

ON THE ABSENCE OF AN INTERFACE: PUTNAM, DIRECT PERCEPTION, AND FREGE'S CONSTRAINT¹

STEPHEN L. WHITE

Department of Philosophy, Tufts University

ABSTRACT

Hilary Putnam and John McDowell have each argued against representational realist theories of perception and in favor of direct realist (or “common-sense realist”) alternatives. I claim that in both cases they beg the question against their representational realist opponents. Moreover, in neither case has any alternative been offered to the representational realist position where the solution to perceptual or demonstrative versions of Frege’s problem is concerned. In this paper I present a transcendental argument that some of our perceptions of external objects must be direct in the sense that we perceive them and there is nothing else we perceive in virtue of which we do so. I also present a reply to standard objections to the claim that transcendental arguments can be used to support conclusions about the world and not simply about our own use of concepts. Finally, I present a theory in terms of which the relevant Frege problems can be solved without appeal to any of the sorts of representations in terms of which representational realism is defined.

Keywords: Hilary Putnam, perception, representation, Frege’s Problem, transcendental argument

1. Is the notion of an interface problematic?

Hilary Putnam, echoing John McDowell, has denied that “there has to be an interface between our cognitive powers and the external world.”² And Putnam has denied as well (the same claim differently expressed) that “our cognitive powers cannot reach all the way to the objects themselves.”³

What, though, does it mean to say that there is no interface, no boundary, between what is internal and external to a subject? Certainly we can stipulate such a boundary—at the surface of the brain, say—or, for particular purposes, around some of its relevant functional sub-units. And if our cognitive powers are thought to be identical to, reducible to, or super-

¹ Some of the arguments presented here are discussed at greater length in White (2007). I am grateful to Akeel Bilgrami, Ned Block, Mario De Caro, Massimo Dell’Utri, David Macarthur, Roberto Pujia, and Susanna Siegel for their comments and suggestions on these topics.

² Putnam 1999, p. 10; McDowell 1994.

³ Putnam 1999, p. 10.

venient upon states or properties of the brain, and if we are considering the question from the objective or third person point of view, then it seems obvious that such an interface exists.

That there is such an interface is not ruled out, even if we hold, as Putnam does, that our cognitive powers are a matter of intentional states and that the contents of such states reach all the way to objects in the world. It is not ruled out, that is, even if we hold, for example, that the contents of beliefs depend on the character of our external environment, i.e., that they have “broad content.” In Putnam’s famous thought experiment, everything on Earth has a molecule-for-molecule duplicate on Twin Earth except water.⁴ On Twin Earth, what falls as “rain,” flows from the faucets, and fills the lakes and reservoirs is a completely different substance from water (i.e., H₂O). What plays the role of water on Twin Earth, though it has the same macro-level properties as water, has a completely different chemical formula (which we abbreviate as XYZ). According to Putnam, subjects on Earth and their Twin Earth duplicates refer to different substances in using the word ‘water’—H₂O and XYZ respectively—solely because of the differences in their external surroundings. But even if we grant that such “broad content” exists, this does not rule out the possibility of there being narrow content as well. That is, granting that our beliefs have broad content that does not supervene on what is inside the subject’s head does not preclude their having a kind of content that does. And narrow content would give us precisely the interface between, on the one hand, beliefs, perceptual states, and so forth understood as internal to the subject and, on the other, the subject’s external environment. Nor does granting the existence of broad content rule out the possibility of our “factoring” such content into a narrow content plus an external causal chain.

Putnam is, of course, well aware of this possibility and has himself offered a possible characterization of narrow content. But in *The Threefold Cord*,⁵ Putnam argues on a number of grounds that no such account is viable. I find these arguments inconclusive or, perhaps it would be more accurate to say, incomplete.

To see the problem, consider the conception of narrow content that I put forward in “Partial Character and the Language of Thought.”⁵ On this conception, the meaning or intension of a term is treated as a two-place function defined for a domain of ordered pairs of possible worlds. The first world of such a pair is a possible context of acquisition of the term and the second world a possible context of evaluation. The function maps such an ordered pair into the extension of the term at the second world when it is acquired at the first. If the term ‘water’, for example, is acquired on Earth, then it picks out bodies of H₂O (if it picks out anything) on Earth, on Twin Earth, and on every other possible world at which we evaluate its extension. Acquired at Twin Earth, it picks out bodies of XYZ at every such possible world. And whereas the meaning of

⁴ Putnam 1975.

⁵ White 1982.

'water' construed as a function from possible worlds to extensions is different for an Earthling and his or her Twin Earth duplicate, the two-place function will be identical for such pairs of subjects. Indeed, it will be identical for pairs of subjects whose relevant functional makeup is "sufficiently similar." Meaning in this sense, then, supervenes on functional makeup, hence on what is "in the head."

It is not my aim to defend the conception of content of a very distant self, a conception that I now reject. And, indeed, Putnam considers a similar view and dismisses it for what I regard as the right reason—that it is parasitic on an account of broad content. But that it is parasitic on such an account was an explicit feature of this conception of content. The factors that determine the association between a term and, say, a natural kind for a particular context of acquisition were like the reference determining factors that figure in typical causal theories. And Putnam provides no argument that narrow content should *not* presuppose broad content in this way. I believe that such an argument is available, but it can only be stated when we have in hand a fully articulated account of the problems raised by Fregean considerations, particularly in demonstrative contexts. This is a topic that neither Putnam nor McDowell takes up, and I shall postpone its discussion until after we consider Frege's constraint.

That there is no interface between our cognitive powers and the external world, then, is far from obvious. McDowell seems to argue that there cannot be one because the existence of such an interface would lead to epistemological skepticism about the external world. But this argument is also seriously incomplete. Many philosophers hold what Michael Williams has dubbed the "new Humean" position. According to this position (held, according to Williams, by Barry Stroud, P. F. Strawson, and Thomas Nagel, among others), we *cannot* answer the skeptic about the external world, but this is of no practical significance.⁶ Thus there should be an *argument* that a view that entails epistemological skepticism about the external world is automatically unacceptable, and none has ever been provided.

2. Can we make sense of direct perception?

Putnam's argument against the existence of an interface is clearly meant to be stronger than McDowell's. Putnam suggests in many places that if there is an interface, we can have no account of how language "hooks onto the world." This consequence is clearly related to what I shall call below "meaning skepticism." And an argument that the existence of an interface leads to meaning skepticism is not one that can be easily dismissed. But what would such an argument look like?

⁶ Williams 1996, pp. 10-17.

As we have seen, not every conception of an interface seems problematic. And this is true even for McDowell's (anti-skeptical) purposes. In fact, it seems, there is only one route from the existence of an interface to epistemological skepticism about the external world. This is the route that proceeds from the first person subjective perspective, or the Cartesian perspective, via Hume's argument.⁷ According to this argument, if one focuses on the character of one's present perceptual experience (an experience, say, as of a room, furniture, etc.), one will recognize that there are other possible explanations of the experience than the one suggested by ordinary common sense. One might, for example, be dreaming, be a brain in a vat, be the victim of an evil demon, or be a subject in a virtual reality setup. There is, then, a logical gap between what one is given in experience and its causal source in the external world. And this gap can only be bridged by an inference. As Hume argues, however, an *a priori* inference couldn't be justified, since we cannot argue *a priori* from effects to causes. Nor, however, could an *a posteriori* inference be justified. Any such inference would have to be grounded in an *a posteriori* principle or generalization connecting the character of our perceptual experiences with features of the external world. And the use of any such principle would clearly beg the question against the skeptic. The conclusion, then, is not simply that we cannot ever *know* the truth of any *a posteriori* proposition about the external world. It is the much stronger claim that we could never be *rationally justified in believing* such a proposition. In other words, we could never have rational grounds for preferring one hypothesis about the external world over any other.

This conclusion is obviously very strong and is, I think, more disturbing (or should be) than new Humeans acknowledge. I am not concerned to argue this claim, however, because, as I have indicated, I believe, (as I think Putnam does) that the real issue is not epistemological skepticism but meaning skepticism. The present point is that we now have a clear conception of the interface that is at least problematic. Suppose our concepts can be factored into components that are "external-world-neutral" and components consisting of mere causal chains to the external world. And suppose that, as a result, everything that is given to the mind—everything available from the first person point of view—can be understood in external-world-neutral terms. Then epistemological skepticism about the external world follows. Suppose on the other hand that we could make sense of the idea of our being given external objects directly, in a way that left no logical gap between what we were given and the existence of such an object. Then no inference from what is given in perception would be necessary, and the Humean argument for epistemological skepticism would be defeated.

Can we make sense of such a notion? I believe that the notion of direct perception is, properly understood, unproblematic. Consider the notion of a basic action. An action is basic if we do something and there is nothing else that we *do* in virtue of which this

⁷ Hume 1988, Section XII, Part I, p. 184; Hume 1896, Book I, Part IV, Section II, p. 212; cf. Stroud 1984, pp. 545-551, reprinted in Alcock 1998, see especially pp. 363-364.

is the case. This is not to suppose that there are no causal antecedents of our basic actions. It is simply to say that at the level of *action* there is nothing more to be said. By analogy, then, a case of direct perception is a case in which we perceive something, and there is nothing we *perceive* in virtue of which this is so. Again this is not to deny that there is a causal chain connecting the perceptual experience to things in the world. It is simply to say that an analysis in terms of *perception* can go no further.

But so far all we have is a conditional. *If* it is to be claimed that there is no interface, then the claim must (if it is to be plausible) be understood in terms of direct perception characterized by analogy with basic action. And, understood in this way, the claim that we perceive external objects directly has a desirable consequence—it provides an answer to the Humean skeptic. Moreover, it is clear that there needn't be an interface understood in the present sense—the sense that gives rise to Humean skepticism—just because, understood in another way—from the objective perspective—there obviously *is* an interface. But this falls short of an argument that there *is* no interface in the sense ruled out by the claim that we sometimes perceive things directly. Is there an argument for this latter claim?

I believe there is. I believe that there is a transcendental argument that *de re* (or object-involving, or Russellian) thought is necessary to our having a meaningful language and that such object-involving thought is irreducible to thought that is non-object involving. (To say that some thought or intentional content is object involving is to say that if the object or objects in question do not exist there is no complete thought; there is nothing that could be evaluated for truth, accuracy, veridicality, or the like.) The view that we must have (some) irreducibly *de re* thought has a claim to being anti-Cartesian and anti-naturalist. It is anti-Cartesian because Descartes' own skeptical arguments require a conception of what we are given in perception as "external-world neutral," even if they do not require a sense-datum conception. It is anti-naturalist since if we had an objective-causal reduction of meaning, then the partial character view would show how we could factor our intentional states into internal and external components.

3. The transcendental argument for irreducibly *de re* thought

The argument is as follows.

1. Language must be grounded subject to Frege's constraint

Frege's constraint says, in effect, that if a subject has two beliefs regarding an object, beliefs that, logically, cannot be true of one and the same thing, then there must be two distinct modes of presentation of that object under which the beliefs are held.⁸ And modes of presentation are precisely the sorts of things that have to be given to subjects

⁸ For Stephen Schiffer's version of Frege's constraint, see Schiffer 1978. Similar versions occur in the work of Gareth Evans, Brian Loar, and Christopher Peacocke. Interestingly, though the notion of object-involving thought figures prominently in the work of John McDowell, Frege's constraint does not.

in the form of conscious experience, not external causal chains. Putnam attributes to William James the claim that “the traditional claim that we must conceive of our sensory experiences as *intermediaries* between us and the world has no sound arguments to support it.”⁹ But if it was James’ view that there was no serious argument for such a conception of experience that any account of these issues would have to address, it was false. The argument stems from the requirement that language be grounded subject to Frege’s constraint.

With a minimal amount of charity we can credit Russell with having been moved by the following considerations in favor of his sense-datum theory, though not with having articulated them explicitly.¹⁰ (And, of course, whether the following argument can be attributed to Russell is irrelevant to its cogency and validity.) The argument is that in addition to inferential roles or word-to-word connections, our linguistic expressions must have some direct connections to nonlinguistic reality. If there were no such word-to-world connections, we would have no more than an uninterpreted formal calculus. And if the connections were not direct (unmediated by any descriptive linguistic content), we would have an infinite backward regress of word-to-word connections. In the absence of such a direct connection to the world, our language would be ungrounded. We must have, then, in addition to lexical definitions, ostensive definitions in which we pick out items in the world directly and in virtue of which our language “hooks onto the world.” The contemporary notion that this is a matter of external causal chains satisfies the requirement of grounding but fails in the second regard—that the grounding be subject to Frege’s constraint. The reason is that such causal chains needn’t be, and normally are not, available to the subject. Hence they can’t provide the modes of presentation that rationalize otherwise incompatible beliefs about the same object. And this is why the partial character conception of narrow content, presupposing as it does (something like) a causal theory of reference, must be rejected.

We do justice to the subject’s rationality, then, only by postulating “modes of presentation” that are available to the subject and that are not modes of presentation of the same object *a priori*. In a case like that of the morning star and the evening star, it is plausible to assume that these modes of presentation will be descriptive. But if all such modes of presentation were descriptive, we would be back to the infinite regress of word-to-word connections, and our language would be ungrounded in any connection to the world. The challenge, then, is to solve the Frege problems in the demonstrative cases. And the question is what the modes of presentation in such cases could be.

It is in this context that Russell’s suggestion that the logically proper names (demonstratives) refer to sense-data looks attractive. Indeed, in this context it can look almost inevitable. Sense-data are nonlinguistic and our access to them via acquaintance is direct in being unmediated by any linguistic-descriptive content. Moreover, sense-data

⁹ Putnam 1999, p. 11.

¹⁰ Russell 1956.

themselves have no hidden sides and are available in principle from only one point of view. Thus they raise no new Frege problems. In Gareth Evans' example of a demonstrative version of Frege's problem, one points out a window to one's right at the bow of a ship and says "that ship was built in Japan."¹¹ One points out a window to one's left at the stern of a ship and says, "that ship was not." And, without realizing it, one has pointed to the same ship twice and said of it logically incompatible things. The problem is solved by saying that the two modes of presentation required to do justice to one's rationality are the two different sets of sense-data associated with the different perspectives on the ship.

2. The attempt to ground a meaningful language in sense-data leads to meaning skepticism

Sense-data solve one problem regarding meaning only by raising another that I shall call the problem of meaning skepticism. If the meanings of our terms like 'table' and 'mountain' are grounded in their connections with sense-data, then the notion of an external object—an object distinct from any pattern (however complex) among our sense-data—will be meaningless. The argument that the attempt to ground language in sense-data leads to meaning skepticism is as follows. Words get their meanings in virtue of their inferential roles (word-to-word connections) and their direct connections to the world (word-to-world connections). How, then, would what purport to be our terms for external objects get their meanings? If it is via their inferential connections with (definitions in terms of) sense-datum terms that are grounded in their direct connections to sense-data, then the conclusion follows. The terms that allegedly pick out external objects will, in fact, pick out nothing over and above patterns and regularities among our sense-data. Such putative external object terms will be definable in terms of, and hence in the strongest sense be reducible to, terms in a sense-datum language. Suppose, then, that they are not so definable or so reducible—that they cannot get their meanings in virtue of their definability in a vocabulary that is itself unproblematically grounded. Then, since we can never in principle be given external objects directly, we can never correlate the terms with the items they purport to pick out. Thus we can never give the terms meaning via acts of ostensive definition, and, it seems, they can never be grounded at all. Thus, whether the external object terms are definable in sense-datum terms or not, it seems that on these assumptions all we can ever talk about meaningfully are patterns and regularities among our sense-data.

Why, though, is there any suggestion here of meaning *skepticism*? The picture of meaning (the one in which external object terms have definitions in a sense-datum language) is just the picture put forward by C. I. Lewis and, more generally, what is definitive of (linguistic) phenomenalism.¹² The answer is that the same arguments that lead to the

¹¹ Evans 1982, p. 84.

¹² Lewis 1946.

conclusion that our external object terms, if they are to be meaningful, must have definitions in a sense-datum language show something stronger. They show that our terms that purport to pick out sense-data in the past or the future, possible or counterfactual sense-data, or the sense-data of others must have definitions in terms of expressions that pick out our own, actual, present sense-data. In this case there seems to be too little even to justify talk of interesting patterns or regularities. Thus it seems we have too thin a “definitional base” for anything we might think of as a genuine language.

How, after all, did the original argument go? We started with an ordinary language with terms that purported to pick out external objects such as tables or mountains. And we argued that such terms were in fact meaningful only if they had translations into a sense-datum language. But what was the basis for this claim? It was that since our access to external objects (as ordinarily understood) cannot be in principle be ostensive or demonstrative (all we can be given directly in perception are sense-data), there is no other way in which such terms could be meaningful. But the same point applies to past and future, possible and counterfactual sense-data, as well as the sense-data of others. These are also things that we cannot be given directly in perception, hence cannot demonstrate. Suppose, then, we ask about the meaningfulness of the terms that the phenomenalist would use in accounting for the meanings of our ordinary terms for external objects. That is, suppose we ask about the meaningfulness of the sense-datum terms themselves. By the phenomenalist’s own arguments, these terms must have definitions or translations in terms that refer only to our own, actual, present sense-data.

It might be thought that this is unfair to phenomenism. For it seems that there is a disanalogy between the two arguments. If the sense-datum intuitions in the context of the necessity of grounding language (subject to Frege’s constraint) are right, then ordinary external objects are simply not the kinds of things we could suppose we were given directly. To suppose they were, it seems, would be to give up trying to satisfy Frege’s constraint and thus to give up on the attempt to do justice to the rationality of subjects. And since there is widespread agreement that rationality is constitutive of the ascription of intentional states, this would be to give up on intentionality and meaning themselves. But, it might be argued, this is not the case where sense-data other than our own, actual, present sense-data are concerned. For these are just the kinds of things it *does* make sense to suppose are given directly. Past sense-data were given directly *in the past*. Future sense-data will be given directly *in the future*. Counterfactual sense-data would be given directly *in appropriate counterfactual circumstances*, the sense-data of others are given directly *to them*, and so forth. Does this get the phenomenalist off the hook?

To suppose that it does is to forget the dialectical context in which sense-data come to seem so attractive. We appeal to sense-data to play a certain role: to provide the modes of presentation necessary to do justice to the rationality of individual subjects. But certainly it would seem absurd to argue that because sense-data are the kinds of things

that are given directly, are non-linguistic, and don't themselves raise Frege problems, we can appeal to the sense-data of other people to play this role in our own case. What rationalize *my* beliefs are the sense-data given to *me*.

Suppose I have irrational beliefs about a given object O. Two of my beliefs about O are contradictory, but I fail to notice this because they normally come to consciousness in different contexts. And suppose that I recognize O to be the same under all the relevant modes of presentation. Thus, were the contradictory nature of the beliefs pointed out to me, I would not say that I had been ignorant of the fact that two modes of presentation of mine were modes of presentation of a single object. Rather I would say that I had indeed been irrational in holding a pair of contradictory beliefs. Surely, then, I am not (and my beliefs are not) rationalized by the fact that *Smith* has sense-data in virtue of which there is a mode of presentation of O that I would not recognize as such

What this means, in effect, is that it is the notion of acquaintance that allows sense-data to play the role for which they have been slated by the phenomenalist response to the demonstrative versions of Frege's problem. And we can no more be acquainted now with our past sense-data than we can be acquainted with the sense-data of another subject. What rationalizes (or fails to rationalize) our present beliefs about the past is not the totality of our sense-data (past, present, and future) but our present sense-data associated with our present memories. Analogous points are true of our future, possible, and counterfactual sense-data.

It might be thought that we needn't talk of past sense-data, for example, but merely of sense-data *in the past*. But we can no more suppose that we have direct or unmediated access to (direct perception of or acquaintance with) the past than we can, on phenomenalist principles, suppose that we have it to the external world. Indeed, the past seems even less like the sort of thing that could be perceived directly. On such principles, our access to the past is mediated by our present sense-data, primarily the sense-data in virtue of which we have memories. On phenomenalist principles, however, what makes them memories is their (alleged) connection with past sense-data. Hence the phenomenalist needs to be able to explain the meaning of this notion. And unless the notion of a past sense-datum can be defined in terms of patterns and regularities among our present sense-data, it seems that this cannot be done.

Pursuing the analogy with the problem of other minds may help in this context. As is well known, the problem of other minds has been, at least since Wittgenstein, a problem of meaning. What does it *mean* to talk about a sensation such as pain that I myself don't experience? Following Kripke we might see the problem as one of answering the following question. Suppose that the meaning of my term 'pain' is exhausted by my ability to recognize pains when *I* have them. What, then, is there in what I mean that could be detached and applied in the case of another subject?¹³ Any conception such as

¹³ Kripke 1982, pp. 114-145.

functional role or physical realization that could in principle be detached and applied in the case of another seems irrelevant to what I mean—namely, “the subjective feeling of pain (to me)”. And all we have explained where the meaning of this expression is concerned is, of course, what it feels like to me when *I* have it. The suggestion now is that the concept of past sense-data of mine is problematic in the same way as the concept of the sense-data of other subjects. And exploiting the analogies between other subjects and our past and future selves is a standard philosophical move—one that has been used to very telling effect (against Hume!) by both Derek Parfit and Henry Sidgwick in the domain of moral psychology.¹⁴ I conclude that the phenomenalist cannot resist the shrinkage of the “definitional base” to his or her own, present, actual sense-data and thus cannot avoid meaning skepticism.

3. Language must be grounded in demonstrative belief, direct perception, and basic action possibilities

If the foregoing is right, then we have an argument for the strong thesis that Putnam’s views suggest—that a conception of the mental that entails the existence of an interface in the problematic sense is incoherent. The alternative is a conception of the mental according to which *de re* or object-involving content is basic and irreducible. The problem, of course, is that we can have two object-involving beliefs—beliefs in singular propositions in the analytic jargon—in virtue of which we believe logically incompatible things of the same object. In other words, the problem is to reconcile direct realism with Frege’s constraint. The account I have suggested in a number of places involves the appeal to what I call packages of basic action possibilities. Consider again Evans’ ship. What are the modes of presentation in virtue of which the incompatible beliefs about the ship do not compromise one’s rationality?

It is tempting to throw up one’s hands and say the following. In such cases we have two distinct direct perceptual experiences of the same object, one in virtue of which we can point to the bow, and one in virtue of which we can point to the stern. That they are distinct shows that their contents cannot be represented as singular propositions (ordered pairs of objects and properties) because, so understood, they would represent the subject as ascribing incompatible properties to the same object. Nor could we add a third element—external-world-neutral modes of presentation (whether sense-data or not). For this would simply be another instance of the factoring of the internal and external that leads to the problematic conception of the interface and meaning skepticism. The alternative seems to be to say that direct perceptual experiences individuated in such a way that we can have different direct experiences of the same object are simply *sui generis*. And this is to say that there is no explanation of *what it is in virtue of which they are different*—hence, no explanation of *how* such experiences could satisfy Frege’s constraint.

¹⁴ Parfit 1984, pp. 117-194; Sidgwick 1981, pp. 418-422.

Desperate situations call for desperate measures, but not this desperate. There is quite a lot that we can say about our direct access to the world. And we can do it without supposing that direct access has an analysis into an internal and external component, hence without supposing that it is reducible to indirect access. The strategy is simply to explain some forms of direct access in terms of others and to avoid the charge of vicious circularity by making the circle large and informative. Evans' ship, for example, presents different packages of basic action possibilities, depending on whether one is looking at the bow or the stern. In either case, one can point, trace the outline in the air, move closer for a better view, aim a laser sight, give directions to a crane operator loading cargo, and do an indefinite range of other things, all as basic actions. That they are basic actions means that there is no external-world-neutral mode of presentation in virtue of which the action possibilities hook onto the ship in question—one simply performs the actions directly on “that ship.” Such actions are, nonetheless, different actions depending on whether one is looking at the vessel's bow or stern. And they clearly help to explain how one could believe incompatible things of the same ship.

There is a sense, then, in which our two modes of presentation of the ship are a matter of know how, and this is crucial to the claim they are object involving. But it might be objected that this attempted reduction of the perceptual to the agential leaves out precisely what is essential to the perceptual states—that there is something it is like to have them. And this thought threatens to reintroduce sense-data and with them the problematic conception of an interface. The objection, however, ignores the reference to holism and to non-vicious circularity and so misconstrues the suggestion. Consider what I have called elsewhere the perfect blindsight example and the example of the passive subject.¹⁵ Imagine first a modification of the well-known empirical example of blindsight. In the original example, subjects with certain kinds of brain lesions report a blind area in their visual fields. The subjects, however, are asked to guess the features of images projected on the part of a screen that falls within the so-called blind area. And many such subjects reliably discriminate a range of simple visual features, while maintaining that they see nothing. Imagine now such subjects becoming better and better in their discriminations until, spontaneously and without prompting, they can make all the same discriminations as a normal subject, as well as performing all the same actions with the same degree of confidence and reliability. Imagine in one's own case driving in downtown Boston, avoiding pedestrians and careless drivers, and driving around new construction and serious potholes, while succeeding in finding suitable parking. Can one really imagine doing this while being, nonetheless, blind? It seems not, particularly in light of the fact that one could reproduce the same sense-datum description of one's visual experience as a normally sighted subject. Asked why one parked in a no-parking zone, one could say that from where one was, the no-parking sign was occluded by a parked truck. And such a description could in an obvious way be restated in sense-datum terms.

¹⁵ For the passive subject see White 2004a. For the passive subject and the perfect blindsight example see White 2004b and White 2007.

This first example suggests that just as the phenomenology of direct perception requires agency, agential capacities, if they are as sophisticated as ours, require some phenomenology of experience. But we can go further: there seems to be a distinctive phenomenology of agency itself. Imagine a subject who awakens one morning and claims, with apparent sincerity, not to understand the concept of action. It seems possible to tell a coherent story in which such a subject believes all the same propositions we do when they are couched in non-agential terms. In addition, such a subject, it seems, might know all the things that it is appropriate to *say* about action. At least this might be so in the sense in which an anthropologist might know everything it is appropriate to say about tribal gods in which he or she does not believe. Such a subject, we might say, has retained the inferential roles of our agential expressions, but finds nothing in the world to motivate or justify their application. So described, what the subject lacks, it seems, must be something perceptual. What is missing, we can suppose, is any perception of the world as presenting opportunities for action. And this suggestion is supported by the research tradition stemming from the work of J. J. Gibson, according to which what we are given in experience is not a world of agentially neutral physical objects and structures, but a world of affordances—of shelters and hiding places, doorways, stairs, bridges, and escape routes.¹⁶ And according to Gibson, these are not descriptive interpretations we impose on a neutral perceptual experience, but a part of the experience itself. Were they missing, as they are in the example of our passive subject, our perceptual experience of the world would be radically different.

The conclusion is that an agential account of the modes of presentation of objects in cases of direct perception does not entail that there is no story to be told about the phenomenology of perceptual experience. It would be as accurate to say that the account of agency is a phenomenological one as it is to say that the account of phenomenology is agential. In both cases the intentional states in question are object involving. The perceptual experiences are direct and the opportunities are opportunities for basic actions. And, it seems, there is a clear sense in which direct perceptions and basic actions are made for each other. Roughly, the things we see directly are the things we can manipulate through our basic actions.

The upshot is what I have described as a phenomenology that is both inflationary and deflationary relative to what the sense-datum theory entails.¹⁷ We are, on the inflationary side, given far more than sense-datum theories, at least of the orthodox sort, allow. Far from being given only apparent shapes, colors, and relative sizes, we are given humanly meaningful objects and features of the world—“threats and promises” in Ryle’s terms. But also, I would argue (though I cannot do so here), we are given objects and aspects of the world that are valuable in specific ways. We are given features and things that are beautiful, strong (in an intellectual and aesthetic as well as a physical sense),

¹⁶ Gibson 1986, pp. 33-44.

¹⁷ See White 2004a, White 2004b, and White 2007.

virtuous, elegant, fascinating, cool, hot, bold, subtle, steadfast, brave, and so forth. On the deflationary side, however, we may fail to see any of the things the sense-datum theorist says we must see in order to see anything else. Nothing could be more common than seeing a friend's distress without noticing or being aware of any of the subtleties of the facial geometry in virtue of which the distress is given. *A fortiori* we can do so without seeing those things the perception of which the sense-datum theorist holds to be basic. Again, there is nothing more mysterious here than there is in the fact that we can see a letter on a page without being aware of any of the details of the typeface in which it is printed. This is not to deny that the shape of the letter is *causally* relevant to our recognition of it. It is merely to deny that, in addition to the letter, the details of the shape must be registered in our conscious awareness. To say that an adequate phenomenology of experience is both inflationary and deflationary relative to the sense-datum theory, then, is to say that our experience is both richer and more transparent than that theory allows. Moreover, these two features go hand in hand. It is precisely because we are not given what the sense-datum theory requires that we can be given our friend's distress directly.

4. Thus the Humean claim that any *a posteriori* belief about the external world would have to be inferred (illegitimately) from what we are given in external-world-neutral perception is false: We must have some object-involving perceptual beliefs in order to have a meaningful language, and such beliefs give us ontological commitments.

Therefore, we cannot intelligibly step back, as the Humean skeptic requires us to do, to a point from which we can call all of our external-world beliefs into question simultaneously.

The criterion of ontological commitment that would support this claim would be some appropriate generalization of the Quinean criterion. Something, for example, like:

We are ontologically committed to the existence of those things about which we have to talk in order to say the things we are committed to saying.

Putnam rejects the Quinean criterion,¹⁸ but expresses sympathy for something like the generalized criterion.¹⁹ If something like this is right, then we have the strongest possible answer to the Humean skeptic. We have a transcendental argument that the skeptic's claims are incompatible with the possibility of our having a meaningful language.

¹⁸ Putnam 1999, p. 179.

¹⁹ Putnam 1999, p. 120.

4. The legitimacy of transcendental arguments

In response to such a claim, two objections will be raised, both of which have been made familiar by Barry Stroud. First, and most obviously, it will be objected that even if such an argument works, it can show at most something about us—that we must apply certain concepts or the like. What will be denied is that such an argument could, in principle, show us anything about reality outside the mind. What I hope the preceding discussion has shown is that this objection to transcendental arguments, long regarded as conclusive, can in fact be dismissed. For the objection presupposes precisely the factoring into external-world-neutral facts about our psychologies and facts about the external world that it has been the burden of that discussion and the transcendental argument it supports to undermine. The transcendental argument, then, provides more than just a conclusion about us—that we have certain commitments and that in principle we cannot step back from them all simultaneously. It also provides a reply to the standard objection to the legitimacy of transcendental arguments themselves. To put the reply another way, one cannot have genuine commitments regarding the external world and say, nonetheless, “but that’s just me.” One cannot meaningfully thump one’s chest but keep one’s fingers crossed behind one’s back.

The second objection is not to transcendental arguments *per se*, but to the implication that what the skeptic says is incoherent. We obviously *do* understand the skeptic’s claim that we might be brains in vats, it will be said, and any suggestion to the contrary will presuppose a theory of meaning that we should reject on that very ground. What possible justification could there be, it might be asked, for the claim that we literally don’t understand *The Matrix*, or *Vanilla Sky*? And it may be added that the only theory of meaning with the required consequence is verificationism, a theory known to be defective on independent grounds.

It should be clear, of course, that the approach being defended is nothing like the classic versions of verificationism, which presuppose exactly the interface being rejected. In fact, it seems that the theory on offer has as much right to be regarded as a form of realism as any other—as indeed the connections between direct perception and direct or naive realism suggest. The real answer to the objection, however, is that it trades on an equivocation. It is indeed difficult to deny that we understand the suggestion that we might, like Neo, or the Tom Cruise character in *Vanilla Sky*, be in a virtual reality setup. We can, after all, imagine the experiences Neo has in coming out. But to suppose that our capacity to imagine such experiences from the inside is what our understanding of the skeptic’s claims consists in is to confuse two senses of ‘skepticism’. In the first sense (“Humean” or “philosophical” skepticism), nothing could in principle count for or against any hypothesis about the external world. In particular, nothing could count as evidence that one was (or had been) in a virtual reality setup. On this view, all we could have in principle as evidence would be what we can be given directly in experience—sense-data—and no inference from sense-data to the external world

could ever be justified. This, however, is clearly not how we understand the examples. Neo's experience in which he seems to come out of the Matrix is taken as virtually conclusive evidence that he was once in it and now no longer is. And so is the Tom Cruise character's experience with an uncanny stranger who has knowledge of him that no one but he could have and apparently supernatural powers that no ordinary understanding of the world could explain.

There is, then, another understanding of the usual skeptical hypotheses ("real" skepticism), according to which we can quite easily imagine evidence that would count in their favor. And if we can imagine evidence that would count in their favor, then the lack of such evidence is clearly some evidence against them. More generally, we can say that this latter way of treating the skeptical hypotheses neither demands nor allows what the Humean understanding of hypotheses that sound exactly the same requires—that we entertain these hypotheses from a perspective in which we have put into doubt or "bracketed" all of our external-world beliefs simultaneously. The conclusion is that the apparent ease with which we understand the standard skeptical scenarios—i.e., understand *real* skepticism—provides no support for Stroud's claim: that there can be no transcendental argument against skepticism in its Humean or philosophical form.

5. Conclusion

That we can make sense of real skeptical scenarios, so understood, is difficult to deny. And yet it is no small concession. I shall conclude by spelling out some of the costs of the concession—costs that I am more than willing to bear, but that others may not be. First, though I have said a great deal about the phenomenology of our agential perspective, I have said nothing about what our agential capacities are. Both naturalistic accounts and dualistic accounts, if they are to figure in a philosophy of mind and language that does justice to Frege's constraint, are equally problematic.

Consider the question whether such capacities supervene on the physical. To suppose they do is to invite the question how large a supervenience base is needed. In particular it is to ask whether it must be any larger than the brain. But now consider two examples. In the first example, one is in perceptual contact with the external world in the usual way. One has, however, a duplicate in a virtual reality setup whose brain is molecule-for-molecule identical to one's own and is receiving the same electrical inputs. Suppose that in the duplicate's case the inputs are all completely artificial, and that the duplicate is out of touch with the external world. In the second example, the source of the electrical inputs to one's own brain switches back and forth (seamlessly) between the real world and an artificial source. Suppose it does so in such a way that one has no idea that such switches are taking place. And suppose that when one is told about the switches, one has no idea when or how often they occur. Considerations stemming from the role of Frege's constraint in demonstrative contexts (together with

the requirement that we say *how* it is satisfied in particular cases) suggest that we should opt for supervenience on what is intrinsic to the brain. The notion of seamless switching, after all, suggests that, temporal considerations aside, there is no difference from one's own point of view between the way things are given before and after the switch. Thus the cases of seamless switching suggest that events outside the brain make no difference to what things are like for one subjectively. The desire to do justice to the phenomenological fact that there is something it is like to have the experience one is having now—something shared by one's duplicate and the person-stages of oneself that are out of touch with the world—seems to point in the same direction. And yet if the arguments above are roughly correct, to opt for the supervenience of our relevant agential capacities on intrinsic features of the brain is to embrace one more version of the interface between the internal and the external that generates Humean skepticism and its meaning-skeptical consequence. Thus opting for such a supervenience thesis can only be self-defeating. However, the nonsupervenience of our conscious experiences on the intrinsic features of events in the brain, together with their supervenience on the physical threatens to make facts about what it is like to have one's present experiences mysterious. And the nonsupervenience of such facts on the totality of physical objects, facts, and events makes it mysterious how we could ever get in touch with mental events. Thus it leads to apparently intractable problems with self-knowledge and knowledge of other minds—problems that generate their own versions of meaning skepticism.

For completely independent reasons, I believe this problem of supervenience cannot arise. Very briefly, the agential perspective and the objective-causal perspective are, on my view, incommensurable in the Kuhnian/Hansonian sense.²⁰ If the foregoing phenomenological points are correct, each perspective generates its own perceptual base, in a way that prevents their coming into direct or immediate competition or conflict. This is not to say, of course, that there are no serious prospects of conflict at all. I believe, rather, that the relation between the two perspectives is dialectical—each grows naturally out of the other, and each, if carried to its logical extreme, undermines (or threatens to undermine) the other. As to the first point, the objective scheme grows naturally out of the agential scheme in virtue of our attempts to improve our agential capacities by viewing ourselves as objective entities and objects of appropriate sciences. And the agential scheme grows out of our need to do justice to the meaningfulness of our scientific terms via the account of “meaning as use” already sketched. As to the second point, there is the familiar argument that no causal-objective scheme, whether deterministic or indeterministic, seems to leave any room for genuine agency. And a resolutely agential perspective, particularly as regards the practice of science, seems impossible to reconcile with the assumption that scientific objectivity delivers the ultimate perspective on the world. Suppose, as is generally believed, that doing justice to agency involves taking one or more of the essential indexicals—‘I’, ‘here’, or ‘now’

²⁰ Kuhn 1962; Hanson 1961.

as ineliminable. Then doing justice to agency will compromise either the third person objectivity of the objective scientific scheme or the assumption that such a scheme provides the ultimate perspective on the world. These considerations support the idea of the incommensurability of the agential and objective schemes. Indeed, they support the idea of incommensurability in a much stronger sense than Kuhn or Hanson had in mind. And if these claims about incommensurability are correct, the question of supervenience cannot arise, since there is no common ground from which to ask it. The position, then, is not supervenience or nonsupervenience, but asupervenience.

That such a line rules out the usual causal theories of mind-related domains—for example, causal theories of perception, reference, action, and knowledge—is perhaps a conclusion that Putnam would not find uncongenial. In the case of action, for example, this suggestion has a very Wittgensteinian ring. (“Of course reasons can’t be causes—to suppose they could would be to ignore their normative dimension.”) Other consequences, however, are less familiar, and I mention only one. Imagine Neo, having emerged from the Matrix, recalls apparently thinking to himself the thought he would have expressed by saying “That woman is wearing red.” Since no such woman exists, it seems that neither we nor he could even entertain a thought of the type that seems necessary to do justice to his experience in the Matrix. My own view is that this is indeed a consequence of our taking *de re* intentionality as basic, but I shall end on a positive note. Our access to other subjects, as our access to the world, is less epistemological than practical—less a matter of knowledge than of know how. In the case of others, I believe that this know how involves not only knowing how to deal with them, but knowing how to play their parts, as Olivier knew how to play the part of Hamlet and as we know how to play ourselves. But this is, I think, clearly a topic for another occasion.

REFERENCES

- Alcoff L.M. (1998), *Epistemology: The Big Questions*, Oxford: Blackwell
- Evans G. (1982), *The Varieties of Reference*, New York: Oxford University Press
- De Caro M. & D. Macarthur eds. (2004), *Naturalism in Question*, Cambridge: Harvard University Press
- Gibson J.J. (1986), *The Ecological Approach to Visual Perception*, Hillsdale, NJ: Lawrence Erlbaum
- Gunderson K. ed. (1975), *Language, Mind and Knowledge*, Minnesota Studies in the Philosophy of Science, VII, Minneapolis: University of Minnesota Press
- Hanson N. R. (1961), *Patterns of Discovery*, Cambridge: Cambridge University Press
- Hume D. (1896), *A Treatise of Human Nature*, Oxford: Clarendon Press
- Hume D. (1988), *An Enquiry concerning Human Understanding*, Chicago: Open Court
- Kripke S. (1982), *Wittgenstein on Rules and Private Language*, Cambridge: Harvard University Press
- Kuhn T. (1962), *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press
- Lewis C.I. (1946), *Analysis of Knowledge and Valuation*, La Salle, Ill: Open Court
- Marsh R.C. ed. (1956), *Logic and Knowledge*, London: George Allen and Unwin

- McDowell J. (1994), *Mind and World*, Cambridge: Harvard University Press
- Parfit D. (1984), *Reasons and Persons*, Oxford: Clarendon Press
- Putnam H. (1975), "The Meaning of 'Meaning'" in Gunderson 1975, pp.131-193 reprinted in *Mind, Language and Reality*, Philosophical Papers, Volume 2, Cambridge: Cambridge University Press, 1975
- Putnam H. (1999), *The Threefold Cord: Mind, Body, and World*, New York: Columbia University Press
- Russell B. (1956), *The Philosophy of Logical Atomism*, in Marsh 1956, pp. 177-281
- Schiffer S. (1978) "The Basis of Reference", *Erkenntnis* 13, pp. 171-206
- Sidgwick H. (1981), *The Methods of Ethics*, Indianapolis: Hackett
- Stroud B. (1984), "Skepticism and the Possibility of Knowledge," *The Journal of Philosophy*, LXXXI, pp. 545-551, reprinted in Alcoff 1998, pp. 360-365
- White S. (1982), "Partial Character and the Language of Thought", *Pacific Philosophical Quarterly*, 63, pp. 347-365, reprinted in White 1991, pp. 27-49
- White S. (1991), *The Unity of the Self*, Cambridge: Bradford/MIT Press
- White S. (2004a), "Skepticism, Deflation, and the Rediscovery of the Self", *The Monist*, 87, pp. 275-298
- White S. (2004b), "Subjectivity and the Agential Perspective", in De Caro & Macarthur 2004, pp. 201-227, available at <http://ase.tufts.edu/philosophy/people/white.shtml>
- White S. (2007), "The Transcendental Significance of Phenomenology," *Psyche*, 13, pp. 1-31
<http://psyche.cs.monash.edu.au/symposia/siegel/White.pdf>
- Williams M. (1996), *Unnatural Doubts: Epistemological Realism and the Basis of Skepticism*, Princeton: Princeton University Press

Received: January 9, 2008
Accepted: August 28, 2008

Department of Philosophy
Tufts University
Medford, MA 02155 USA
stephen.white@tufts.edu