) using four straight lines. Of course, in the case of such relatively simple shapes, there is a tendency to think that the problem must be a matter of memory. But even when allowed to look, most people cannot, while standing, estimate with any accuracy the ratio of the apparent length of their hand held four inches from their eyes to the apparent width of their foot. Moreover, there is a clear argument that visual sense-data *could not* have the character that the theory attributes to them. For the visual field as it is given to us seems to have the following properties. On the one hand, it seems to encompass a visual angle of roughly a hundred fifty degrees, and, on the other, shapes of objects seem undistorted, even at the peripheries. In addition, objects seem to maintain their shapes, even as we turn our heads. But images cannot have these three properties at the same time. This is apparent if one watches the shapes of objects outside the middle of the frame in any film in which a wide-angled lens camera pans across the scene. And this is not an empirical point about current lens technology. Rather, it is a point about the constraints that can and cannot be satisfied simultaneously in mapping three-dimensional space onto a plane—and one that was well known to Renaissance theorists (White 2004b p. 211).

On the side of inflation, notice, as Berkeley did, that even depth could not be a sense-datum property, strictly speaking. Richer examples include the fact that a staged fight will look different when one thinks it is real and when one realizes it is not. (Imagine that you see the same set of physical movements twice, once under the first impression and once under the second.) Similarly, a park, viewed from the same angle, will look different when you are completely new to the city in which it is located and when the city and the park's orientation in it have become completely familiar. The most compelling examples for many, however, are those involving the expressive features of other people. We can be, and normally are, given another's distress directly, and not on the basis of an inference from the purely geometric features of the person's face or bodily posture (much less their sense-datum counterparts). This example works as well as an example of deflation, since if we are given the person as distressed, it is difficult to see how we *could* be given the pure geometry of the relevant body parts. There would, it seems, be incompatible determinations of the same part of the visual field—something the sense-datum theorist cannot accommodate. In fact, if this is right, inflation *entails* some form of deflation. If we have inflation, then what we are given directly are richer properties than the strict sense-datum theory allows. But if this is the case, then we are *not* given some of the sense-datum properties that the theory postulates, and we have (some degree of) deflation. Moreover, we should not forget the important research traditions stemming from the work of Michotte and J. J. Gibson (Michotte 1963; Gibson 1986). Empirical studies deriving from the former involve the direct perception of the rich properties constitutive of causation. Those following on the latter concern the direct perception of humanly relevant functional properties of the environment (physical-geometric features may be given directly as hiding places or escape routes, for example), as well as intentional and mentalistic properties.

The third point is intended to clarify the theoretical payoff of the agential approach, while anticipating somewhat the discussion of agency to follow. Recall the general structure of the argument. I argued that language must be grounded subject to Frege's constraint and that this requires that some perceptual experience be both direct and perspectival. This raises the question how any experience could be both. I therefore introduced the agential approach to perception, which will make it possible to appeal to the notion of *know-how* in beginning to answer this question. And shortly I shall argue what I have already claimed in passing—that the agential analysis does not leave out what experience is like because agency has the rich phenomenology I have just been describing. Moreover, it will be shown that this is not merely a disputable intuition about agency, but a requirement of our making sense of action. But then the following will have been shown: We need perception to ground language subject to Frege's constraint, we need action to give us perception that is both direct and perspectival, and we need a rich phenomenology (and, in particular, a rich phenomenology of action) to give us genuine agency. And the conclusion that we need a rich phenomenology of action in order to have a meaningful language was exactly the conclusion we wanted.

What, then, is it about action that entails that it must have a rich phenomenology of the kind I have been gesturing toward? I shall give three arguments, none of which appeals simply to our intuitions about what the phenomenology is like (of which there are many). The first stems from the possibility of what I have called the passive subject (White 2004a, p. 284; White 2004b pp. 201-202). Imagine a person who claims not to

“get” the notion of agency. (He is, at the outset, capable of speaking, and so not completely passive. But we can imagine him rapidly getting worse and eventually falling silent.) In the initial phase, we can imagine that the passive subject is completely competent and comfortable in using the language of objective science to explain and predict events, and that it is only the specifically agential vocabulary (including perhaps evaluative notions used in the specification of the goals of action) that raises a problem. We can, however, suppose that he was once a normal subject and thus that he has, in one sense, a perfectly good grasp of the use of agential terms. The passive subject knows what generalizations it is appropriate to accept (e.g., “One is responsible in the first instance only for what one does”) and even what is appropriate in particular instances. (If someone is the causal source of spilled milk and wasn’t pushed, tripped, etc., then it is appropriate for that person to “own up” and say, “I did it.”) We could say, then, that the passive subject still grasps the inferential role, but only the inferential role, of the agential vocabulary. Alternatively, we could say that his knowledge is like an ideally well-informed anthropologist’s knowledge of the “demonological” language of some group when the anthropologist has no belief in demons. The passive subject finds, we might say, nothing *in the world* that could justify or motivate the specifically agential vocabulary.

What is it that we have and that the passive subject lacks? It will be obvious that the example is so constructed that the answer that strongly suggests itself is something perceptual. For what else could it be? First, we can suppose that the passive subject’s motor systems are all intact, indeed that this has been tested. We can imagine, for example, that while the passive subject still had some capacity for spontaneous action, he was made to think that someone was about to put poison in his intravenous bottle. And we can suppose he jumped up without thinking to knock the poison away and then immediately collapsed back into passivity. Second, recall that we are assuming that the passive subject agrees with us about all the facts that can be characterized without a specifically agential vocabulary and that he still retains the inferential role of the agential expressions. And finally we can assume that there is no lack of motivation, that the passive subject wants desperately to be active again. What is it, then, that the passive subject is missing?

It seems that all the passive subject could lack, given that he retains the inferential roles of all the terms in the specifically agential vocabulary, is something on the side of perception and know-how. It must, it seems, be something at least closely related to what Gibson describes as the perception of affordances—the perception of things as affording certain actions. As we have seen, according to Gibson things like chairs are not normally given simply as solid objects with certain geometric configurations. Rather chairs are given as things one can sit on (or stand on to change a light bulb, etc.), and, more generally, things are given in their agentially relevant aspects. The passive subject, then, lacks the rich phenomenology in virtue of which things are given as doable, situations are given as opportunities, and so forth. And the point of the example is to suggest by contrast what we have and thus the richness of the perceptual experience necessary for agency.

Some may find the example of the passive subject difficult to imagine. But consider the past. For us the past is dead with respect to action—we have no idea what it

would mean if someone were to suggest that we act toward the past. But why? Certainly it seems there are no conceptual difficulties in the idea of backward causal chains. To suppose that instances of backward causation exist simply means that where the direction of time has been established by the preponderance of causal processes, there are local or isolated exceptions that run in the opposite direction. What, then, would it be like to act toward the past? Imagine that you can control the cursor on our computer screen with your mind as a basic action (without touching the mouse, etc.). (This is not currently science fiction.) Suppose the cursor is now in the middle of the screen, moving slowly to the right on a horizontal path. You are asked to make it the case that the cursor reached its midpoint position on the screen from a point above the midpoint. You “tune in” to a point in the past when the cursor was at the left edge of the screen, midway between the top and bottom and moving slowly and horizontally to the right. And you guide the cursor up, toward the top of the screen, and then down to the midpoint that it occupied when the instruction was given.

What, though, would it be to be “given the past” in the right way? One would have to be given some part of the past directly (not via present representations) if we want it to be a basic action in virtue of which we bring about something in the past. But vivid memories with which one is fully engaged, and in which one is given past events and not present images of those events, are experiences in which one *is* given the past directly. One’s being given the past in the right way, then, could be thought of as being like this—without, however, one’s being given the past events necessarily as things one originally experienced oneself.

There is another way in which we might conceive of acting toward the past. We might begin with a common representation in films of the experiences of psychics. Typically the subject has a vision of how events went in the past—a vision of events that he or she did not witness. Imagine having such visions. At first one doesn’t know how to interpret them. With experience, however, it comes about that the events are given as in the past. How does this happen and what is it like? We can imagine it by analogy with the theory-and-practice-laden experience of the trained scientist for whom the equipment is no longer given as metal and wires, but as a particle accelerator. And, given that we have the idea of a direct perception of some part of the past, there is an even simpler way in which one might act toward the past. If the instruction is bring it about that the cursor, now in the middle of the screen, is at a point above the middle at some time in the past, one might focus on an appropriate time and put it there—i.e., move it to that place at that time. This would be no more mysterious than the fact that in moving the cursor as a basic action aimed at the immediate future one can make it appear where one wants it—or that one can do so with one’s arm if one is told to raise it.

With regard to the passive subject, then, we can say this. We can well understand the phenomenology of the experience of the passive subject because we can understand it by analogy with our own experience of the past. Unlike the passive subject, however, we experience ourselves as agents over a broad range of cases. And because we do, we can understand by analogy what it *would be* like to experience some aspects of the past as a domain open to our purposive intervention.

If, as these arguments suggest, the example of the passive subject is intelligible, then we have an *argument* that agency requires a very rich phenomenology of experience

(rich relative to the strict sense-datum theory). And the argument supports what, in any case, we know on experiential grounds. We see things and the world as affording us a great many possible actions. And we do so by seeing things as tools or as means (the rocks in the stream are a bridge), seeing things as opportunities, (the distance between two cars is not a space of so many meters, but an opportunity to change lanes), and, as I shall argue later, seeing things as desirable or valuable. This is on the inflationary side of the inflationary/deflationary phenomenology. On the deflationary side, we may cease to see things at all. The automobile we are driving is, like our own bodies, very often not experienced; rather we experience the world *through* it. The space between cars that is an opportunity to change lanes is an opportunity for us, given that we are driving a mid-size sedan. But it might not be an opportunity (nor given as one) if we were driving a tractor-trailer truck. And to repeat—these affordances must be given to us directly in experience. Were objects given in perception in an agentially neutral way and then interpreted as providing opportunities, say, as the result of a linguistic description, we would have no explanation of what the passive subject lacks. Furthermore, we would have no explanation of the meanings of the expressions in the alleged linguistic interpretations. For such expressions would not be definable in non-agential terms. If they were, we could tell the passive subject what he is missing, and we could simply be told, without appeal to any experiential analogies, what we are missing regarding the possibility of acting toward the past.

The second argument that agency requires a rich phenomenology of experience stems from the apparent inadequacy of current attempts to characterize action. Davidson's original suggestion that a piece of behavior was an action just in case it was caused by an appropriate belief-desire pair is undermined by his well-known climber case. In the example, one climber is suspended by a rope held by another, who is in danger of losing his grip on the rock face. The latter suddenly thinks, and thus believes, that if he drops his companion he can ensure his own survival. And certainly he has the desire to survive himself. As the example is told, however, the recognition that he could even think such a thing so shocks and unnerves him that he accidentally lets go of the rope. Thus an appropriate belief and desire cause the behavior, but, by hypothesis, the climber's dropping his companion is an accident.

The temptation is to require that that the appropriate belief and desire cause the behavior "in the right way," where this is to be specified in neurophysiological terms. But this seems clearly wrong, giving, as it does, the neurophysiologist the last word about our actions. For in so doing, it threatens the notion of action itself by undermining the authority that as agents we normally suppose we have. Suppose we have what looks to the subject, and to fully informed (non-expert) spectators, like a paradigm case where action is concerned. The subject is acting not only on the basis of genuinely held desires, but deep value commitments—commitments that he or she would be prepared to defend in detail. Furthermore, there are no hidden sources of interference of the sort that non-experts could recognize as undermining agency. There are no mad scientists manipulating the subject's brain, no rampant self-deception, and so forth. Is it really plausible to suppose that all the non-expert judgments could be overturned on the basis of the details of the causal chain that connects the belief-desire pair with the behavior? (The discovery that the connection was merely accidental, would be a different story.) We normally know which are our actions without knowing anything about the relevant

neurophysiology, and, in fact, the Davidson climber example would ordinarily be understood as just such a case.

This objection to an appeal to the neurophysiological details of the causal chains connecting belief-desire pairs to behavioral outputs suggests another strategy for characterizing action. Rather than making the characterization of action depend on neurophysiological facts, we might make it depend on facts that are accessible to the agent. In particular, we might require that the causal chains terminating in behavior go through some appropriate piece of practical reasoning in order for the behavior to count as action. But this seems open to a fundamental objection. Consider weak-willed action. At a party, Smith sees the waiter approaching with a Martini tray. He recalls that he will have a long day of important work before him in the morning and that he has had four Martinis already. He reasons correctly that, all things considered, it would be best not to take another drink. Nevertheless, when the tray arrives, he reaches out and takes his fifth Martini.

We ask Smith what happened. Did he forget the salient points in his piece of practical reasoning and reason again on the basis of different premises—for example that he had nothing to do the next day and that there was no reason to refrain from another drink? Or did he entertain two logically incompatible beliefs at the time of the action—that, all things considered, it was best to abstain and that, all things considered, it was best to have the Martini? Smith is indignant. He is not he assures us, despite his four Martinis, lacking in either lucidity or any kind of reasoning ability. He had, he claims, all the relevant beliefs—about the importance of his work the next day, about the debilitating effects of a fifth Martini, etc.—consciously and unwaveringly in mind as he reached for the glass. Nor, he claims, was he guilty of any irrationality of the sort involved in consciously holding two logically incompatible beliefs at the same time about what was best all things considered. The action was, he claims, merely an instance of *practical* irrationality—a weak-willed action, pure and simple.

Weakness of the will thus understood seems entirely possible—indeed, to some of us, all too obviously actual. But now the question arises: What makes Smith's taking the Martini an *action*? And the difficulty here is acute because, far from being supported by an appropriate piece of practical reasoning, this action flies in the face of the relevant practical argument. And again, mere causation by an appropriate desire, even one's strongest relevant desire, cannot by itself make the movement of the arm an action. Such causation is compatible with the arm moving, like Dr. Strangelove's, as though with a will of its own. And certainly such alienated movement is incompatible with action.

Indeed, we can imagine an example of an alienated subject as an extension of the passive subject example. Suppose that while action remains incomprehensible to the passive subject, his beliefs and desires come to cause appropriate bodily movements. This is compatible, as we have seen, with his being completely alienated from the movements—not because he disvalues their consequences, but because they are given as though they were happenings on a par with knee jerks and involuntary twitches. Even if such movements happened (accidentally as far as the passive subject is concerned) to satisfy his strongest desires, he could experience such events while still claiming sincerely not to find action intelligible. In such a case he would refuse to acknowledge such events as actions, and we, it seems, would have to concur. Furthermore, even adding

privileged knowledge, including privileged knowledge in advance, of the movements in question would not be enough. This is clear, first of all, because such privileged knowledge is still something the passive subject could have while still finding the concept of action unintelligible. In addition, however, the passive subject would have special, privileged knowledge of what bodily movements would occur only under their descriptions as bodily movements, whereas the privileged perspective of the normal agent would involve action descriptions. In fact, the normal agent might be almost as ignorant of the movements described as movements or behavior as the passive subject is about them described as actions.

What is missing, both for the alienated subject and for the weak-willed subject, as was the case for the passive subject, is something perceptual. The argument for this claim that was given in the case of the passive subject applies to the alienated subject, who, like the original passive subject, fails to find the concept of action intelligible. Certainly the perception of affordances or the direct perception of opportunities goes a long way toward eliminating the most blatant forms of alienated behavior. And it helps to ground the agential vocabulary for the passive and alienated subjects who understand its inferential role, but, nonetheless, fail to grasp its meaning. But what the example of weakness of the will suggests is that this is not enough. For what seems to be missing is any sense of the *teleological structure* of action. This includes the fact that we aim at an outcome as a goal, rather than being pushed toward it from behind, and that the agent's "take" on the object at which he or she aims rationalizes, rationally motivates, and makes intelligible the agent's aiming at it. In the case of weakness of the will, there is no possibility of such a "take" being given in linguistic terms. For the agent's considered belief, linguistically expressed, is that the qualities of the Martini do *not* make aiming at it rational, rationally motivated, or in the relevant sense intelligible. We have to suppose that such a take that *does* rationally motivate and make intelligible the subject's action will be given not in linguistic, but in perceptual terms. This again is an argument that agency requires a rich phenomenology of perception.

That this requirement that the teleological structure necessary to rationalize (rationally motivate, make intelligible) action is not completely captured by the perception of affordances leads to the third argument. Suppose we ask how the agent's desires are given to him or her. We can see immediately that an objective representation, say as a state of the brain, will not do. Represented in this way, there is no necessary connection between and agent's having (and being aware of having) desires and being rationally motivated to satisfy them. Suppose, for example, that you were told by a neurophysiologist that your strongest desire (not yet, but soon to be, manifest to you) was for great wealth. Then, given that you were planning an academic career, a perfectly rational response would be to try to get rid of it, rather than making any plans to increase the likelihood of its satisfaction. And, it seems, *no* characterization that focuses on the desire will reveal any internal connection between the desire and rational motivation. Hence no such characterization could show how the presence of a desire could rationalize a subject's action, and so no such characterization could reveal its teleological character. What is necessary, it seems, is that we ask how the *object* of the desire is given to the subject when it is given "through" the desire—how it is given when it is given as the desire presents it. This, then, is what it seems we have to ask in order to answer the question of how an object must be presented to the subject in order to be intelligible as

the object of an action aimed at securing it. (And such action is, as we have seen, to be contrasted with mere behavior—even behavior caused by an appropriate belief and desire—that has the same effect)

The only answer, it seems, that captures the teleological character of action as it is experienced by the subject is that the object or goal of the action is given as in some sense valuable or desirable. What this means in most cases is that the object is given in some specific way that entails its value or desirability. Such an object will be given, e.g., as beautiful, fascinating, delicious, captivating, virtuous, compelling, right, morally right, cool, hot, perfect, alluring, exemplary, intoxicating etc. It will not, then, be given as merely desired, or as desired by oneself. The experience of the object must have, to use John Mackie's term, "to-be-pursuedness" built into it (Mackie 1977, p. 40). The reason is, as we have seen, that desires by themselves don't automatically or necessarily make action aimed at their objects intelligible. Consider again the person who is told that his strongest desire, which is as yet not manifest to him, is for great wealth. Suppose he is aware that the desire can be removed safely, effectively, at minimal cost, and without unwanted side effects. If, nonetheless, he immediately begins applying to business schools, he cannot explain this action by saying that he has it on the best authority that great wealth is what he most strongly desires. Given the alternatives, we want to know what makes wealth *desirable*—and we won't accept an explanation of the action until its object is portrayed as having some "desirability feature."

We are now in a position to see the upshot of this treatment of action for the analysis of weakness of the will and the implications of that analysis for both parts of the original transcendental argument: the argument for a phenomenology of experience and the argument for a distinctive phenomenology of agency. In the case of weakness of the will, the desirability feature is completely missing from the subject's linguistically expressed thinking and reasoning. Therefore, the weak-willed action can only genuinely be an action on the basis of some desirability feature given directly in the agent's perceptual experience. In the case of Smith, the Martini is given in perception as manifesting some feature relevantly similar to those listed—a fact that would normally be elicited in a subtle and detailed description of the Smith's subjective experience.

The intelligibility of agency, then, presupposes a rich phenomenology of agency. That is, it presupposes one that includes not simply the direct perception of affordances, opportunities, action possibilities and the like, but also the direct perception of a wide range of desirability features. And this perceptual treatment of action should defuse the worry that what looked like an agential treatment of perception would leave out the facts about what it is like to have the perceptual experience. Neither analysis is more basic than the other, and, in fact, basic action and direct perception go hand-in-hand. Roughly, basic actions require that their objects be given in a direct and rich perceptual experience, and direct perceptions are those that are intimately tied to appropriate packages of basic action possibilities. Furthermore, as we have seen, meaningful language presupposes the possibility of grounding subject to Frege's constraint. Since grounding involves a connection between words and the world unmediated by further language, such grounding requires perception. The perception required, however, is perception whose content is specified neither in linguistic terms, nor in terms of sense-data. For the first

would generate an infinite regress and preclude grounding, and the second, as we have seen, leads to epistemological skepticism and its meaning-skeptical counterpart.

We need, then, access to the world that is direct in the sense that it is unmediated by any mode of presentation that can be fully characterized independently of the object presented. In the absence of such access, the Humean argument for skepticism goes through. This access, however, must, nonetheless, be perspectival. And the agential treatment of perception seems to have just these features. Different modes of presentation involve different action possibilities, so the differences are clearly intentional and perspectival differences. But recall that the phenomenology of agency and perception is both inflationary and *deflationary* relative to the strict sense-datum theory. Often, when we act directly on an object, there is *no* individuating characterization independent of the object itself. Thus had the object not existed, there would have been no complete thoughts or intentional contents among or associated with the relevant states—i.e., the intentional states in which the object is given directly.

Consider again, for example, that we may see someone's distress in the person's face without having access to any features of the face that would be sufficient to distinguish distress from an enormous range of other possible states and emotions. Nor need we have access to any features that would explain (at the personal level) why we immediately respond to distress rather than to any of those other possibilities. If the person in question isn't really distressed, there is no way of capturing, either for an observer, or for the subject himself or herself, the content of the experience. "I was having a perceptual experience in which it seemed that so-and-so was distressed" won't do, because we can give no account of what the experience of the person's seeming to be distressed was like.

Here is another example. Imagine that in the Matrix Neo has an experience that he would describe as seeing a woman in red. Suppose that he says in the course of it, "That woman is wearing red." Then imagine that he comes out of the Matrix and realizes that there was no such woman. Can he fully characterize the experience or the content of the experience he had inside the matrix from his new perspective outside? It seems not. Certainly he could say "It seemed to me that I was seeing a woman in red." But 'a woman in red' clearly doesn't capture the singularity and particularity of the experience, which is, we may assume, irreducibly demonstrative (See Austin 1990). Also, one could have the demonstrative experience that Neo has in the Matrix, without believing the content "It seems to me that I am seeing a woman in red." The reason is that one could have the perceptual experience without having the conceptual sophistication required to express a content that is neutral as between a real seeing and an apparent seeing. In fact, it seems that one could have the perceptual experience without having the conceptual sophistication to express a content about an apparent seeing—indeed, one might not need anything like a full-blown concept of perception at all.

As a final example of the directness and transparency of most perceptual experience, imagine that you are operating a robotic device in an alien environment (under water, inside a volcano, on another planet, etc.). Images on your screen are not photographic, but are computer generated on the basis of information from light, heat, chemical analyses, and so forth, and look nothing like images of natural features with which you are familiar. You learn to manoeuvre the robot in order to complete the

assigned tasks, while avoiding the most serious risks of damage to the machine itself, and initially you do so by memorizing associations between two-dimensional images on the screen and real features of the robot's environment. Quickly, however, you begin to recognize the real environmental features directly, without conscious awareness of the 2-D images or any inferences based upon them. If you are then asked to describe what you are doing, you will say that you are traversing a certain area, avoiding certain obstacles, monitoring the levels of certain dangerous chemicals, and so forth. And if it turns out that there is no such environment, you will be at a loss to provide an adequate characterization of your experience. (Compare: How long was the parking space the last time you parked? "Just long enough to accommodate the car." Now imagine that what you took to be your driving experience was one long nightmare—perhaps not that difficult in Boston.) It seems, then, that in the absence of the objects of such direct perceptions, there is no way to capture the contents of such experiences, where contents are (as usual) individuated in terms of cognitive significance.

Such contents may be called *de re*. And understanding them in terms of an inflationary/deflationary phenomenology, we can see how contents could be both *de re* or direct and perspectival. There may be no way of characterizing the contents of our experiences when we operate in the alien environment that is independent of the existence of the environment itself. But that is not to say that that content is adequately captured by a singular proposition, where such a proposition is represented, if we focus on the subject place alone, as an ordered pair of an object and a property that we ascribe to it. The reason is that in operating in the alien environment, one may see the same obstacle from two perspectives (once directly and once in a mirror, say) and fail to recognize it as the same. (This is the demonstrative version of Frege's problem again.) Moreover, what we might call our perceptual-agential perspectives and capacities understood in terms of the inflationary/deflationary phenomenology have exactly the combination of features in question. They are direct or *de re*, and they are perspectival. Hence they are the appropriate intentional entities to ground our capacity for meaningful language, subject, as always, to Frege's constraint. Such a capacity for meaningful language presupposes them, but such perspectives and capacities per se presuppose no such linguistic capacity. (Of course, since on this view much or most of our actual perception will be theory laden, many of our *actual* perceptual-agential capacities will have just such a linguistic presupposition.) Let us refer to this as "the primitiveness of *de re* thought."

We are now in a position to see how the transcendental argument is concluded. If the meaningfulness of our language presupposes such *de re* thoughts or attitudes, then there is no possibility of our stepping back from all such attitudes simultaneously—of our putting them all into doubt simultaneously and entertaining the possibility that they are the illusory products of a coherent dream or an evil demon. For there is no place to step back to, no way to describe our experience in noncommittal, external-world-neutral terms. If we are to have a meaningful language, we must be talking about such things as tables and chairs, even to characterize that subjective experience necessary to make our words something more than elements in a formal calculus. And the generalized Quinean criterion of ontological commitment that says, roughly, that we are committed to the existence of those things about which we must talk in order to say the things we are committed to saying entails that we cannot be neutral about all such external objects

simultaneously. Thus we have our transcendental argument against Humean skepticism about the external world. And the rich phenomenology of agency—more generally the inflationary/deflationary phenomenology—plays an ineliminable role.

6. Conclusion.

The standard objection to transcendental arguments—that even if they work they show at most that we must use certain concepts, or that if we use some we must use others—is widely assumed to be conclusive. Transcendental arguments, it is supposed, though they might tell us something important about ourselves, could never tell us anything about the external world (Stroud 1999). But whether or not this objection is effective against transcendental arguments such as those of Kant, Strawson, or Davidson (Strawson 1959; Davidson 2001), it is not one that can be made against the argument I have given. For the objection itself presupposes a division between, on the one hand, a subjective domain that is external-world neutral (a domain that includes our “uses” of words and concepts), and, on the other, the external world. And this is precisely the division that the transcendental argument—in particular the conclusion that the primitiveness of *de re* thought is presupposed by the meaningfulness of our language—shows to be incoherent. Moreover, this demonstration of the incoherence of the domain-counter-domain division that the skeptic needs does exactly what transcendental arguments have always purported to do. It shows that something presupposed by the skeptic’s argument undermines the possibility of his or her employing a meaningful language.

But how does this show that we aren’t, for example, brains in vats or subjects in the Matrix? The answer is that it doesn’t. It forces us instead to reconceptualize the nature of such skeptical hypotheses. Or, better, it forces us to recognize that the only meaningful hypotheses in the offing are the real ones. That is, the meaningful hypotheses are the ones for which we could have real empirical evidence and the ones that, for lack of such evidence, we could be justified in giving very little credence. Can we really make sense, though, of the idea of empirical evidence for, say, a virtual reality hypothesis? In *Vanilla Sky* the character played by Tom Cruise is approached by a stranger, who displays uncanny knowledge and unexplained powers—knowledge to which no one besides the Cruise character could have access and powers that would be supernatural on the hypothesis that his experience was of an ordinary external world. This is surely *some* evidence against the ordinary world hypothesis, and we can easily imagine further evidence (forthcoming in the film) that would make the virtual reality hypothesis more attractive than a radical revision of the laws of physics.

It follows from the transcendental argument that Humean skepticism can be ruled out on philosophical or *a priori* grounds as presupposing a false philosophical theory of perception. The Humean argument for skepticism about the external world does not, as we have seen, presuppose a sense-datum theory. But it does presuppose that we have access to a domain, or the possibility of a stance, that is external-world neutral and within which we can understand the skeptic’s argument. And if our having a meaningful language means we must in fact have direct perception of the world—have perceptual states of a kind that if their (external-world) objects don’t exist there is no complete intentional content—then no such domain exists. On the account derived from the transcendental argument, our experience involves commitments that have taken us all the

way out to the world already. And the fact that some such experience is presupposed by the meaningfulness of our language implies that there is no stepping back as far as the skeptic's argument requires.

There is a sense, then, in which the skeptic's hypotheses are incoherent. But there is no mystery as to why it seems that the skeptic is saying something perfectly intelligible. First, many philosophical theses that are demonstrably false on purely conceptual grounds are, nonetheless, intelligible in any ordinary sense. And, indeed, many such theses have a perennial appeal. (Example: all actions are aimed at the satisfaction, pleasure, or happiness of the agent.)

Second, virtually every skeptical scenario understood as making a philosophical or Humean point has a corresponding real scenario expressed in exactly the same words. Imagine that you are about to enter the Library when one of your students runs up, says, "Professor, you may be under the influence of an hallucinogenic drug," and rushes off without explanation. We can easily hear this remark in two ways. (1) The student is simply pursuing, outside of class, the thread of an earlier philosophical discussion with somewhat greater than average enthusiasm. (2) He has just come from a meeting with the Dean in which his roommate has been accused of spiking the coffee in the Faculty Club. The possibility of such real-world interpretations of standard skeptical scenarios suggests again that we could, as I have argued, be brains in vats or virtual reality subjects (though those who would consider seriously whether their beliefs leave them out of touch with the world would be better advised to read Chomsky on the media than Hume on perception). If so, the fact that we seem to understand the skeptic poses no problem for the view I am maintaining.

Finally, that anti-skeptical strategies based on claims about meaning count as senseless things we seem to understand is tied closely to another fact about them—that they have a long history of having proved too much. In failing to make the distinction between philosophical or transcendental skepticism and real skepticism, even such transcendental arguments as Davidson's that most of our beliefs must be true (Davidson 2001) have seemed to rule out real possibilities. And they have seemed to do so merely through reflection on the internal character of our experience. Thus transcendental arguments have seemed to divide between those that are too ambitious to be believed and those that—since they tell us nothing substantive and non-psychological—are (if the truth be told) too modest to be very interesting. The distinction between philosophical and real skepticism, however, reveals a third possibility: that transcendental arguments might defeat skeptical hypotheses understood as philosophical, while remaining silent on their phonological and orthographic counterparts understood as real. Such arguments need not be thought to do the impossible—to give us substantive, non-psychological conclusions on the basis of the internal character of our experience alone. Nor, however, are they so modest as to tell us only that if in fact we think some things, we must think others. Such arguments rule out as illegitimate what we might otherwise be tempted to say: "Of course we have commitments regarding the external world—but that's just us!" (I would not want to be accused of sloganeering, but a useful mnemonic aid might be "transcendental arguments against transcendental skepticism, empirical arguments against real skepticism.")

This understanding of transcendental arguments is, as we have seen, grounded in the thesis of the primacy of *de re* thought, which is, in turn, grounded in the inflationary/deflationary phenomenology. And this understanding is relevant to the standard objection to coherentism—that a perfectly coherent set of beliefs and experiences might amount to nothing more than a coherent dream (alternatively, that such beliefs and experiences could “float free” of any connection with reality, or that nothing could justify the assumption that coherence and truth are connected). Another way to put the point is to say that coherentist theories involve an infinite regress. Something is justified by its coherence with something else only if the latter is justified. But if all we can point to in the case of any belief is that it coheres with some others, then it seems that we could never regard any belief as justified at all.

Again, however, these arguments presuppose exactly the picture that I have argued against—that of a domain of belief and experience that involves no external-world commitments. If we suppose, as I have argued we should, that as a matter of conceptual necessity we cannot ever avoid such commitments, the situation changes. We are not, as the objection suggests, trying to justify some beliefs in terms of their coherence with beliefs that themselves are equally open to question. Rather, we are trying to justify those beliefs whose truth and justification are open to doubt by reference to their coherence with beliefs that, in context in question, are *not* up for grabs—beliefs that represent, at least in that context, our commitments.

There is, then, in this answer to the standard objection to coherentism, a foundationalist element. But it does not yield the absolute foundationalism that sense-datum theories have always been thought to support. It is rather what we might call *context-relative foundationalism*. There are no empirical propositions that are immune to either doubt or scrutiny. But such scrutiny, involving as it must the use of a meaningful language, can only take place against a background of beliefs that are fixed—and beliefs that, in virtue of the capacities we must have for direct perception and basic action, take us all the way to the world. Such beliefs thereby give us what we must have as well: external-world commitments. And since such direct perception and basic action presuppose a rich and distinctive phenomenology of action and perception, I conclude that such a phenomenology has the transcendental significance I have claimed.⁴

References

- Austin, D. (1990). *What's the Meaning of "This"?: A Puzzle about Demonstrative Belief*. Cornell University Press.
- Davidson, D. (1973). Radical Interpretation. *Dialectica* 27, 313-28.
- Davidson, D. (1983). A Coherence Theory of Truth and Knowledge. In D. Henrich (ed.), *Kant oder Hegel*. Klett-Cotta, Stuttgart. Reprinted in D. Davidson, *Subjective, Intersubjective, Objective*. Clarendon Press, 2001.
- Evans, G. (1982). *The Varieties of Reference*. Oxford University Press.
- Gibson, J. (1986). *The Ecological Approach to Visual Perception*. Laurence Erlbaum.

- Hume, D. (1968). *A Treatise of Human Nature*. Clarendon Press.
- Hume, D. (2001). *An Enquiry concerning Human Understanding*. Open Court, Chicago.
- Lewis, D. (1974). Radical Interpretation. *Synthese* 23, 331-44.
- Mackie, J. (1977). *Ethics: Inventing Right and Wrong*. Penguin.
- Michotte, A. (1963). *The Perception of Causality*. Basic Books.
- Schiffer, S. (1978). The Basis of Reference. *Erkenntnis* 13, 171-206.
- Searle, J. (1983). The Intentionality of Perception. In J. Searle, *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.
- Strawson, P. (1959). *Individuals: An Essay in Descriptive Metaphysics*. Doubleday.
- Stroud, B. (1984). *The Significance of Philosophical Skepticism*. Clarendon Press.
- Stroud, B. (1999). The Goal of Transcendental Arguments. In R. Stern (ed.), *Transcendental Arguments: Problems and Prospects*. Clarendon Press.
- White, S. (2004a). Skepticism, Deflation, and the Rediscovery of the Self. *The Monist* 87, 275-298.
- White, S. (2004b). Subjectivity and the Agential Perspective. In M. De Caro and D. Macarthur (eds.), *Naturalism in Question*. Harvard University Press.
- White, S. (2007). Property Dualism, Phenomenal Concepts, and the Semantic Premise. In T. Alter and S. Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge*. Oxford University Press.
- Williams, M. (1996). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Princeton University Press.

Notes

1. Those whom Michael Williams calls “new Humeans” are committed to our being able to do so. Williams includes P. F. Strawson, Barry Stroud, Thomas Nagel, and, possibly, Stanley Cavell in this category (Williams 1996, Ch. 1).
2. For suitably characterized sets of propositions corresponding to the domain and counter-domain. Since I shall be arguing on other grounds that this example cannot generate Humean skepticism, the complications in applying the independence condition needn’t concern us here.
3. For his doubts regarding the first assumption, see Stroud 1999, pp. 160-163. For his reservations on the second score, see Stroud 1984, Chs. IV and V.
4. I am grateful to Massimo Dell’Utri for stimulating discussions of these issues and to Ned Block and Susanna Siegel for comments and suggestions on an earlier draft.