

Non-Conceptual Content and the ‘Space of Reasons’*

Richard G. Heck, Jr.

1 Opening

In *The Varieties of Reference*, Gareth Evans (1982) argues that the content of perceptual experience is non-conceptual, in a sense I shall explain momentarily.¹ More recently, in his book *Mind and World*, John McDowell (1996) has argued that the reasons Evans gives for this claim are not compelling and, moreover, that Evans’s view is a version of “the Myth of the Given”: More precisely, Evans’s view is alleged to suffer from the same sorts of problems that plague sense-datum theories of perception. In particular, McDowell argues that perceptual experience must be within “the space of reasons”, that perception must be able to give us reasons for, that is, to justify, our beliefs about the world: And, according to him, no state which does not have conceptual content can be a reason for a belief. Now, there are many ways in which Evans’s basic idea, that perceptual content is non-conceptual, might be developed; some of these, I shall argue, would be vulnerable to the objections McDowell brings against him. But I shall also argue that there is a way of developing it which is not vulnerable to these objections.

The view I shall defend here is not one I am entirely comfortable attributing to Evans—nor one I am particularly comfortable claiming as my own. Because Evans does not say very much about the nature of non-conceptual content, nor about the relation between perceptual states and beliefs, the text is simply too thin to support the attribution to him of any specific, developed version of the view that perceptual content is non-conceptual. There is textual evidence that elements of the view I am about to develop were accepted by Evans, and I shall discuss these

*First published in *Philosophical Review* 109 (2000), pp. 483–523.

¹ See, in particular, §§5.1, 5.2, and 7.4.

passages in due course. But there are other passages which suggest that parts of my position were, if not explicitly denied by Evans, at least no part of the view he succeeded in enunciating (or thinking through) in the time he had for the task. Nonetheless, I think that the position to which I shall commit myself is close enough in spirit to Evans's, and sufficiently consistent with it, that—as much in acknowledgment of his pervasive influence on my work, as for any other reason—I shall often speak of “Evans's view”, without pausing to defend the interpretive claim implicit in my so speaking.

Before I discuss McDowell's arguments, however, I should first say a few words about how I understand the claim that perceptual content is non-conceptual. The terminology suggests a view according to which perceptual states have a particular kind of content, non-conceptual content, as opposed to the conceptual content that beliefs are supposed to have. But the claim has not always been understood in this way.

To see why, consider one sort of reason Evans gives in favor of it, namely, that there is a difference between perceptual states, on the one hand, and cognitive states, like belief, on the other: What beliefs a given thinker can have depends upon what concepts she possesses; she cannot have any belief whose content essentially involves a concept she does not possess. But, says Evans, there is no such constraint on the contents someone's perceptual states might have: In particular, a perceiver can be in perceptual states a faithful report of whose content would necessarily employ concepts she did not possess.

Assuming that there is such a difference, how shall we account for it? What we seem to have here is, in the first instance, a distinction between two sorts of states—we might call them ‘concept-dependent’ and ‘concept-independent’. And one might suggest that what is revealed by the difference to which Evans draws our attention is just that there are these two sorts of states.² In particular, one might think that there is no

² I do not think the view I am about to explain is always distinguished from the one I intend to defend—so I do not think one will often find authors explicitly adopting one rather than the other. On the other hand, Tim Crane (1992, p. 155) says explicitly that a belief and a perception can have the same “whole” content; concepts for him are, in effect, abstracted from whole contents, in a way I shall discuss in note 7.

Crispin Wright (1996, pp. 244–5) suggests, as a possible response to McDowell, that one could hold that the content of experience can be conceptual—in the sense that it will, as it were, present the world to any agent equipped with the relevant concepts as instantiating those concepts—without agreeing that the content of experience is limited by the concepts available to the subject. This too might be taken to be a version of the state view: But Wright does not develop the idea enough for me to be sure what he

reason we must distinguish the kinds of contents beliefs and perceptions have: Whatever one might take the contents of beliefs to be—Fregean Thoughts, say—there is no reason that perceptions cannot have the same sorts of things as their contents; it is just that the contents of a thinker's perceptual states can, while the contents of her beliefs cannot, involve concepts she does not possess. But if this were one's view, it would be misleading to summarize it by saying that perceptual content is non-conceptual.³ There is, on this view, which we might call 'the state view', nothing unusual, as it were, about perceptual content. Perception is just a state of a different sort from belief: A non-conceptual, or concept-independent, state, as opposed to a conceptual, or concept-dependent, state. Since Evans (1982, p. 227) does speak, quite explicitly, of perceptual states as having non-conceptual content, I think we cannot interpret him as having intended to defend the state view. His view was what we might call 'the content view', that the content of perceptual states is different in kind from that of cognitive states like belief: The former is non-conceptual; the latter, conceptual.⁴

It is, in any event, the content view that I intend to defend here.⁵

To understand the content view, we need to understand what is meant by saying that beliefs have conceptual content, that is, what is being denied when one says that perceptual states do not have such content. Now, the less substantial the claim that beliefs have conceptual content—the less robust the notion of conceptual content—the less contentious the claim that perceptual states have conceptual content. Thus, for example, one gets the impression, from some authors, that they regard the claim that beliefs have conceptual content as being trivial, as if 'conceptual content' were just short for 'representational content': And if one uses the term 'conceptual content' in that way, 'perceptual states have conceptual content' will express a claim with which many—including Evans—would agree. So Evans cannot be using 'conceptual content' in that way. Indeed, as I understand him, when Evans claims that Thoughts—the contents of beliefs—are conceptual, he does not regard himself as emphasizing a

intends, and it is unclear whether he means to endorse it.

³ Unless, of course, one held that all content is non-conceptual—though, then, the claim so stated would not register any difference between belief and perception.

⁴ Jim Pryor suggested to me that it might be better to speak of 'conceptually-constituted content' and 'non-conceptually-constituted content'. This seems reasonable. But, as the terminology used in the paper is well-established, I am afraid we are stuck with it.

⁵ I suspect that the state view is indefensible—even incoherent, if coupled with the claim that the contents of beliefs are conceptual—but I shall not argue this point here.

triviality; he means to be making a substantive claim about the nature of belief, one which can be, and has been, denied.⁶

Unfortunately, however, it is not obvious what substantive claim Evans means to be making, and he does not say a great deal about the matter. But it is clear that the Generality Constraint is central to his thought on this subject, and Evans's usual way of stating the Generality Constraint provides us with a key to what he might mean by the claim that beliefs have conceptual content. According to the Generality Constraint, no thinker is capable of entertaining a Thought with a particular structure unless she is able to recombine the elements of that structure so as to form other, related Thoughts. Note the mention here of a Thought's structure: I take it that it is this idea, that Thoughts are, in a certain sense, structured, that Evans intends to express by saying that Thoughts are conceptual. If we call the elements of the structure 'concepts', then we can put this point by saying that Thoughts are composed of concepts or, better, that Thoughts are conceptually articulated.

Evans sometimes speaks of entertaining a Thought with a given structure as essentially involving the exercise of distinct, but interdependent, cognitive capacities: To entertain the Thought that John is tall, for example, a thinker would have to exercise a capacity to think of John and a capacity to think of a thing as being tall. And to have these cognitive capacities is, on Evans's view, to grasp the concepts of which the Thought is composed. So, in effect, what Evans is claiming is that no one is capable of entertaining a given Thought, conceived as structured in a particular way, who does not apprehend it as being so structured: To entertain the Thought that John is tall, one must exercise certain capacities, possession of which constitutes grasp of the concepts of which that Thought is composed; so, in entertaining this Thought, one must make use of one's grasp of these concepts, must conceive the Thought as composed of those concepts in a particular way. And it is, on Evans's view, because entertaining the Thought that John is tall presupposes

⁶ For example, by Robert Stalnaker (1998), who argues, pace McDowell, that the contents of beliefs and perceptions are both non-conceptual. Stalnaker understands very well that, once the nature of conceptual content has been properly understood, the claim that beliefs have conceptual content can be seen to be controversial; the claim that perceptions do, even more so.

Of course, it is an open question whether McDowell understands this thesis as Evans does: But McDowell certainly does not use the term 'conceptual content' in such a way that the claim that beliefs have conceptual content is trivial. Failure to grasp this point will make certain aspects of *Mind and World* hard to understand.

possession of such cognitive capacities—presupposes a grasp of certain concepts—that an ability to entertain that Thought requires the ability, mentioned in the Generality Constraint, to “recombine” its elements: If one really is able to entertain the Thought that John is tall, one must have a capacity to think of (or must have a concept of) John; if one also has a capacity to think of (or has a concept of) a thing as being bald, there can be no cognitive obstacle to one’s thinking the Thought that John is bald, too.

Suppose, then, that we understand the claim that beliefs have conceptual content in this way. The thesis that perceptual states do not have conceptual content then amounts to this: That perceptual content is not conceptually articulated, in the relevant sense; that being in a perceptual state with a given content does not presuppose possession of particular cognitive capacities; that being in a perceptual state with a given content need not require one to grasp certain concepts (as determined by its content, namely, those of which its content is composed). So the thesis that perceptual content is non-conceptual, as I understand it, implies the claim, highlighted above, that one can be in perceptual states an adequate specification of whose content would necessarily employ concepts one does not possess. For, if the content of perceptual experience were conceptually articulated, a thinker’s being in a perceptual state with a given content, composed of various concepts, would involve the joint exercise of cognitive capacities whose possession constitutes grasp of those concepts. But if, as Evans claims, a thinker need not even have the concepts that would figure in a “conceptual” report of the content of one of her perceptual states, she need not have the relevant cognitive capacities, either—in which case she will hardly be able to exercise them. The content of perceptual states would therefore have to be of some other, non-conceptual kind; their content would have to lack the conceptual articulation characteristic of Thoughts (as, say, a set of possible worlds does).

Obviously, there is much more to be said about what the claim that beliefs have conceptual content means, and I have done nothing to defend it. Nor have I said much about how the notion of non-conceptual content might be understood. But fortunately, we need not pursue these matters further here.⁷ The question whether beliefs have conceptual content is

⁷ I think Evans actually meant more by the claim that beliefs have conceptual content. In particular, someone who took the contents of beliefs to be sets of possible worlds could agree with what has been said so far—could profess to accept the Generality Constraint—so long as she conceived of grasp of a concept as an abstraction from an

not at issue between McDowell and Evans. And, as will become clear, McDowell's objections to Evans's view are, in effect, objections to any view on which perceptions and beliefs have different kinds of content, not just to a particular view which explains that difference in a particular way. (McDowell's objections must be of this level of generality: Evans provides no detailed explanation of the difference between conceptual and non-conceptual content.) Hence, if I am to answer McDowell's objections at the level at which they are put—as is my ambition—I cannot rely upon any particular way of understanding the difference in question. That is, my defense of Evans will not consist in developing a theory of the nature of conceptual and non-conceptual content: It will consist in getting clear about the epistemological relation between perception and belief; doing so will not require a commitment to any particular way of understanding their different sorts of contents.

But before we consider McDowell's charge that Evans's view is a version of the "Myth of the Given", it is worth discussing his objection to Evans's argument for the claim that perceptual states do not have conceptual content: that is, his objections to the Richness Argument.

2 The Richness Argument

Consider your current perceptual state—and now imagine what a complete description of the way the world appears to you at this moment might be like. Surely a thousand words would hardly begin to do the job. And it is not just that the description would be long: Rather, it seems hard to imagine that your perceptual state, as it is now, has any specific articulation corresponding to the conceptual articulation of a particular one of the many different Thoughts which might capture its content; and it seems at least as hard to imagine that you now possess all the concepts that would be expressed by the words occurring in such a description, even if one could be framed. Before me now, for example, are arranged

ability to entertain certain sorts of Thoughts (and of concepts themselves as abstractions from a class of Thoughts). For Evans, though, possession of the cognitive capacities which constitutes grasp of concepts is supposed in some sense to be prior to, or independent of, one's capacity to entertain such Thoughts as these: Possession of such capacities—grasp of certain concepts—is supposed to figure in an explanation of one's ability to entertain the Thoughts one can. One might say, then, that, for Evans, Thoughts are conceptual in so far as they are intrinsically conceptually articulated.

This sort of issue was discussed, some time ago, in connection with the question what justifies the attribution of knowledge of a structured, axiomatic theory of meaning to a competent speaker (Evans, 1985; Wright, 1993; Davies, 1987).

various objects with various shapes and colors, of which, it might seem, I have no concept. My desk exhibits a whole host of shades of brown, for which I have no names. The speakers to the sides of my computer are not quite flat, but have curved faces; I could not begin to describe their shape in anything like adequate terms. The leaves on the trees outside my window are fluttering back and forth, randomly, as it seems to me, as the wind passes over them.—Yet my experience of these things represents them far more precisely than that, far more distinctively, it would seem, than any characterization I could hope to formulate, for myself or for others, in terms of the concepts I presently possess. The problem is not lack of time, but lack of descriptive resources, that is, lack of the appropriate concepts.

McDowell insists that I do have concepts, available to me now, for describing the scene before me: These are demonstrative concepts, of shapes and colors—and, presumably, patterns of fluttering. This is a clever and appealing response to Evans's somewhat rhetorical question whether we even understand the idea that I have concepts of all the shapes and colors presently exhibited in my experience (Evans, 1982, p. 229). But let us be clear about why McDowell's response is so appealing: It is, in large part, because of the allusion to a familiar way in which experience provides us with concepts we otherwise would not have; it is because it embodies a suggestion about how these concepts are made available to us, namely, in the familiar way that demonstrative thought makes concepts available in other cases. To see why this is important, imagine that McDowell had omitted all mention of demonstration and had instead responded as follows:

It is true enough that the content of a perceptual state cannot always be given in terms of concepts antecedently available to its subject. But that doesn't matter: When we enjoy a perceptual experience of a particular scene, we have concepts we otherwise would not have. You now have, for example, concepts of all the colors and shapes arrayed before you.

This remark, though it may rightly be called a 'reaction' to Evans's question, hardly merits the name 'response', because it leaves it obscure how perception makes these concepts available to me, if I did not possess them before I had the experience. It does not answer Evans's implicit claim that we simply do not understand how I could have all these concepts: It just flatly asserts that we do have them. What McDowell actually says in reply to Evans does not have this shortcoming, and

the reason is that he takes demonstration as his model: The colors and shapes are there, and I see them, and so I can talk, or think, about that color or that shape (and acquire demonstrative concepts of them), in much the same way I can talk, or think, about that person or that tree (and acquire demonstrative concepts of them).

McDowell's response thus amounts to a charge that Evans has wrongly restricted the concepts in terms of which I might characterize the world, as I now experience it, to ones I have independently of, and prior to, my having this very experience. As McDowell (1996, p. 58) notes, the option he brings to our attention is not so much as considered by Evans. But it seems unlikely, to me at least, that as concerned as Evans is with demonstratives, he could just have overlooked it; I therefore suspect that there is a good reason Evans tacitly imposed the restriction McDowell questions.

First of all, note that, even if we do have (or can have) demonstrative concepts of the colors and shapes presented to us in experience, that does not imply that the content of experience is conceptual. Now, McDowell does not, so far as I can tell, claim that it does: But it might nonetheless seem surprising that Evans can accept McDowell's observation. Indeed, for reasons I shall mention below, even if one does not hold that perceptual content is conceptual, one should still regard McDowell as having made an important contribution here, because he has made it clear that there is a way we can conceptualize the content of a particular experience, without acquiring a concept that is independent of that experience (as are most of the concepts that we have prior to our having the experience, such as *blue* and *square*), namely, by forming certain sorts of demonstrative concepts.

The claim that we have such demonstrative concepts will threaten the view that the content of perceptual experience is non-conceptual only if it is impossible to explain how I might come to have such demonstrative concepts, whenever I have certain sorts of experiences, without supposing these concepts to be part of the content of the experiences which make them available to me. Even if I do now have (or could now form) a demonstrative concept of the color of a particular part of my desk, it does not follow that the content of my experience essentially involves that very concept: There may be a story to be told about how my having an experience, with a content in which a particular concept does not figure, makes it possible for me to form that concept.⁸ It is another matter

⁸ Michael Martin (1992) makes some remarks in a similar spirit.

to tell such a story: But Evans's ideas about how experience makes demonstrative concepts of objects available to us, which I shall mention shortly, could perhaps be deployed for this purpose, too.

These remarks suggest an objection to McDowell. Demonstrative concepts of the sort he discusses are ones I have only because I am presently enjoying (or have recently enjoyed) an experience of a certain kind. Or again, what explains my having these concepts is my having (had) an experience with a certain sort of content.⁹ But, if that is right, it is hard to see how these demonstrative concepts could be part of the content of my experience. Suppose we say, with McDowell, that my having certain demonstrative concepts is partially constitutive of the world's appearing to me in a particular way.¹⁰ How then can my having that concept be explained by my having such an experience? There would not seem to be sufficient distance between my having the experience and my possessing the concept for the former to explain the latter. So, if such an explanation were wanted, the content of the experience would have to be stated in terms which did not involve these demonstrative concepts or any other concepts of the colors and shapes in question, since the examples the Richness Argument exploits presuppose that I have no such other concepts: The content of perceptual experience would have to be treated as non-conceptual (in the relevant respects).

Although I personally find the argument just sketched compelling—extremely compelling—it cannot be conclusive, for McDowell could respond by just denying the causal (or explanatory) intuitions upon which it rests. If the demonstrative concepts in which we are interested here really are part of the content of my present experience, should we really be surprised if we find ourselves tempted to say we have them because we have such experiences? There is a perfectly good, non-causal sense in which my having such an experience makes the concepts available to me: I would not have them but for having it. And if one wants to ask, with Evans, whether we really understand how this sort of thing could possibly be true, the response will be that the concepts in question are

⁹ McDowell is concerned to emphasize that my demonstrative concept need not die with my experience: This is the reason for the parentheticals. But, having inserted them here, I shall omit them below, as it will be enough for us to focus on the situation which gives rise to the concept in the first place.

¹⁰ If to say that it appears as if that part of my desk is that color is to give a completely faithful report of (part of) the content of my experience, and if perceptual content is conceptual, I could not have such an experience unless I had the demonstrative concept that color.

demonstrative concepts, which experience makes available to us in much the same way it makes, say, a demonstrative concept of the glass to my right available to me.

The real trouble with this position is that what it takes for granted—that experience makes demonstrative concepts of objects available to us—is itself in need of explanation. In Evans’s own writings, the notion of non-conceptual content is introduced, not to resolve problems in the theory of perception (which is hardly Evans’s central focus), but because it is needed in his theory of demonstrative reference: Without the claim that perceptual content is non-conceptual, Evans could not give the sort of account he does of what fixes the contents of demonstratives such as ‘that object’. Very roughly, Evans’s story goes like this: For a demonstrative concept to be of a particular object x is for one’s attitudes toward contents containing that concept to be sensitive, in the right sort of way, to information about x , information which is, in central cases, delivered by perception (1982, pp. 145ff). It should be obvious that this explanation would be viciously circular if the information to which one was supposed to be sensitive (the content of the relevant perceptual experiences) had the conceptual articulation of ‘That object is F ’: One cannot have information with such a conceptual articulation without already having the demonstrative concept that object.¹¹ But the circle is broken if the information is (typically, and in relevant respects) non-conceptual.

Evans would thus have had good reason to reject McDowell’s suggestion that demonstrative concepts are part of the content of our experience: Such a claim would make it impossible to say anything substantial about what fixes the contents of such concepts; it is because it would have this consequence, I think, that Evans does not so much as consider McDowell’s proposal. McDowell, of course, would be unlikely to be moved by such considerations, since he regards as misguided any attempt to explain what fixes perceptual content (or any other sort of content). But we are not, despite appearances, at stalemate.

Suppose that I say, pointing twice at the same part of my desk (which part has a quite determinate color), ‘That part of my desk is that color’. True or false? True, of course—and this is something you know quite independently of any knowledge you might have about what color that

¹¹ One could have such information without having such a concept, if the state view were correct. (This is one sort of reason it might seem appealing.) But I mean to be setting that view aside here.

part of my desk is.¹² Though my utterance expresses a contingent truth, given that the references of the two occurrences of ‘that’ are linked in the way I said they are, the sentence simply has to be true, whatever it might express. The reason is that the demonstrative ‘that color’ refers to the color that part of my desk actually has (the color of the part of my desk to which I was referring when I uttered the words ‘that part of my desk’): The reference of this demonstrative, like that of demonstratives generally, is fixed by the world.¹³ As McDowell puts it:

In the throes of an experience of the kind that putatively transcends one’s conceptual powers—an experience that *ex hypothesi* affords a suitable sample—one can give linguistic expression to a concept that is exactly as fine-grained as the experience, by uttering a phrase like ‘that shade’, in which the demonstrative *exploits the presence of the sample*. (McDowell, 1996, pp. 56–7; my emphasis)

Presumably, the demonstrative concept “exploits the presence of the sample” in the sense that the particular shade it is a concept of is determined by the sample of the shade (the sample being the color of my desk, a color out in the world, which I am currently perceiving).

But this implies that the demonstrative expression ‘that color’ need not have the same content as my present experience of the relevant part of my desk, for you cannot know, just like that, as it were, whether I am perceiving it veridically, that is, whether I am experiencing it as having the color it in fact has or some other color. Perhaps the desk appears to me to be a redder shade of brown than it in fact is. Nonetheless, when I say that its color is that color, what I say is true. So, in the imagined

¹² Independently, maybe, even of your understanding what I have said, since, on Evans’s view, you do not understand what I have said if you are not, yourself, in a position to perceive the relevant part of my desk (or have some other informational link to it), though you do know what kind of thing I have said.

¹³ There are cases in which the demonstrative need not refer to the color of the desk. If there is a transparent film of color covering the desk, the demonstrative ‘that color’ could refer to the color of the film, in which case my utterance of ‘That part of the desk is that color’ need not be true. But I am interested here in what we might call ‘pure’ cases of misperception, cases in which I simply misperceive the color of the desk, not because of some external interference, such as films of color, funny lighting, etc. I find, for example, that it is frequently the case that, if I shut one eye, colors take on a slightly greenish tint, whereas, if I shut the other, they take on a comparatively reddish tint. That is the kind of thing I have in mind. (Such “pure” misperception can occur in film cases, too, of course: If I misperceive the color of the film, no “external” color need be available to determine the content of my experience.)

circumstances, my experience would not have the same content as ‘That part of my desk is that color’, since my experience would not represent the desk as being that color at all (since it would not represent it as being the color it in fact is). The content of my experience, that is to say, cannot, in general, be characterized in terms of demonstrative concepts: In particular, it cannot be so characterized when my perceptual state is non-veridical.

This argument may seem facetious. Perhaps it is true that the utterance ‘That part of my desk is that color’ cannot capture the content of my experience; perhaps it is true that the demonstrative expression ‘that color’ must refer to the color my desk actually has. Waive that: It doesn’t matter. Surely I am capable of entertaining a Thought regarding the color of my desk, one which need not be true: Surely I can form a concept of the color my desk appears to me to have. Is that not the demonstrative concept McDowell needs?—Yes and no. Yes, in the sense that I am inclined to believe that we can have demonstrative concepts of this kind.¹⁴ But no, because McDowell cannot respond to my argument in this way. McDowell claims not just that we can form concepts of the colors presented to us in experience, but that their references are fixed by a “sample” of the color in the world. To allow that my concept *that color* might not denote the color my desk actually has, but the one it appears to me to have, is obviously to deny that its reference is fixed by a “sample”. And that claim is not one McDowell can easily abandon, for what alternative conception of how the concept’s reference is fixed is available to him? The most obvious alternative would be to say that its reference is fixed by the content of my perceptual experience. But to say that would be to appeal to a level of perceptual representation McDowell does not want: If the content of my perceptual experience is to fix the content of my demonstrative concept of the color experience presents to me, my concept of that color cannot also be part of the content of that experience. If it were, the content of the demonstrative concept would be fixed by the content of that same concept.

McDowell might suggest that the argument I have given is but a veiled form of the argument from illusion. The crucial step involves

¹⁴ Just as Evans would insist that I cannot have a demonstrative concept of the little green man who appears to be sitting on top of my monitor, if there is no such man there, perhaps he would also insist I can have no demonstrative concept of the color the relevant part of my desk appears to me to be, if I am misperceiving its color. But the intuition is very strong that, if my experience presents my desk as being a certain shade of brown, I can form a concept of that particular shade. It will not do simply to deny it.

consideration of a case of misperception, and McDowell is suspicious of all such manoeuvres. Sometimes, in fact, he goes so far as to suggest that he is licensed to ignore the fact that perception is on occasion non-veridical (McDowell, 1996, pp. 9, 111–3). But misperception is a relatively common phenomenon: One does not need to wheel out sophisticated perceptual illusions to get cases in which we misperceive our environment; one need only pay careful attention to what, at this very moment, one's senses are actually telling one about the world.¹⁵ We should not build our theory of perception in such a way as to make misperception mysterious; we must not be so dismissive of misperception as inadvertently to preclude its very possibility. But that is what will happen if we refuse to consider cases of misperception, if we refuse to allow our theory of perceptual content to be shaped by them: And if we insist that the content of my experience of the color of the relevant part of my desk must be given in part by a demonstrative concept whose content is fixed by the world, it will indeed be impossible that I should (simply) misperceive its color.

McDowell's writings do contain some remarks about how he might handle cases of misperception, and one might suggest that these ideas should be invoked here. McDowell proposes a disjunctive account of seeming to perceive that p : One will seem to perceive that p just in case, either one really does perceive that p (veridically) or one merely seems to perceive that p (non-veridically). Crucially, these two states are supposed to have no "common factor": There is to be no single sort of perceptual state which can be both veridical and non-veridical.—In this same spirit, then, McDowell might deny that we need to give a single, unified account of what determines perceptual content, one which works both in veridical and non-veridical cases: The fact that demonstrative concepts of the sort we have been discussing can play no role in a general account of perceptual content—since they can play no role in cases of non-veridical perception—would then be no objection to our making use of them in giving an account of perceptual content in veridical cases.¹⁶

¹⁵ Perhaps the beliefs I form on the basis of my experience are, by and large, true. But I see no reason to think my total perceptual state is usually veridical. My total perceptual state almost always represents the world as containing something that is ringing at a particular frequency, which there rarely is, and my visual experience misrepresents depth with an astonishing regularity. And because I wear glasses, the periphery of my vision almost always represents the world in a most peculiar fashion. (It is also worth remembering that illusions, like the Müller-Lyer illusion, highlight quite ordinary features of perceptual experience. That is why they are interesting.)

¹⁶ Of course, the question what to do about the richness of non-veridical perception would remain open: But perhaps that would not trouble McDowell.

The disjunctive account of perceptual knowledge seems incredible to me. But it is not my purpose to examine the arguments McDowell uses to motivate it, nor to evaluate the position itself. For whatever one thinks of such accounts in general, a disjunctive account of what fixes representational content is plainly a non-starter. Compare the case of belief: Suppose we were offered a disjunctive account of the contents of beliefs, one which might say, for example, that some belief of mine, if true, had the content that snow is white; if false, that snow is pink.¹⁷ Such a belief has no definite content at all, and such a theory about how the contents of beliefs are fixed is viciously circular. We cannot say what content my “belief” has unless we know whether it is true or false—but we cannot answer the question whether it is true, because we do not know what its content is—because we do not know whether it is true—etc., etc. A similar point applies to disjunctive accounts of perceptual content: If we do not know what the content of any perceptual state is until we know whether it is veridical, then we have no grip on what it is for such a state to be veridical; we cannot first decide whether a perceptual state is veridical and then go on to say what its content is. And if a disjunctive account of the contents of perceptual states is hopeless, then a disjunctive account of what fixes their contents is equally hopeless, and for the very same reason.

Demonstrative concepts of colors, shapes, and the like, can therefore play no role in the theory of perceptual content. Of course, it does not follow from this that I do not have concepts of all the colors and shapes presently exhibited in my experience. I may have “experience-dependent” concepts of these things (that is, concepts which come into my possession whenever I have certain sorts of experiences), even if these concepts are not rightly called ‘demonstrative’. But even if that is true, it is of no help to McDowell: It is, as I said earlier, in no conflict with the claim that perceptual content is non-conceptual. For there to be such a conflict, these “experience-dependent” concepts would have to figure in the content of the experience itself, and we have been given no reason to believe they must do so.

Let me concede, then, that the richness argument does not prove that perceptual content is non-conceptual: But that should have been obvious all along; one can always just insist, flatly, that the concepts

¹⁷ Note that this is different from saying that my belief has the following content: That, if it is true, then snow is white, and, if it is false, then snow is pink. We certainly can, and I certainly do, have beliefs with that sort of content.

which figure in the content of a given experience, even if we do not have them prior to the experience, enter our consciousness along with it. So the Richness Argument does not refute the view that perceptual content is conceptual: Rather, it challenges anyone who would hold that it is to explain how experience provides us with a grasp of all the various concepts that would figure in a conceptual report of its content. Such an explanation must be offered if Evans's question—whether we so much as understand this sort of proposal—is so much as to be addressed, let alone answered. That is why the analogy with demonstration, which might otherwise have seemed a tantalizing add-on, is in fact essential to McDowell's position: Without it, his response to Evans is nothing but an emphatic re-iteration of the claim that perceptual content is conceptual.

Of course, McDowell thinks that perceptual content has to be conceptual, if we are to give a proper account of the rational relations between perception and belief; if so, there must presumably be some answer to the challenge posed by the Richness Argument, even if we do not now know what it is. So let us turn to the question whether Evans's view is compatible with the existence of the sorts of rational relations between perception and belief on which McDowell focuses our attention.

3 Sense-Data versus Non-Conceptual Content

As said earlier, McDowell argues that Evans's position, and any other committed to the claim that perceptual content is non-conceptual, is but a new version of the Myth of the Given. The content of this colorful claim is that any such view will be vulnerable to objections which are but re-fashionings of objections to sense-datum theories of perception. To evaluate this charge, it seems best to start by asking what McDowell takes the real problem with these much maligned theories to be.

The distinctive claim of the sense-datum theory, he says, is that perceptual experience is purely sensational, a merely subjective happening, with none but a causal connection to the outside world. One might put this by saying that my current perceptual experience, though it has "content", has only sensational content, not representational content: It is, as one might put it in an undergraduate course, not about anything. It is, for that reason, obscure how, on this view, my current perception is supposed to justify my believing, say, that there is a brown desk in front of me. If my current perceptual state has no representational content, then it can no more stand in a semantic relationship to a belief than

a rock can: It can neither imply it nor make it probable. But that my perceptual states surely must do if they are to justify my beliefs about the world, for, very roughly, part of the story about how a particular perception justifies a belief will be that the content of the perception underwrites that of the belief it justifies, in the sense that it logically implies it, or makes it probable, or maybe just in the sense that an inference from the perception to the belief would, in the present context, be reliable.¹⁸

This is one of the problems on which McDowell focuses: That we lack any plausible story about how, on the sense-datum view, perception might provide us with reasons for (that is, justifications of) the beliefs we take to be based on our experience. McDowell considers two ways a sense-datum theorist might respond to this problem. Classical sense-datum theorists adopted a quasi-observational, “Cartesian theater” sort of picture: On this view, our beliefs about the world are justified by beliefs we form *about* our sense-data. The phenomenal properties represented in our experience are conceived, not as features of perceptual states, but as the immediate objects of perception; since sense-data are immediately present in experience, we can form various sorts of beliefs about them, and then infer from these to beliefs about the world. But even if the quasi-observational model could be defended, the question would remain how experience provides us with justifications for our beliefs. Something with intentional content, namely, my beliefs about my sense-data, are in the picture now, so something is in play which is of the right kind to be a reason for a belief. But nonetheless, it seems quite impossible to answer the question how our beliefs about mere subjective occurrences, however well-founded, could provide us with any reason to believe one thing rather than another about the world. This is not just because it seems difficult to say (for all that phenomenologists have tried) how we might “translate” statements about sense-data into statements about the world. It is, more fundamentally, because it is part and parcel of classical versions of the sense-datum view that these subjective occurrences have, in and of themselves, nothing to do with the world.

¹⁸ Some theorists of perception have denied this claim, namely, those who think that perceptual knowledge is to be justified by an inference to the best explanation: The existence of an external world causing them is supposed to be the best explanation of our having the sense-data we do (Jackson, 1977). I take it that McDowell’s objection to such views would be one to be mentioned in note 20 below, namely, that they leave it obscure how we could so much as have concepts of external objects and properties. But I do not intend to evaluate this claim here.

Despairing of this strategy, a sense-datum theorist might instead abandon the idea that perceptual experience justifies our beliefs and insist, with Davidson (1986), that it simply causes them.¹⁹ But this view, or so McDowell argues, is untenable, since experience must provide us with justifications for our beliefs about the world and not just “exculpations”.

There are many questions that could be raised about this part of McDowell’s argument; I am not sure I fully understand it. It is natural to think that it depends upon substantive claims about the nature of justification: I would certainly understand the complaint that McDowell does not take externalist theories of perceptual justification seriously enough. But McDowell does not present the argument as depending upon any very specific claims about justification; I do not think he sees it as doing so, and I do not think it needs to do so. Rather, the argument is driven by a much more general claim, really, by an intuition, namely, that I do not just find myself having certain beliefs, such as that there is a brown desk in front of me, having no idea where they came from; it is not as if perceptual experience gives rise to perceptual beliefs in the same sort of way a bump on the head might cause me to believe that I am Napoleon. On the contrary, I see the desk, and I believe that it is there *for that reason*, not just because I see it there. To be only slightly more precise: The formation of perceptual beliefs is a rational (not just a causal) process.

As it happens, I find myself sharing McDowell’s core intuition. But, as we shall see, he thinks this already causes trouble for me: McDowell’s argument against Evans is precisely that Evans cannot respect this intuition, without tottering over to a view that takes the formation of perceptual beliefs to be mediated by judgements about one’s experience (as the classical sense-datum theory does).²⁰ What I want to do now is to explore the basis of McDowell’s suspicion.

¹⁹ McDowell presents Davidson’s view as an alternative to the sense-datum theory, which is certainly how Davidson himself conceives it. But it seems to me that, if the objection McDowell (1996, pp. 13–8) brings against Davidson works at all, it works against this view, too.

²⁰ One thing that is distinctive of McDowell’s treatment of these issues is his connecting them with questions about how, if the sense-datum theory were correct, we could so much as have beliefs, not just about our sense-data, but about the world itself—let alone how these beliefs might be justified: How, McDowell wants to ask, could we so much as have a concept of a square object if that concept never made contact with square objects? indeed, with anything outside our minds? But I shall not discuss this aspect of his presentation in any detail, since I do not think it bears directly upon the concerns of this paper.

McDowell's own way out of the bind into which the sense-datum theory forces us begins with the suggestion that we should reject the picture of experience as a mere subjective occurrence, with none but sensational content, and insist that experience—as it is in itself, so to speak, and not just as we make it in thought about it (or in reaction to it)—is structured by concepts. This appears to promise a neat resolution of the problems that plague the sense-datum view: If experience presents us with conceptual contents, it is relatively easy to see how our perceptions might stand in semantic relations to our beliefs; perhaps it will not be much harder to see how they can stand in rational relations to beliefs.

Even at first blush, however, this response seems like an over-reaction. It is, in particular, hard to see how Evans can be vulnerable to these sorts of objections when he writes:²¹

...[A]s regards perception, it is now widely realized that the traditional conception gets things impossibly the wrong way round. The only events that can conceivably be regarded as data for a conscious, reasoning subject are seemings—events, that is, already imbued with (apparent) objective significance, and with a necessary, though resistible, propensity to influence our actions. (Evans, 1982, p. 123)

Sense-datum views are vulnerable to the objections rehearsed above because they deny that experience has intentional, or representational, content: The objections exploit the fact that, on such views, experiences are mere subjective happenings with no intrinsic bearing on the world. *Prima facie*, then, it would seem that one could avoid these objections by insisting that perceptual experience does have representational content. But non-conceptual content *is* representational content, so it is far from clear why the additional claim that experience has conceptual content should be required.

Let me be clear about what is being proposed here: That particular perceptual experiences have representational content, though that content is not conceptual. This view is to be distinguished from the following one, sometimes called the 'representational theory of perception': According to it, the immediate objects of experience are not a mere wash of color (like an expressionist painting); they have intentional properties. That is,

²¹ McDowell himself describes Evans's view as being, for just this reason, "deceptively innocent" (McDowell, 1996, p. 53) and admits that his charge that it is a version of the Myth of the Given "can seem difficult to sustain" (McDowell, 1996, p. 162).

the sense-data we immediately perceive represent the world (so they are more like a photograph). To this view, one might rightly object that it no more helps us to understand how perceptual experiences might justify beliefs than does the sense-datum theory in its more pure form. Allow that sense-data have intentional content, that is, that the immediate objects of perception represent the world. How does the content of that representation find its way into my thought? On this view, sense-data do have intentional content: But that does not, of itself, imply that my perceptual experience has any intentional content at all; only the immediate object of the experience has such content. So it remains obscure how my perceptual state could possibly provide me with a reason for a belief. On the representational view, then, if we are to explain how perception gives rise to belief, either we shall have to say that perceptions merely cause beliefs—which McDowell will say amounts to abandoning any rational connection between them—or we shall have to posit intermediate beliefs we form about our sense-data. These beliefs will not, in this case, be about what phenomenal properties are present (though that might be needed as yet another step), but about what these represent: I will have to form such beliefs as that the sense-data I am presently experiencing represent a brown desk, etc., etc., and then use that as my reason to believe what I do about the world.

Perhaps the representational view is the one McDowell has in mind when he says that Evans's view is but a warmed-over version of the sense-datum view. If so, then I agree, at least to this extent: The representational view *is* but a warmed-over version of the sense-datum view; it retains the idea that the immediate objects of perception are (not objects in the world but) sense-data, adding only that these sense-data represent the world. But the representational theory of perception does *not* embody the idea that experience has non-conceptual content: The representational view is not that our perceptual states have representational content, but that our sense-data do. And conversely, the claim that perceptual states have representational content, non-conceptual or otherwise, does not imply that the immediate objects of experience are not objects in the world, but instead are representations. If particular perceptual experiences have representational content, then there must presumably be such things as representations which have the representational content, but it simply does not follow that these representations are the immediate objects of perception. They might be nothing but certain sorts of structures in the brain. In any event, no perceptual intermediary is needed: My perceptual state can have representational

content all by itself, as it were, whether or not that content is conceptual.

The view I am defending, and the view Evans meant to defend, is thus that perceptual states themselves have representational content, though that content is non-conceptual. Of course, there are a number of ways that this content might be conceived, and I do not want anything here to depend upon how it is, ultimately, to be conceived. But it may be helpful if we focus for a moment upon Evans's suggestion that the content of a perceptual state be understood as the information it carries, in the broad sense in which that term is now used. The notion of information is, in the first instance, a causal notion, not a cognitive one: What information a given state carries will, in general, be a function of its causal ancestry and, presumably, of more general facts about the organism in which it occurs, including facts about its design (or evolutionary history). What information a state carries is, roughly speaking, a matter of its place in the causal nexus, not of its place in some rational order: So it seems likely that, if perceptual content is to be explained in terms of information, in this sense, it is going to be non-conceptual.²²

The question McDowell wants to press is: How might my enjoying a particular experience justify a particular belief, say, that there is a desk in front of me? What we have seen, so far, is that Evans's view satisfies a necessary precondition for answering this question, not yet that it can answer it. Experiences, on his view, have representational content: They are at least the kinds of things that might justify beliefs. If, for example, the information carried by a given perceptual state is a scenario, a set of ways in which the space around the observer might be arranged, as on the view of Christopher Peacocke (1992) view, there will be no bar whatsoever to perceptions' standing in semantic relations with beliefs: Some beliefs about how space is arranged will be inconsistent with its being arranged in one of the ways the scenario includes; others, required by it; others, made probable by it; others, in the context, could be reliably inferred from it; and so on.

What, then, is McDowell's problem with this sort of view? If it is

²² I have heard it objected that an organism will only be able to make use of the information that experience provides in so far as it can conceptualize it, where the conclusion is supposed to be that the content of the experience—the content it has for the organism—is limited by the concepts the organism possesses. But the inference from premise to conclusion is questionable (it ignores the distinction between concepts one has independent of the experience and experience-dependent concepts which it might make available) and the premise is independently doubtful (see note 25 and the text to which it is attached).

possible for me to have a perceptual experience whose content entails that of a particular belief, what obstacle is there to seeing how that experience might justify the belief? The answer is that, according to McDowell, if perception is to justify belief, then my enjoying a particular perceptual experience must give me a *reason* to hold a certain belief; and, or so he claims, *only something with conceptual content can be (or provide me with) a reason*. This claim—which McDowell puts by saying that the “space of reasons” is contained within the “space of concepts”—is, as I understand his argument, its most central premise: But we have yet to uncover any argument for this premise, as opposed to the correct, though near trivial, claim that the “space of reasons” is contained within the space of (representational) content.

Now, McDowell’s suspicion is clearly that the only answers Evans can give to the question how perceptual experience justifies beliefs about the world—or how our experiences might give us reasons for our beliefs—will be versions of answers vainly tried by sense-datum theorists. Evans could try saying that what we form, in the first instance, are beliefs about the contents of our perceptual states; he could try saying that perceptions do not justify beliefs at all, but merely cause them. But McDowell has reason to be dissatisfied with those views. And McDowell’s own view, that the content of a perceptual experience can, so to speak, simply be taken up as the content of a belief, is unavailable, since, on Evans’s view, although perceptions and beliefs both have content, they do not have the same kind of content. So one can, at least to this extent, understand why McDowell thinks Evans will prove unable to cope with the rational relations between perception and belief.

Still, all we have so far is an argument that none of these three sorts of answers is both workable and available to Evans: We have not seen that there are no other options. To make further progress, we need to focus more clearly upon the question McDowell wants to press: Whether the mere fact that experiences, with their non-conceptual contents, can stand in semantic relations to beliefs is sufficient to justify our saying “that one term in such a relation can be someone’s reason for another” (McDowell, 1996, p. 53). One can begin to see what is bothering McDowell here by noting that, if we are simply caused to have certain beliefs, whenever we enjoy perceptual experiences with certain non-conceptual contents, then perception does not provide us with justifications for our beliefs: Not because it is without content, but because we have a justification for a belief only where we have a reason, and we have a reason for holding a belief only when our coming to hold that belief for that reason is a

rational process (not merely a causal one).

If that seems compelling—as it does to me—then the question whether Evans can make sense of the perceptual justification of belief reduces to these ones: What is required if the formation of beliefs on the basis of experience is to be a rational, and not merely a causal, process? Can we understand what it might mean to say that someone’s reason for holding a particular belief is that she is presently enjoying an experience with a particular non-conceptual content?

4 The Presentational Character of Perceptual Experience

McDowell’s view is that only if we say, with him, that experience has conceptual content, can we escape the epistemological problems which plague sense-datum theories of perception. But, or so I shall now argue, saying that perceptual experience has conceptual content will not, by itself, enable us—that is, it will not enable McDowell—to answer the objections he brings against the sense-datum view. This may well seem surprising. McDowell emphasizes, time and again, that what I come to believe on the basis of my perceptual experience is just what I see: If so, the content of the perception can simply be taken up as the content of a belief. This does make it clear how one might come to hold a particular belief, given that one is having a particular experience: So one of the problems Evans has to face—that of explaining how we “translate” between different types of representational content—will not arise for McDowell.

But that problem is, rightly, not the one McDowell regards as fatal to Evans’s position.²³ The fatal objection is supposed to be that Evans cannot explain how a particular perceptual experience might justify a particular belief, not that he cannot explain how perception might cause beliefs whose content is appropriately related to that of the perception. And, or so I shall now argue, the claim that perceptual content is conceptual is not sufficient to allow McDowell to resolve this problem: To make intelligible perception’s role as the rational basis of our beliefs about the empirical world, we must make a different sort of adjustment to the picture of perception with which the skeptic attempts to shackle us. This is an adjustment one can make so long as one regards perception as having representational content—and so, in particular, whatever one’s

²³ See note 29 and the text to which it is attached for some further remarks on this issue.

view about whether it has conceptual content.

McDowell likes to say that, in experience, we are saddled with content. But what is important is not just that we are saddled, but how. All sorts of mental states have content: Desires and intentions, as well as beliefs; even mere considerings, entertainings of Thoughts. But none of these can be a reason for a belief. So, if a sense-datum theorist were to be converted to the view that perceptual experience has conceptual content, but held that it was like the entertaining of a Thought, she would not be much closer to a resolution of the epistemological problems about perception than she was before. That is, if perception were like a little voice telling me, 'Consider the Thought (the possibility) that there is a desk in front of you', then it might be clear enough how I might get from this Thought to the belief that there is a desk in front of me: No translation would be required. But it would remain obscure why I *ought* so to believe.

I take it to be a well-established, and familiar, point that perceptions are not beliefs. This is unfortunate, for if they were, we would have a relatively easy answer to the question how experience justifies beliefs about the world. I do not just mean that perceptions would then be beliefs about the world: They would, of course, but the question how these "perceptual beliefs" might justify other sorts of beliefs (non-perceptual ones, for example) would also be less troubling. Though one can certainly raise questions about how some beliefs justify other beliefs (how the beliefs we now hold give us reasons to hold other ones), these sorts of questions seem relatively tractable—much more tractable, anyway, than questions about how perceptions justify beliefs. It is just in the nature of beliefs to stand in justificatory relations with other beliefs: If one belief is entailed by another, and if I recognize that it is and, for that reason, infer the one from the other, it is hard to see on what ground one might deny that the latter belief is my reason for the former, or that my new belief is justified by my old one. Not all cases are so simple: It is very hard to explain when one belief underwrites another (to use terminology from above), that is, what weaker relations than entailment support knowledge in such cases. But one does not find oneself wondering whether beliefs really do justify other beliefs, as one really may find oneself wondering whether perception justifies belief. Only the most extreme skeptic would be tempted by the view that beliefs only cause, and cannot justify, other beliefs.

Perception is not belief. But no-one, so far as I know, has ever been so much as tempted to say that perceptions are desires, intentions, or

entertainings; only beliefs are liable to be confused with perceptions. The reason is that, as different as perception may be from belief—as isolated in certain ways as perceptual experience is from the influence of our beliefs—there is yet something similar: Both purport to represent how the world is; both, we might say (borrowing some terminology from the philosophy of language) have assertoric force.²⁴ Even when the world appears to be a way I know it not to be—when a stick I know to be straight looks to be bent when I partially immerse it in water (to use a tired example)—it still *looks as if the stick is bent*. That is to say, my experience represents the world as containing a bent stick: In a different way, to be sure, than my beliefs would were I to believe that the stick was bent, but it represents it as being that way nonetheless. Indeed, though I know that the end of the stick is not where it looks to be, if I want to grasp something I have just touched with it, that does not seem to stop me from moving my hand towards the end of the stick, as its location is represented in my experience.²⁵

It is only if we recognize this feature of perception, what we might call its presentational²⁶ aspect, that we can begin to unlock the skeptic's shackles. But how should we incorporate an acknowledgment of it into a theory of perception? Familiar ways in which perception can be conceived as relational might well be thought to be attempts to come to terms with it: Perceptions as dispositions or inclinations to believe. But I suspect that no such account will do justice to the presentational aspect of perception, which seems to me, at least, to be independent of its tendency to give rise to belief (Evans, 1982, pp. 229–30). But even if I am wrong about that, the references to belief, while understandable, are quite unnecessary. Perceptions are not beliefs: But they may yet be attitudes of some other kind, even if they are not desires, intentions,

²⁴ Although the comparison with testimony can be misleading, if not treated with care—it tends to suggest the representational theory considered in the last section—we might say that perception is more like a little voice saying, 'There is a desk in front of you' (Evans, 1982, pp. 122–3).

²⁵ This suggests that perceptual experience can influence action otherwise than by giving rise to beliefs. (Evans might have had something like this in mind when he wrote of connections between the informational system and action which are "more primitive" than its connections to belief.)

²⁶ I use this term, rather than 'representational', because the term 'representational content' is in such common use. (The term has also been used by John Searle (1983).) Indeed, when people speak of perception as having representational content, this seems to me to run together the two aspects of perceptual experience I am trying to distinguish. One does not easily speak of intentions as having representational content.

or entertainings. In fact, I suggest, perceptions are attitudes, attitudes which are like beliefs in so far as to be in a perceptual state is to hold an assertive, or presentational, attitude towards a certain content.²⁷

If that is right, then we can begin to answer the question how perceptions justify beliefs, and how they provide us with reasons for belief, thus: Pretty much the same way beliefs do—whatever that may be. We must not allow the fact that perception is fallible to obscure this point. The fact that not all of my beliefs are true gives me no reason not to infer other beliefs from them. So when the skeptic asks, “How can you form beliefs on the basis of perceptions you know sometimes to be misleading?” I want not just, as McDowell (1996, p. 143) would have it, to shrug my shoulders, but to ask: Would you also suggest that we abandon the practice of inference just because some of our beliefs are false? How else would you have us find out which of them are false?

Any serious skeptic would presumably answer the first question affirmatively—and shrug *his* shoulders at the latter. But I am not suggesting that all the epistemological problems about the perceptual justification of belief will vanish if we treat perceptual states as being presentational. I am suggesting that, once the presentational aspect of perceptual experience has been properly recognized, the problems which remain concern how, or under what circumstances, a given perceptual experience justifies a belief; that the problems which remain are similar to problems about how, or under what circumstances, beliefs justify other beliefs; and so that no serious doubt need remain about whether our perceptual experience *can* justify our beliefs.

The claim that perceptual states are presentational is, it should be clear, neutral with respect to the questions at issue in the present paper. McDowell could accept it, of course, and say that perceptions are appearances—lookings, soundings, feelings, or what have you—emphasizing, as I have, that they have assertoric force, in the sense that beliefs do. But, by the same token, the move is available to Evans (who makes it tacitly, when he says that they are seemings) and to just about everyone else—except sense-datum theorists, and their close kin, the representational theorists, since it is a prerequisite that one accept that perceptual states

²⁷ Of course, it would be nice to know more about what it means to say that perceptual states are “assertive”. Unfortunately, I do not know how to explain this. Nor, however, do I know how to explain what is involved in a belief’s being assertive: What exactly does it mean to say that beliefs purport to represent how the world is? I think an answer to this question could be converted into an answer for the case of perception, too. At least, that is the point of my relying upon this analogy here.

have representational content. But, although the claim that perceptual states are presentational is, in that sense, neutral with respect to what is at issue between McDowell and Evans, it is *not* neutral with respect to McDowell's arguments against Evans: Once we have acknowledged the presentational character of perception, we are at last in a position to answer McDowell's charge that the only story Evans will be able to tell about how perception justifies belief will be a variant of one of the sense-datum theorists' own failed attempts.

I look about me, and as a result I am in a state which has a certain non-conceptual content. I now form the belief that there is a brown desk in front of me. How? Is it merely caused? Or do I have first to think, "I am in a perceptual state with such-and-such non-conceptual content", and then use that as my reason? Those ways lie old traps. Do I instead think something like, "That object is that color, and that color is brown, and that object is a desk"? Maybe that would do, but probably not: Can we not just ask what justifies my belief that that object is that color? That is just the same sort of question. Instead, let us say this: When I look around, the result is, not just that I come to be in a state with a certain content, but that I come to have a presentational attitude towards a particular non-conceptual content; then there is an inference²⁸ or transition—call it whatever you like—from my perceptual state to some belief I recognize it to underwrite.²⁹ My belief will then be justified

²⁸ The occurrence of this word should set off no alarms: When people deny that perceptual knowledge is inferentially justified, what they mean, typically, is that it is not inferred from any prior beliefs about my experiential states. The view I am developing can acknowledge this point: If it is right to say that perceptual beliefs are inferred, they are inferred from prior presentational attitudes about the world, that is, from perceptions.

Although I have no special interest in retaining the word "inferred", it seems to me to be useful. In making an inference, although there is a sense in which I must recognize that the premise entails the conclusion, if the inference is to yield knowledge, my recognizing this need not, and in general cannot (on pain of Lewis Carroll problems), involve my having the belief that it does so. So room has to be made, even here, for my recognizing the existence of such a semantic relation, without my forming a belief that it obtains.

²⁹ I would certainly understand a request to be told more about the nature of this transition—and I wish I had something positive to say about it, that is, something to say about conceptualization, which is what the translation between types of representational content involves. But I do not see any reason for skepticism that there is something to be said here: The problem is not one that should be underestimated, but Peacocke (1992, esp. §1.2 and pp. 79–80) has done much to explain in what a grasp of certain basic, observational concepts might consist and how possession of them is related to the capacity to form beliefs on the basis of perceptual experience. McDowell (1996, pp.

by the perception on which I base it, in much the same way it might have been justified by another belief upon which I based it. Moreover, my perception is my reason for my belief: At least, we have not yet seen any reason it should not be.

5 Reason and the Nature of Appearance

McDowell writes:

If [the semantic] relations [between perceptual states and beliefs] are to be genuinely recognizable as reason-constituting, we cannot confine spontaneity within a boundary across which the relations are supposed to hold. The relations themselves must be able to come under the self-scrutiny of active thinking. (McDowell, 1996, p. 53)

I find it hard to be sure just what is involved in spontaneity's imprisonment. However that may be, though, we can extract from this remark a line of thought that may appear to pose a serious threat to the view that perceptions, if their content is non-conceptual, can give us reasons for our beliefs. But let me emphasize that I am far from certain that I understand what McDowell has in mind here. For that reason, what follows is more speculation than interpretation.

McDowell's central idea (I think) is that, if perception is to give me reasons for my beliefs, the force, relevance, and weight of these reasons must themselves be rationally evaluable. Suppose that I am presently perceiving a stick, partially immersed in water, which appears to be bent; yet I do not form a belief that the stick is bent. Why not? Presumably, it is because I have learned from past experience that a perceptual state with this kind of content does not give me sufficiently good reason to hold such a belief. But now, McDowell might ask, what kind of content does a perceptual state with "this" kind of content have? Aren't perceptual states with "this" kind of content precisely ones in which it appears to me that a stick immersed in water is bent? and so states with conceptual content? In general, it would seem, if I am rationally to evaluate the reasons perception gives me to believe various things, I have at least to be able to think about those reasons. But I will not be able to do

162–3) discusses and criticizes certain of Peacocke's proposals, but his objections mostly concern Peacocke's right to speak of a subject's reasons. It is that right, of course, that I am defending.

so if the reasons perception provides have non-conceptual content, for non-conceptual contents cannot figure in my thoughts.

These reflections can appear very powerful, and I think there are lessons that emerge from them. But they do not, on their own, amount to an argument that perceptual content is conceptual. If we do try to construe them as constituting an argument for that conclusion, it seems to me that we can do no better than this:

1. If my perceptual states are to justify my beliefs about the world, they must provide me with reasons for those beliefs.
2. If perception gives me reasons for my beliefs, then I must be able rationally to evaluate the force of those reasons. (Their force “must be able to come under the self-scrutiny of active thinking”.)
3. In order rationally to evaluate the reasons perception gives me for holding a particular belief, I shall have to consider how the world appears to me, how it appeared to me on previous occasions, and so forth: I shall have, that is to say, to entertain such Thoughts as that it presently appears to me as if p , that it has previously appeared to me as if p , when in fact not- p , and so on.
4. Therefore, perceptual content must be conceptual.

The first two premises I am prepared to accept, and the third is established by the reflections (and, in any event, it seems plausible enough). But no more follows from these premises than that, if I am to have a perceptually justified belief that p , I must be able to entertain such thoughts as that it presently appears (or previously has appeared) to me as if p , and so forth. An additional premise is needed if the reflections are to support an argument for the conclusion that perceptual content is conceptual.

But what, one might ask, is the difference between my reflecting on the reasons perception gives me for forming a particular belief and my considering how the world appears to me? When I consider how the stick appears to me—when, that is, I think about the fact that it appears to be bent—am I not reflecting directly on my (potential) reason for forming the belief that the stick is bent? “directly”, in the sense that my reason for forming that belief just is that it appears to me that the stick is bent? And isn’t the judgement that it appears to me as if the stick is bent true just in case my current experience has as part of its content that the stick in question is bent?—Since such a line of thought seems as if it

might be attractive to McDowell, one might suggest that the additional premise needed to complete the argument outlined above is this one:

The content of a judgement about how things appear to me, when such a judgement is correctly made, is the same as the content of one of my perceptual states.

And, certainly, if the content of my judgement about how things appear to me is the same as the content of my perceptual state, then, since the former is conceptual (it is part of the content of a judgement), the latter too must be conceptual.

But this new premise is useless in the context of an argument against the claim that perceptual content is non-conceptual, because it immediately entails the negation of that claim, with no additional premises. (None of the previously mentioned premises was used in the one-sentence argument given at the end of the last paragraph.) And I know of no better candidate. I therefore think it best to construe the various considerations we have been discussing, not as constituting an argument, but as formulating a challenge to explain how, if perceptual content is non-conceptual, judgements about how things appear can play a central role in the rational evaluation of the reasons perception gives us for our beliefs.

Suppose we were to say, with McDowell, that (at least in certain, perhaps central, cases) statements of the form 'It appears to X as if p ' simply record the content of X's perceptual state, in the strong sense that, if it does appear to X as if p , then X's perceptual state must have as (part of) its content the Thought that p . If that were right, it would be obvious how reflection on how things appear to me might constitute reflection on the contents of my perceptual states, and so constitute reflection on the reasons perception provides for various beliefs: When I truly judge that it appears to me as if p , (part of) the content of my perceptual state is precisely that p .—On the other hand, if the content of my perceptual experience is non-conceptual, then, obviously, for me to think about how the stick appears to me cannot be for me to think literally, if I may put it that way, about the content of my perceptual experience: My perceptual state cannot have the content that p ; perceptual experience has the wrong kind of content for that to be possible. Why then should, or how then can, judgements about how things appear play a role in the evaluation of the reasons perception gives me for my beliefs? That is, how can such judgements play the central role the preceding considerations suggest they play?

That this question should arise here signals that Evans has an explanatory commitment that McDowell does not have. The “challenge” is to discharge this commitment.

What we need to explain, to answer McDowell’s challenge, is how, even though my judgements about how things appear cannot even have the same kind of content as my perceptual states, I can yet be thinking about the reasons perception gives me for forming various beliefs—that is, be thinking about how my perceptual experience presents the world as being—when I think about how things appear to me. To give such an explanation, we need to understand what it is to make a judgement about how things appear: More precisely, we need to understand how, although judgements about how things appear do not simply record the contents of our perceptual states, they can nonetheless reflect the contents of those states, so that our judgements about how things appear might still concern how our perceptions present the world as being.

As it happens, Evans himself offers an account of judgements of appearance with these features:

... [A] subject can gain knowledge of his internal informational [e.g., perceptual] states in a very simple way: by re-using precisely those skills of conceptualization that he uses to make judgements about the world. Here is how he can do it. He goes through exactly the same procedure as he would go through if he were trying to make a judgement about how it is at this place now, but excluding any knowledge he has of an *extraneous kind*. (That is, he seeks to determine what he would judge if he did not have such extraneous information.) The result will necessarily be closely correlated with the content of the informational state which he is in at that time. ... This is a way of producing in himself... a cognitive state [i.e., a judgement of appearance] whose content is *systematically* dependent upon the content of the informational [e.g., perceptual] state... (Evans, 1982, pp. 227–8; Evans’s emphasis)

On Evans’s view, the judgements I make about how things appear to me—just like my judgements about how the world is—are based upon my perceptual experience: To say how things now appear to me is to say how I would judge the world to be if I were to judge purely on the basis of my current experience, that is, in such a way that the judgement would be *prima facie*, though defeasibly, justified by that very experience.

It follows that the contents of my judgements about how things appear, when I make them correctly, will track the contents of my perceptual states.³⁰ So, though the (non-conceptual) contents of my perceptual states do not themselves figure in my thought, in reflecting on how things appear, I still reflect on the contents of my perceptual states, on how they present the world as being. I am, moreover, reflecting on what they give me reason to believe, since judgements about how things appear concern the beliefs my experience gives me reason to hold.³¹

According to Evans, then, for me to say, or think, 'It appears to me as if p ' is for me to report, or make, a judgement,³² one that requires me to conceptualize the non-conceptual content of my experience, just as making judgements about the world does: In making such judgements, I exercise the very "skills of conceptualization" I exercise in making judgements about the world. So, if reflection on how things appear necessarily plays a central role in the rational evaluation of the reasons perception gives me for my beliefs; and if, to treat perceptual states as giving me reasons for belief, I must be able rationally to evaluate their bearing on my beliefs; then, if my perceptual states are to give me reasons for my beliefs, I must be able to exercise, and must therefore have, certain conceptual capacities, that is, must possess certain concepts. So the above reflections do show that thinkers must possess certain concepts if their perceptual states are to provide them with reasons for their beliefs. One might think that this establishes, or at least constitutes some evidence for, McDowell's main claim, that perceptual content is conceptual. But this conclusion is of no help to him at all, and we did not need a sophisticated argument to arrive at it. I cannot so much as have a particular belief unless I have certain conceptual capacities, that is, unless I possess the concepts that figure in that belief: It should be no surprise that my evaluating the reasons perception

³⁰ Evans's claim that their content will systematically depend upon that of my perceptual states is too weak: I take this just to be a slip.

³¹ So, on Evans's view, there is something right in the idea that statements of the form 'It appears to X as if p ' mean that X has prima facie justification for believing that p . I do not think we should take that as an analysis of judgements of appearance, though, since I think this analysis would miss the fact that such judgements do make reference to the underlying perceptual states themselves. See note 35, and the text to which it is attached, for more on this issue.

³² There is no reason Evans needs to insist that judgements about how things appear are incorrigible, especially certain, or what have you: So his view need not inherit any of the problems often thought to attach to views which allow real judgements to have such properties.

provides for believing a particular Thought requires me to exercise the same conceptual capacities I must have if I am so much as to entertain that Thought.³³

The answer I am giving to the question, why I don't form the belief that the stick is bent when I am looking at the stick in the water, is thus very close to McDowell's own answer. I think to myself: If I were to judge purely on the basis of my current perceptual state, I would judge that the stick is bent (that is, it appears to me that the stick is bent); however, in the past, I have been in perceptual states in which other sticks, which were partially immersed in water, appeared bent, but in fact were not bent (that is, I have been in states in which I would have so judged, falsely, were I to have judged purely on the basis of my then current perceptual state); moreover, I have learned that straight sticks partially immersed in water always appear to be bent (as do most bent ones); so my being in such a perceptual state does not, in fact, give me sufficient reason, all things considered, to believe that this particular stick is bent.—The only difference between this story and McDowell's, as far as I can see, is that, for him, my judgements about how things now appear, and previously have appeared, record (part of) the contents of certain of my perceptions; on my view, the contents of such judgements conceptualize (parts of) the contents of those perceptions. It wants argument that this difference is epistemologically significant.

Let me be clear that my answer to McDowell's challenge does not rest upon the claim that the reasons perception gives me for forming various beliefs are always of the form that it appears to me as if *p* (for example, that my reason for my belief that the stick is bent is that it appears to be bent).³⁴ One might have been tempted by a reply to McDowell's reflections which would differ from mine in this respect. On this view, forming a belief on the basis of perceptual experience would be a two-step process: First one would form the belief that the stick appears to be bent, and then go on to form the belief that it is bent. This view is tempting because, if it were correct, then for me to reflect on how things appear would be for me to reflect "directly" on my reason for my belief, since my immediate reason would be my judgement about how things appear.

³³ Of course, one might also exercise other conceptual capacities in evaluating these reasons—one might as well bring one's full cognitive capabilities to bear—and one typically will. But this point is of no help to McDowell either.

³⁴ A diagnosis of McDowell's reflections not unlike mine is given by Alex Byrne (1995, pp. 267–8). McDowell (1995, pp. 298ff) responds in a way that suggests he takes Byrne's diagnosis to rest upon the assumption I am here disowning.

But this sort of view is unacceptable. The picture it offers, of how belief is justified by perception, is precisely analogous to the picture the representational theory offers us: I start by making a judgement about my perceptual state (about what its content is, about how it presents the world as being) and then use that as my reason for my belief. But we need no epistemic intermediary between our perceptions and our beliefs, and it would not help if we had one: If there is a problem about how I can form justified beliefs about the world on the basis of my perceptions—one allegedly solved by letting me form them on the basis of judgements about how things appear to me—why is there not a similar problem about how I can form justified beliefs about how things appear? If judgements about how things appear can justifiably be made without any intermediary, why can't judgements about how things are justifiably be made without one, too?

Maybe, then, it is this view—one which treats judgements of appearance as epistemic intermediaries—to which McDowell thinks Evans is forced. Maybe it is Evans's commitment to such a view which constitutes his unintentional commitment to the Myth of the Given. But neither Evans nor I need be sympathetic with this picture. On the contrary, on Evans's treatment of judgements of appearance, which I have adopted, there is no reason for judgements of appearance to be treated as epistemic intermediaries between my perceptual states and the beliefs I form on their basis. According to Evans, when I determine how things appear to me, I determine how I would judge were I to judge simply on the basis of my experience. Now, suppose I make such a judgement: I judge, say, that, were I to judge solely on the basis of present experience, I would judge that my desk is brown. Perhaps, under certain circumstances, I might take this fact to justify me in believing that my desk is, indeed, brown. But typically, that is not what I do: I just judge, on the basis of my present experience, that the desk is brown. It should be clear, on minimal reflection, that Evans's position demands that it should be possible for me to make such a perceptual judgement directly, and immediately, from my perceptual state: Otherwise, the question what I would judge if I were to judge solely on the basis of my current experience would be without content; it would be impossible to make any such judgement.

It is important to understand that what I have just denied is that our judgements about how things appear give us reasons for our perceptual beliefs—not, paradoxically, that appearances do. We must distinguish the fact that it appears to me as if p both from my so judging and from what is required for me to be able to make such judgements. What it

is for it to appear to me as if p is for me to be in a perceptual state which presents the world as being a certain way: The perceptual state cannot itself have the content that p , strictly speaking, since its content is non-conceptual; rather, it appears to me as if p just in case I am in a state on whose basis I would judge that p , were I to judge solely on that basis.³⁵ For me to judge that it appears to me as if p , on the other hand, is for me to recognize that I have perceptual justification for the belief that p .³⁶ The capacity to make judgements of appearance is thus more sophisticated than, and is not (so far as I can see) a necessary precondition of, the capacity to be appeared to, that is, to be in perceptual states. So there is no contradiction in my asserting that (in the usual sorts of circumstances) judgements about how things appear do not, though facts of appearance do, give us reasons for perceptual beliefs.³⁷

There is, though, one further misgiving to which these remarks might give rise. What has just been said commits me to saying that there is a certain sort of gap between the reasons we have for our beliefs and the reasons we can communicate to others: If someone asks me why I believe, say, that there is a desk in front of me, I can do no better than say that it appears to me that there is a desk in front of me (or, less formally, that I see it there); the fact that it does so appear to me, that is, my being in a certain perceptual state, is what gives me a reason for my belief, but I cannot tell you (or myself, for that matter) exactly what content my perceptual state has (that is, exactly how the world does appear to me).—There is a sense, then, in which I cannot really tell you what my

³⁵ It seems to me that one can truly say that it appears to me as if p , even if I lack the concepts required if I am to form that judgement, so long as I would so judge, if I did have those concepts (Wright, 1996, pp. 244–5). But it does not seem to me that it matters for our purposes what view one takes here.

³⁶ I hesitate to say that it is for me to judge that I am in a state which gives me reason to hold the belief that p , because it is not clear to me that Evans's account of judgements of appearance should be taken as an analysis of them. It is not, in particular, clear to me that, in order to be able to make judgements of appearance, a creature has to have such concepts as that of a reason, has to be able to think about how she would judge under certain circumstances, and so on and so forth. Compare note 31.

³⁷ This difficulty may be caused by a linguistic tangle: We speak of beliefs as being reasons for holding other beliefs; but, if I infer q from p , my reason is not that I believe that p , but just that p , where that is something I believe. So it is OK to say that my belief is my reason, so long as it is clear that my reason is *what* I believe, not the fact *that* I believe (though, of course, what I believe it can only be a reason for me so long as I do believe). And similarly, if we say that my perceptual state is my reason, my reason is not *that I am in that state* (i.e., the fact that it appears to me that p) but how things appear.

reason is: I can only tell you that I am currently enjoying a perceptual experience I take to justify that belief (and invite you to get yourself into a similar state).—But it is not as if my reason is utterly inexpressible: I can make it known by talking about how things appear to me, since the contents of my judgements about how things appear, when correctly made, track the contents of my perceptual states. Nor need we worry that, since the content of perception is non-conceptual, I might lack the concepts even to say how things appear to me: Just as McDowell says, I can make use of demonstratives. That is why it is important that Evans is not barred from accepting McDowell's observation that such demonstrative concepts are made available by our experience.

McDowell speaks, at one point, of “a time-honored connection between reason and discourse”, apparently suggesting that only what I can articulate in language can be a reason for me. Perhaps the idea that the reasons we have must be reasons we can communicate is behind this thought. But, whatever its source, I agree with McDowell that the presumption that there is such a connection is part of “the reflective tradition we belong to” (McDowell, 1996, p. 165), especially if by that he means the analytic tradition in philosophy.³⁸ But we will never be able to make sense of perception and its relation to belief until we abandon the idea that there is a constitutive connection between reason and discourse, however time-honored it may be. I do not say that living without this idea will be easy: But a large part of Evans's purpose in *Varieties of Reference* was to show us how we might start to do so, and it is no objection to his project that he was an innovator.

6 Closing

McDowell's objections, both to the Richness Argument and to the overall coherence of the view that the contents of perceptual states are non-conceptual, should therefore be rejected. But I think that it is worth emphasizing that answering them, especially the latter, requires us to learn to take the representational character of non-conceptual content seriously—and, almost as a corollary to that, to recognize at least one attitude whose content is non-conceptual, thereby disposing of the idea

³⁸ The idea that there is such a connection is so deeply embedded in that tradition that—in the form of a thesis that an account of thought can only be given via an account of language—it was once enshrined by Sir Michael Dummett (1993, p. 171) as “the fundamental principle of analytic philosophy”. Alexander George (1997) has some interesting things to say on this point.

that there is a constitutive connection between reason and language. That may require some restructuring of my, and perhaps others', ways of thinking about these matters, and, admittedly, I have hardly indicated what form that might take. Still, I see no reason to think that the necessary re-orientation cannot be accomplished.

In closing, then, let me consider, in a more direct fashion, McDowell's charge that the view developed here is a version of the Myth of the Given. In one place, Evans himself expresses a conception of how perception justifies belief not unlike that sketched above:

The informational states which a subject acquires through perception are *non-conceptual*, or *non-conceptualized*. Judgements *based upon* such states necessarily involve conceptualization: in moving from a perceptual experience to a judgement about the world (usually expressible in some verbal form), one will be exercising basic conceptual skills. But this formulation (in terms of moving from an experience to a judgement) must not be allowed to obscure the general picture. Although the subject's judgements are *based upon* his experience (i.e. upon the unconceptualized information available to him), his judgements are not *about* the informational state. The process of conceptualization or judgement takes the subject from his being in one kind of informational state (with a content of a certain kind, namely, non-conceptual content) to his being in another kind of cognitive state (with a content of a different kind, namely, conceptual content). So when the subject wishes to make absolutely sure that his judgement is correct, he gazes again *at the world* (thereby producing, or reproducing, an informational state in himself); he does not in any sense gaze at, or concentrate upon, his internal state. (Evans, 1982, p. 227; Evans's emphasis)

In a sense, something is "given" on this view: My perceptions, in particular, are something with which I just find myself "saddled". They are merely caused: They are not something I can justify; but nor are they something which it makes sense to ask me to justify (or to expect me to revise). But that does not prevent my perceptions from providing me with reasons for my beliefs about the world, nor from being what justify them. So there is a sense in which Evans treats perception as the (or at least a) foundation for our knowledge of the empirical world. But the foundation is a shaky one, not the foundation of (certain) knowledge that

classical sense-datum epistemologists sought. The reason is not that knowledge is an attitude whose content is conceptual. It is that my perceptual states are about the world: When circumstances are unfavorable, they will be false and so will not constitute (or even be able to yield) knowledge, for familiar reasons.

The real worry about this view, I think, is that my struggle to avoid the Myth of the Given has driven me to the other pole of McDowell's "oscillation", that the view I have outlined is but a trivial variation on Davidson's coherentism:³⁹ Ultimately, it might seem, we have nothing but "exculpations", since our perceptual experience, which is the foundation of much of our knowledge of the empirical world, is merely caused. Faced with this objection, I think we can do no more than insist (not just admit) that we are hostages to the proper workings of our perceptual systems; that their working properly—their representing the world (for the most part) correctly—depends upon our being in the kind of environment in which they were designed to function; that we are, therefore, only less likely than Dennett's frog⁴⁰ to be massively and fundamentally deceived (since our perceptual systems are more sophisticated); and so that, although our cognitive position may in some sense be special—in that we have mental states with conceptual content, whereas (some) other animals may not—that does not make our *epistemological* position special. We must insist, that is to say, that we learn to live with the admittedly uncomfortable thought that, if our perceptual systems do, when all goes well, "lay the world open to us" (and I see no reason I cannot wave McDowell's banner), their doing so is a contingent matter—for which we have only God, evolution, or both to thank—and so something which could, for all we would ever know, be made not to obtain. We can have reasons for what we believe *and* exculpations in the event of massive deception.⁴¹

³⁹ One which includes "experiences" among the things that must cohere, as Wright (1996, pp. 241–2) speculatively suggests.

⁴⁰ Whose tongue lashes out at any moving black thing, and not just at flies.

⁴¹ I owe great debts to Alex Byrne, Jim Pryor, and Susanna Siegel, whose extensive comments on earlier versions of this material led to significant changes (and, I hope, improvements). Thanks also to Tyler Burge, Steven Gross, Eliezer Malkiel, Charles Parsons, Sally Sedgwick, Bob Stalnaker, and Jason Stanley, for their helpful reactions to earlier drafts, and to two anonymous referees, whose comments were unusually useful.

I benefitted from discussions of these ideas with Justin Broackes, Michael Glanzberg, Alva Noë, and Hilary Putnam. Thanks as well to all who attended a reading group at MIT, at which this paper, and Bob Stalnaker's "What Might Non-conceptual Content Be?", were discussed.

References

- Byrne, A. (1995). 'Spin control: Comment on John McDowell's *Mind and World*', *Philosophical Issues* 7: 261–73.
- Crane, T. (1992). 'The nonconceptual content of experience', in T. Crane (ed.), *The Contents of Experience*. Cambridge, Cambridge University Press, 136–57.
- Davidson, D. (1986). 'A coherence theory of truth and knowledge'. Oxford, Basil Blackwell, 307–19.
- Davies, M. (1987). 'Tacit knowledge and semantic theory: Can a five per cent difference matter?', *Mind* 441–62.
- Dummett, M. (1993). 'Language and communication', in *The Seas of Language*. Oxford, Clarendon Press, 166–187.
- Evans, G. (1982). *The Varieties of Reference*, McDowell, J., ed. Oxford, Clarendon Press.
- (1985). 'Semantic theory and tacit knowledge', in *Collected Papers*. Oxford, Clarendon Press, 322–342.
- George, A. (1997). 'Has Dummett over-salted his Frege?', in R. Heck (ed.), *Language, Thought, and Logic: Essays in Honour of Michael Dummett*. Oxford, Clarendon Press, xxx–xx.
- Jackson, F. (1977). *Perception: A Representative Theory*. New York, Cambridge University Press.
- Martin, M. G. F. (1992). 'Perception, concepts, and memory', *Philosophical Review* 101: 745–63.
- McDowell, J. (1995). 'Reply to Gibson, Byrne, and Brandom', *Philosophical Issues* 7: 283–300.
- (1996). *Mind and World*. Cambridge MA, Harvard University Press.
- Peacocke, C. (1992). *A Study of Concepts*. Cambridge MA, MIT Press.

Many of the views expressed here were developed while I taught, jointly with Alison Simmons, an undergraduate seminar entitled "Consciousness and Content", at Harvard University, in the Spring of 1997. Special thanks are due to Alison for discussions about perception she and I have had over the last few years, but especially for those which occurred during that course.

- Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge, Cambridge University Press.
- Stalnaker, R. (1998). 'What might nonconceptual content be?', in E. Villanueva (ed.), *Concepts*. Atascadero CA, Ridgeview Publishing, 339–52.
- Wright, C. (1993). 'Theories of meaning and speakers' knowledge', in *Realism, Meaning, and Truth*. Cambridge, Blackwell, 204–238. 2d ed.
- (1996). 'Human nature?', *European Journal of Philosophy* 4: 235–54.