

Justified vs. Warranted Perceptual Belief: Resisting Disjunctivism

JUAN COMESAÑA

University of Wisconsin—Madison

In this paper I argue that McDowell's brand of disjunctivism about perceptual knowledge is ill-motivated. First, I present a reconstruction of one main motivation for disjunctivism, in the form of an argument that theories that posit a "highest common factor" between veridical and non-veridical experiences must be wrong. Then I show that the argument owes its plausibility to a failure to distinguish between justification and warrant (where "warrant" is understood as whatever has to be added to true belief to yield knowledge).

1. Introduction

"Disjunctivism" about perceptual knowledge (roughly, the thesis that there is nothing epistemically relevant shared by cases of veridical and non-veridical experience) is a perplexing theory. It seems to entail, for example, that perceptual knowledge cannot be based on non-conclusive reasons (more on this entailment below). Disjunctivism is presented as a reaction to theories that posit a "highest common factor" between cases of veridical and non-veridical experience (in what follows, "HCF theories").¹ In this paper I argue that McDowell's brand of disjunctivism about perceptual knowledge is ill-motivated. First, I present a motivation for disjunctivism, in the form of an argument that HCF theories must be wrong. The interpretation of disjunctivism that results is charitable in that it offers plausible interpretations of some of its most seemingly implausible claims, and also in that it offers arguments for some of the disjunctivist theses. Plausible and interesting they are, but false and unsound. A main thesis of this paper is that one main argument in favor of disjunctivism owes its plausibility to a failure to distinguish

¹ Disjunctivism is sometimes intended (by McDowell, for example) as a general view, encompassing knowledge in general and not only *perceptual* knowledge. In this paper I focus on perceptual knowledge because it is also disjunctivists' usual focus, and because there is no reason to think that the conclusions reached in the case of perceptual knowledge cannot be generalized.

between epistemic justification and warrant (understood as whatever turns true belief into knowledge).²

In the next section I present an HCF theory—because disjunctivism is best seen against the background of HCF theories and the thoughts that motivate them. Then, in sections 3 through 5, I characterize McDowell's disjunctivism and I reconstruct the main arguments for being a disjunctivist in this sense. Finally, in the last section I explain why I think those arguments don't work.

2. Highest common factor theories

One of the thoughts that motivates HCF theories is that there must be something in common between cases of veridical and non-veridical experience, given that some veridical experiences are phenomenologically indistinguishable from some non-veridical experiences. That thought is captured by the following characterization:

HCF (first pass): For any subject *S* and any pair of experiences of *S*, *x* and *y*, if *x* and *y* are phenomenologically indistinguishable, then there is something in common between *x* and *y*.

Note that, on pain of making HCF trivial, the thing that phenomenologically indistinguishable experiences share must be something different from, and not entailed by, their being phenomenologically indistinguishable. Everybody should accept that if *x* and *y* are phenomenologically indistinguishable then they share at least two properties: the property of being phenomenologically indistinguishable from *x* and the property of being phenomenologically indistinguishable from *y*. The proponent of HCF wishes to go beyond this platitude, to something that *explains* the phenomenological similitude.

There are two different kinds of things that proponents of HCF have advanced as fit to play the role of the "highest common factor," the thing that is shared by phenomenologically indistinguishable experiences: objects, such as sense-data or ideas, on the one hand, and states, such as ways of being appeared to or ways things look, on the other. In what follows, I will focus on ways things look. Consequently, I will replace the above formulation of HCF with:

HCF1: For any subject *S* and any pair of experiences of *S*, *x* and *y*, if *x* and *y* are phenomenologically indistinguishable for *S*, then there is a way things look to *S* that is the same in *x* as in *y*.

²

I will be focusing on McDowell as the main disjunctivist. Another attack on HCF theories comes from Hilary Putnam. Cf. his Dewey Lectures, "Sense, Nonsense, and the Senses: An Inquiry Into the Powers of the Human Mind," reprinted in his *The Threefold Cord: Mind, Body, and World* (New York: Columbia University Press, 1999).

HCF1 is an ontological thesis, but what bothers disjunctivists about HCF1 theories is that it enables an epistemological commitment: that the factors identified in HCF1 as common to veridical and non-veridical experiences constitute the epistemic basis of our perceptual beliefs. According to this epistemological use of the HCF theory, whether I genuinely perceive something, say, a tomato or, instead, I am having a phenomenologically indistinguishable hallucination, if I form the belief that there is a tomato in front of me based on the fact that it looks to me as if there is a tomato in front of me, then my belief is *prima facie* justified. It need not be the case that I consciously (or even unconsciously) *infer* that there is a tomato in front of me from the presence of those factors, or that I am in any sense aware of the fact that the factors obtain, but just that those factors will *explain* the belief in question. We have here, then, another thesis related to the idea that there is a highest common factor:

HCF2: The way things look to a subject is the epistemic basis for his perceptual beliefs.

The two highest common factor theses, together with intuitively plausible claims, entail that perceptual knowledge is fallible. I can, for example, know that there is a tomato in the refrigerator based on something that doesn't *entail* that there is a tomato in the refrigerator. The implication is more or less obvious, but it will be instructive to spell it out in some detail. HCF2 entails the following:³

(1) If a subject S perceptually knows that p, then his belief that p is based on the fact that it looks to S as if p.

Moreover, from HCF1 and the fact that the subject could have a non-veridical experience as if p which is indistinguishable from a veridical experience that p, it follows that:

(2) It is possible that it looks to S as if p while p is false.

And now, (1) and (2) entail that perceptual knowledge is fallible, in the following sense:

(3) If a subject perceptually knows that p, then his belief that p is based on something compatible with the falsity of p.

³ I am assuming here that the content of experience can be described propositionally—which need not be equivalent to assuming that experience has propositional content.

When two cases are related such that one is a genuine perception that *p* and the other is a phenomenologically indistinguishable hallucination that *p*, I will call the first case a “good case” and the second one a “bad case.” With that in mind, the above considerations can be summarized in the following table (letting ‘*p*’ stand for any perceptually acquired proposition, and where an ‘X’ means that the situation described on the left obtains in the case named above):

| | Good Case | Bad Case |
|-------------------------------------|-----------|----------|
| p | X | |
| It looks to me as if p | X | X |
| I justifiably believe that p | X | X |
| I know that p | X | |

Table 1

3. The case against the existence of a highest common factor: perceptual knowledge and warrant

We might be tempted to think that the following proposition follows from (3) (or even that it is just another way of formulating the fallibility of perceptual knowledge that (3) expresses):

(4) The truth of the belief involved can be the only epistemically relevant difference between a case of knowledge and a case where knowledge is absent.

(4) might seem obviously true, on the face of (3) or perhaps on the face of table 1. But McDowell thinks that (4) is false. Now, if (4) is false and it follows from (3), then we must reject (3). This argument against (3) is one of McDowell's main weapons against HCF theories.⁴ What McDowell thinks is that HCF theories lead inevitably to some kind of fallibilism like that expressed by (3), and that this is bad because (3) entails (4), and (4) is false. In what follows, I will argue that McDowell is right in rejecting (4). However, I will argue that (4) *doesn't* follow from (3). Thus, even if (4) is false,

⁴ See John McDowell, “Criteria, Defeasibility, and Knowledge,” *Proceedings of the British Academy* 68 (1982), reprinted in his *Meaning, Knowledge, and Reality* (Cambridge, MA: Harvard University Press, 1998), pp. 369-94 (see esp. p. 371 and p. 389); “Knowledge and the Internal,” *Philosophy and Phenomenological Research*, Vol. LV, No. 4 (1995), reprinted in *Meaning, Knowledge, and Reality*, pp. 395-413 (see esp. p. 399); and *Mind and World* (Cambridge, MA: Harvard University Press, 1996) (see esp. Lecture VI, pp. 111-13). In section 6 below, I consider a different interpretation of these passages (especially those from “Criteria, Defeasibility, and Knowledge”). That other interpretation might well be closer to McDowell's intentions, but it also makes his position weaker.

(3) can still be true, and so McDowell's misgivings about HCF theories are misplaced.

It is important to realize that the comparison between the good case and the bad case does *not* constitute a proof of (4). Consider a third kind of case, very much like the bad case but with an important difference: suppose that my friends wish to play a joke on me and put a papier-mâché tomato in my refrigerator that is phenomenologically indistinguishable from a real tomato. But suppose also that they don't remove the original tomato from the refrigerator, but instead they put it behind the bottle of orange juice, outside of my visual field. In this case, my belief that there is a tomato in the refrigerator is true, but I still don't know it. For obvious reasons, I'll call this case "the Gettier case."⁵ Let's add it to our table:

| | Good Case | Bad Case | Gettier Case |
|-------------------------------------|-----------|----------|--------------|
| p | X | | X |
| It looks to me as if p | X | X | X |
| I justifiedly believe that p | X | X | X |
| I know that p | X | | |

Table 2

The existence of this Gettier case shows that the truth of my belief cannot be the only epistemically relevant factor that distinguishes the good case from the bad case. For, if it were so, adding truth to the bad case would have to result in a case of knowledge, but that is not always so, because the Gettier case is not a case of knowledge.

So far, I have only shown that, in one particular instance, the comparison between the good and the bad case doesn't establish (4). That certainly shows that one reason on which (4) might be based is a bad reason, but it doesn't show that (4) is false. But the case against (4) can be strengthened if it could be shown that, given any good case/bad case pair, it is always possible to construct a third, Gettier-like case. If so, we would have shown that, for any good case/bad case pair, the difference between them must include factors other than the truth of the belief in question.

An argument along those lines can be constructed against (4). As Plato noted, not every true belief constitutes knowledge. As Gettier noted, moreover, not every justified true belief constitutes knowledge. But surely *some*

⁵ Cf. Edmund Gettier, "Is Justified True Belief Knowledge?," *Analysis* 23, no. 6 (1963), pp.121-3. Surprisingly, Gettier is not mentioned in any of the works by McDowell that I cite in this paper. By contrast, he will play a large role both in my exposition of the argument against (4) and in my defense of HCF theories.

true beliefs constitute knowledge! Therefore, there must be some property that true beliefs can have such that, if a true belief has that property, then it is known. Let's call that property 'warrant.'⁶ (By saying that warrant is a property I do not mean to imply that it is not disjunctive, perhaps highly so, or otherwise gerrymandered. In fact, it is almost surely gerrymandered.) Knowledge, then, by definition, is warranted true belief. In particular, perceptual knowledge is warranted and true perceptual belief.

The argument against (4) is the following:

- (i) (4) The truth of the belief involved can be the only epistemically relevant difference between a case of knowledge and a case where knowledge is absent. (Assumption for reductio)
- (ii) Knowledge is warranted true belief. (Definition of 'warrant')
- (iii) Whether a belief is warranted or not is epistemically relevant. (Premise)
- (iv) There can be false but warranted beliefs. (From (i), (ii), and (iii))
- (v) There cannot be false but warranted beliefs. (Premise)
- (vi) There can be false but warranted beliefs and there cannot be false but warranted beliefs. (Conjunction (iv) and (v))
- (vii) (4) is false. (Reductio ad absurdum, (i)-(vi))

Obviously, (v) is the crucial step, and I'll turn to it in a moment. First, let me clarify the only other nontrivial step of the argument, the inference to (iv). (i) says that it is possible for there to be two cases, let's call them case *a* and case *b*, such that: *a* is a case of knowledge, *b* is not a case of knowledge, and the only epistemically relevant difference between *a* and *b* is the truth of the belief involved—that means, *inter alia*, that the belief involved is the same. Now, given that, by (ii), if S knows that *p* then S's belief that *p* is warranted, it follows that, in *b*, the belief is warranted—for the *only* epistemically relevant difference between the cases was the truth of the belief involved, and, by (iii), whether a belief is warranted or not is epistemically relevant—and, therefore, it also follows that there can be false but warranted beliefs.

⁶ The terminology is from Plantinga, *Warrant and Proper Function* (New York: Oxford University Press, 1993). Other philosophers use the term 'warrant' to refer to a kind of positive epistemic status that is different, in some sense or other, from justification, but that is not necessarily such that it transforms a true belief into knowledge. In this work, I use 'warrant' exclusively in Plantinga's sense.

Could the defender of (4) reject (iii), the idea that whether a belief is warranted or not is not epistemically relevant? Hardly so: there seems to be no interesting sense of 'epistemic' in which truth counts as epistemically relevant but warrant doesn't, and the defender of (4) is committed to the claim that truth is epistemically relevant.

Let's now turn to (v), the claim that necessarily, if a belief is warranted, then it is true. Trenton Merricks has presented an intriguing argument for this thesis,⁷ which goes as follows. Suppose that it is possible for there to be a warranted but false belief. Suppose, in particular, that S's belief that p is warranted but false. S can infer that either p or q, where q is any truth for which S has no independent warrant. It seems, then, that S's belief that p or q is now warranted and true, but yet it is not knowledge. But that is absurd.

The argument relies on the principle that warrant can be transferred across (at least some cases of) known logical implication, a principle that has been denied. But Merricks offers another argument which doesn't rely on this principle. Suppose that on Monday I believe that Jones owns a Ford, and that my belief is warranted but false. On Tuesday, I still have the same belief, but, unbeknownst to me, Jones's aunt has died and left him a Ford. On Tuesday, I have a warranted true belief, but yet it is not knowledge. But that is absurd.

Merricks's guiding idea is that if warrant didn't entail truth, then it would be possible to construct a Gettier case for warrant: a case where we have a warranted belief that is only *accidentally* true (roughly, it is true for reasons that are not connected in the appropriate way to our holding the belief): either the same belief at some later time, or a belief inferred from it. But that doesn't make sense, because warrant is simply *defined* as whatever turns true belief into knowledge.

Now, as I said, it has been denied that warrant can be transferred even across *known* logical implication. Against Merricks's second argument, the one that doesn't rely on implications, the defender of the possibility of false but warranted beliefs can hold that what is really impossible are warranted but *accidentally true* beliefs. A false belief can be warranted, but if it were true (in another possible world), or if it were to become true (at a later time), then it would either lose its warrant or its truth would be connected in an appropriate way to its being believed. That is, the defender of the possibility of false but warranted beliefs would have to deny a principle that Merricks says he has to accept:

⁷ Trenton Merricks, "Warrant Entails Truth," *Philosophy and Phenomenological Research*, Vol. LV, No. 4 (1995), pp. 841-55.

(5) If the belief that *p* is warranted, false, and possibly accidentally true, then it is possible that the belief that *p* is warranted and accidentally true.⁸

Merricks's argument for the truth of (5) is that a false but warranted belief has the same connection to the truth of the belief as an accidentally true but warranted belief would have, namely, none, and so that it couldn't be that only one of them is really warranted. But it seems to me that Merrick overlooks the possibility that a belief's being warranted entails that if it were true then its truth would have to be connected in some way to its being believed (though if it is false, of course, its truth cannot be connected in *any* way to its being believed). Given that, as Merricks himself reminds us, "warrant may have some odd properties,"⁹ it would be unwise to suggest that it couldn't possibly behave such as to make (5) false.

So, unlike Merricks, I don't think that it follows from the very idea of warrant that there cannot be false but unwarranted beliefs—that is, I don't think that warrant *entails* truth. Nevertheless, as a substantive independent thesis, I think that the claim that there cannot be warranted but false beliefs is very plausible. Indeed, many different accounts of warrant have as a consequence that there cannot be false but warranted beliefs: Merricks cites the works of Pollock, Lehrer, Nozick, Dretske, Goldman, and Tomberlin, and Sosa could be added to the list.¹⁰ So while I disagree with Merricks's defense of premise (v), I still think that the claim is more plausible than its negation. In conclusion, though he doesn't use them, McDowell has at his disposal good reasons against (4).

4. An alternative to highest common factor theories: mere appearances vs. facts that make themselves perceptually manifest

Let me restate, for ease of reference, some relevant propositions:

HCF1: For any subject *S* and any pair of experiences of *S*, *x* and *y*, if *x* and *y* are phenomenologically indistinguishable for *S*, then there is a way things look that is the same in *x* as in *y*.

⁸ Merricks, *op. cit.*, p. 851. The principle is numbered (7) in Merricks's paper.

⁹ Merricks, *op. cit.*, p. 849.

¹⁰ Cf. John Pollock, *Contemporary Theories of Knowledge* (Totowa, New Jersey: Rowman and Littlefield, 1986); Keith Lehrer, "Knowledge Reconsidered," in M. Clay and K. Lehrer (eds.), *Knowledge and Skepticism* (Boulder: Westview, 1989); Alvin Goldman, "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73 (1976), pp. 771-91; James Tomberlin, "Critical Review of Carl Ginet's *Knowledge, Perception, and Memory*," *Nous* 14 (1980), pp. 239-50; and Ernest Sosa, "Skepticism and Contextualism," *Philosophical Issues* 10 (2000), pp. 1-18, and his "Replies" in the same volume, pp. 38-42; see esp. pp. 39-40, where Sosa explicitly acknowledges, in reply to comments by Tomberlin, that his notion of warrant entails truth.

HCF2: The way things look to a subject is the epistemic basis for his perceptual beliefs.

(1) If a subject S perceptually knows that p, then his belief that p is based on the fact that it looks to S as if p.

(2) It is possible that it looks to S as if p while p is false.

(3) If a subject perceptually knows that p, then his belief that p is based on something compatible with the falsity of p.

(4) The truth of the belief involved can be the only epistemically relevant difference between a case of knowledge and a case where knowledge is absent.

Holding that (4) is false is the first step in the larger argument for disjunctivism that I am considering. The second step is to say that, given that (4) follows from (3), we must reject (3). This rejection of (3) would lead us, ultimately, to reject either HCF1 or HCF2, or both—or so McDowell hopes.

Now, given that (1) and (2) entail (3), it seems that McDowell is forced to reject at least one of (1) and (2). But McDowell, it seems, can have his cake and eat it too. For there is a certain ambiguity in (1) and (2), and McDowell can hold that on the disambiguation that makes them true, they don't entail (3). So McDowell is going to reject HCF1, HCF2 and (3), but he is going to accept (1) and (2), at least under a certain disambiguation. Here's how.

When McDowell rejects HCF1, he does it because he thinks that a veridical experience can be phenomenologically indistinguishable from a non-veridical one, but he also thinks that there is nothing (no *thing* or *state*) in common between them. If a subject is having a veridical experience that p, McDowell will say that “the fact that p is making itself perceptually manifest” to the subject, and if a subject is having a non-veridical experience, then he will say that the subject is having “a mere appearance” as if p.¹¹ So McDowell thinks that it can be phenomenologically indistinguishable for a subject whether the fact that p is making itself perceptually manifest to him or whether he is having a mere appearance as if p.¹² But the subject perceptually knows that p if, and only if, he has a perceptual belief that p and the fact that p is making itself perceptually manifest to him; and, if he is having a mere appearance as if p, then he doesn't perceptually know that p. So, apparently, McDowell thinks that all of the following are necessary truths:

¹¹ Cf., for example, McDowell, “Criteria, Defeasibility, and Knowledge,” p. 386.

¹² Cf. McDowell, “Criteria, Defeasibility, and Knowledge,” p. 389.

(6) If a subject *S* has the perceptual belief that *p*, then either he is having a mere appearance as if *p* or the fact that *p* is making itself perceptually manifest to *S*.

(7) A subject *S* perceptually knows that *p* if and only if he has a perceptual belief that *p* and the fact that *p* is making itself perceptually manifest to *S*.

(8) If a subject *S* is having a mere appearance as if *p*, then *S* doesn't know that *p*.

“The way things look” can refer either to facts making themselves manifest or to mere appearances, and while it is true that the epistemic basis of perceptual beliefs that amount to knowledge are facts making themselves manifest, McDowell thinks that mere appearances don't provide an epistemic basis for beliefs. Of course, mere appearances can explain why we have certain beliefs, but the explanation will be, for McDowell, more like an explanation in terms of a brain tumor that produces random beliefs than like a rationalizing explanation.

Now, having rejected HCF1 and HCF2, there is no need to interpret “it looks to *S* as if *p*” univocally in (1) and (2). Indeed, having distinguished between facts making themselves perceptually manifest to a subject and the subject's having a mere appearance, McDowell can now show that (1) and (2) can each mean two different things:

(1a) If a subject *S* perceptually knows that *p*, then his belief that *p* is based on the fact that *p* is making itself perceptually manifest to *S*.

(1b) If a subject *S* perceptually knows that *p*, then his belief that *p* is based on the fact that *S* is having a mere appearance as if *p*.

(2a) It is possible that *S* is having a mere appearance as if *p* while *p* is false.

(2b) It is possible that the fact that *p* is making itself perceptually manifest to *S* while *p* is false.

It is now clear that (1) and (2) entail (3) only if they are disambiguated in the same way. Now, given (7), McDowell has to reject (2b), and he has no problem in accepting (2a). But, given (8), he must reject (1b), though he will be happy to accept (1a). So McDowell has shown that the rejection of (3) doesn't commit him to rejecting either (1) or (2)—at least on certain interpre-

tations of them, interpretations that are made available by the distinction between facts making themselves perceptually manifest and mere appearances.

5. Taking stock

The dialectical status of the discussion between McDowell and the defender of HCF theories is now complicated, so let me state how I see it. HCF1 and HCF2 are initially appealing—they seem to correctly describe the phenomenology of good and bad cases. These HCF theses imply, via (1) and (2), a principle of fallibility: (3). That principle of fallibility might seem to imply, or even be equivalent to, (4). But there is a powerful argument against (4), which uses the crucial premise that there cannot be warranted but false beliefs. A disjunctivist like McDowell can use this argument to cast doubt on (3) itself, and so, a bit more indirectly, on HCF1 and HCF2. But now (1) and (2) enter the fray. Given the high plausibility of these theses, they play a double role: they support (3) by implying it, and they also support HCF1 and HCF2 by being implied by them (i.e., by playing the role of premises in an argument to the best explanation). McDowell can now dig in his heels and protest that the argument against (4) upsets any plausibility that (1) and (2) can have. But he can also show that he can accept (1) and (2), at least when they are disambiguated according to the distinction between the two ways in which a subject can have an appearance as if *p* that he wishes to introduce. The question now is whether (1) and (2) are at most as plausible as (1a) and (2a) or whether their plausibility has to be explained in other terms, perhaps in terms of HCF1 and HCF2.

The plan for the last section of the paper is as follows. I will argue, against McDowell, that (3) doesn't imply (4). Therefore, even if, as I believe, (4) is false, that doesn't have any tendency to show that (3) is false, and, hence, there is not here even an indirect argument against HCF1 and HCF2 and no need to distinguish between (1a) and (2a) on the one hand and (1b) and (2b) on the other. Showing this involves appealing once again to the distinction between justification and warrant, a distinction that McDowell might find dubious because of a certain kind of epistemic internalism that apparently grounds it and which, according to McDowell, leads to an unsatisfactory position in the dialectic against the skeptic. My answer to this charge will be double: first, the distinction between justification and warrant is not really grounded on any kind of internalism about justification; and second, even if it were so grounded, McDowell's attack on that kind of internalism is misguided.

6. The highest common factor defended: justification, warrant, and skepticism

McDowell thinks that (4) is false, and I agree with him, at least if “epistemic relevance” is understood so that, for example, truth can be epistemically relevant. Now, a crucial premise in the argument against (4) that I used on behalf of McDowell is the idea that there cannot be warranted but false beliefs. But I also think that (3) is true, that is, that the basis of perceptual beliefs is compatible with the falsity of those beliefs. These two commitments entail a third: that the basis of perceptual knowledge that *p* must be different from the warrant for *p*. I am happy to accept that third commitment, for there is justification aside from warrant.

We all have a pre-theoretical understanding of what epistemic justification amounts to. That pre-theoretical understanding is, for example, clear enough to make Gettier cases possible. Of course, it would be nice to have a *theory* of epistemic justification, but (aside from a few remarks that I will make in a moment) I can't present such a theory here, and I don't need to present it, for every claim that I will make can be tested against that pre-theoretical understanding.

I propose to interpret HCF2 in terms of justification. That interpretation will in turn be transmitted to (1) and (3). More specifically, I suggest that they be interpreted as follows:

HCF2*: When a subject *S* has a justified perceptual belief that *p*, his justification for it consists in that it looks to him as if *p*.

(1*) A subject *S* can know a perceptual proposition that *p* with his justification being that it looks to *S* as if *p*.

(3*) A subject can know a perceptual proposition that *p* with his justification being compatible with the falsity of *p*.

But now (3*) would entail (4) only if knowledge were justified true belief, which we know is not.

Let us compare McDowell's notions of facts making themselves perceptually manifest and enjoying a mere appearance as if *p*, on the one hand, with the notions of justification and warrant on the other. McDowell's notion of a fact making itself perceptually manifest and the notion of warrant that I have been using are importantly tied together. For instance, both being warranted in believing that *p* and the fact that *p* making itself perceptually manifest to you cannot occur without *p* being true. Now, HCF theories, at least as I am characterizing them here, hold that, although justification doesn't entail warrant, it is, of course, compatible with it. On the other hand, McDowell's disjunctivism entails that enjoying a mere appearance as if *p* is incompatible

with knowing that p —and so, McDowell seems to think, not really epistemically interesting. McDowell's main concern in epistemology is with a guaranteeing notion of positive epistemic status, and he thinks little of any other putative epistemic status that can fall short of truth.

This disdain for any non-guaranteeing notion of justification derives from the fact that McDowell's main concern is not epistemological (not, at least, if epistemology is narrowly construed), but lies rather in the philosophy of mind and language; more precisely, McDowell's main concern is with the idea of "openness to the world," or with what it is for a mental episode to be "directed towards the world," with what is sometimes called "the problem of intentionality."¹³ In McDowell's mind, any notion of justification that doesn't entail truth will leave it a mystery how thought about the world is possible. It will leave it a mystery, McDowell sometimes says, because such a notion of justification will have to be an internalist notion, and internalism leads to skepticism, and the main result of skeptical considerations is, precisely, that thought about the world comes to seem impossible.¹⁴ As I said above, I will argue that the notion of justification that I have attributed to HCF theories need not be internalist, but also that, if it were internalist, that wouldn't make HCF theories more vulnerable to skepticism than McDowell's own disjunctivism.

¹³ Cf. *Mind and World*, *passim*, but especially the Introduction, section 3.

¹⁴ This is the McDowell of "Knowledge and the Internal," where he says that highest common factor theories can arise when "the Sellarsian [...] image of standings in the space of reasons undergoes a certain deformation [...] The deformation is an interiorization of the space of reasons, a withdrawal of it from the external world" (p. 395). The McDowell of "Criteria, Defeasibility, and Knowledge," on the other hand, seemed attracted to some form of internalism: "When someone has a fact made manifest to him, the obtaining of the fact contributes to his epistemic standing on the question. *But the obtaining of the fact is precisely not blankly external to his subjectivity [...]*" (pp. 390-1, emphasis added). This latter quote suggests the idea that McDowell is making an unlikely combination of epistemic internalism with radical content externalism. But it seems to me that McDowell is rejecting precisely such an interpretation in "Knowledge and the Internal," where, in response to a suggestion by Blackburn "according to which, instead of a shrinkage in what can be known, the mind [...] expands to 'embrace' all sorts of worldly states of affairs," he says: "This idea, which Blackburn rightly finds bizarre, has nothing to do with what I am proposing here, and was proposing in the work that Blackburn is discussing ("Criteria, Defeasibility, and Knowledge")," p. 406, n. 15. (Blackburn's paper is "Knowledge, Truth, and Reliability," *Proceedings of the British Academy* LXX (1984), pp. 167-87.) There is a further puzzle, for McDowell is clearly unapologetic about the spatial imagery—as he should be, for he had used "embrace" in "Criteria, Defeasibility, and Knowledge," p. 386, implying quite clearly that the mind does embrace worldly states of affairs—but he says that Blackburn's gloss on it depends on "a gross missing of its point" ("Knowledge and the Internal," p. 408, n. 20). But that the world must do us a favor if we are to enjoy "factive standings in the space of reason" (the point that, supposedly, Blackburn is missing) doesn't make it any less (or more) bizarre that those standings consist in our minds embracing worldly states of affairs.

Is the notion of justification that I am using an “internalist” notion? Internalism is said in many ways, but mainly in two.¹⁵ “Accessibilism” is the internalist view that the factors that justify a person’s beliefs must be factors to which the person has some special sort of epistemic access. “Mentalism,” on the other hand, is the internalist view that the factors that justify a person’s beliefs must be internal to (supervene upon) that person’s mental life. So, is the notion of justification that I need a version of either accessibilism or mentalism? Not necessarily. There are at least two accounts of what it is about the way things look that enables them to justify beliefs about how things are, and in one of those accounts the justificatory power of the way things look supervenes upon factors that are *external* to the subject’s mental life.

Both accounts agree in that the fact that things look to S as if p justifies S’s belief that p. One approach, Conee and Feldman’s “Evidentialism,” takes this justificatory connection to be a primitive fact, which doesn’t admit of (nor requires an) explanation.¹⁶ Evidentialism is indeed an internalist theory of the mentalist kind. But the other approach takes the justificatory connection between how things look to S and what S believes to supervene upon, and be explained by, the reliable connection that exists between how things look to S and how things are. Very roughly put, it is because things look pretty much the way they are that it would be rational for you to believe that things are thus and so on the basis of the fact that it looks to you as if things were thus and so.¹⁷ This second, reliabilist approach to perceptual justification is entirely compatible with HCF theories, and, if correct, it would show that those theories need not be internalist. This reliabilist account is indeed the account that I favor, but it would be instructive to see that McDowell’s attack on HCF theories fails even if he were right in thinking that those theories must be, in the end, internalist.

¹⁵ Cf. Earl Conee and Richard Feldman, “Internalism Defended,” *American Philosophical Quarterly* 38 (2001), pp. 1-18. The characterizations of Accessibilism and Mentalism are paraphrases of their words in that paper. Ernest Sosa makes a similar distinction between what he calls Cartesian and Chisholmian internalisms in his “Skepticism and the Internal/External Divide,” in Greco and Sosa (eds.), *The Blackwell Guide to Epistemology* (Oxford: Blackwell, 1999), pp. 145-57.

¹⁶ See Conee and Feldman, “Evidentialism,” *Philosophical Studies* 48 (1985), pp. 15-34.

¹⁷ A problem arises for this kind of reliabilism. Namely, the reliabilist approach would have to either account for or explain away the intuition that, e.g. an evil demon of the kind Descartes imagined could deprive someone of *knowledge*, though it won’t deprive her of *justification*—as opposed to the case of the brainwashed, who lack justification as well as knowledge—see Stewart Cohen, “Justification and Truth,” *Philosophical Studies* 46 (1984), pp. 279-95. I present a version of reliabilism that is immune to this problem (and that can also help us explain the appeal of skeptical arguments) in “The Diagonal and the Demon,” *Philosophical Studies* 110 (2002), pp. 249-66. This version of reliabilism does count as mentalist in the supervenience sense, but, as I argue elsewhere (“We Are (Almost) All Externalists Now,” *Philosophical Perspectives* 19 (2005), pp. 59-76), this doesn’t mean that it is an internalist theory.

What is it about internalism that McDowell finds objectionable? Metaphorically put, that it creates a gap between mind and world that will be impossible to close. The gap, if there is one, is not ontological, however, for in defending the existence of a highest common factor I have not defended any kind of dualism. 'Internalism' might be a misleading term here, for it indeed evokes the idea that what justifies your beliefs is locked inside something, your mind perhaps, which is of a different nature than the rest of things. But internalism, of course, is compatible with reductionism about the mind. It is also compatible, needless to say, with dualism, but it doesn't imply it—nor, I think, does it even suggest it.

If there is any gap, therefore, it is an epistemic gap: that is, there are justified true beliefs that are not knowledge. McDowell might now accept the offer of an epistemic gap and try to widen it. Given that a belief can be fully justified and yet fall short of the fact itself, giving this notion of justification a place in epistemology, McDowell might say, invites skepticism. When we take ourselves to be in touch with how things are because of how things look, whether we *are* in touch or not is something that is outside our control, in the sense that, whatever error-preventing measures we take, there is nothing we can do to *guarantee* that we are in touch.¹⁸ In McDowell's own words:

[I]n view of the defeasibility of the inferential relation [between, e.g., how things look and how things are]... we have no real explanation of how the relation could be knowledge sustaining; for even in the most favourable cases it remains possible, for all one knows, that the beliefs about what is said and done [or about how things look] which are to stand as one's basis (in virtue of their role as criteria) are both true and known to be true, while the conclusions which rest on them are false and hence not known at all.¹⁹

And, a few pages later in the same paper, McDowell asks a rhetorical question:

¹⁸ This is not intended to imply that we can guarantee that we are *justified* by properly taking account of all the error-preventing measures available to us—this idea is rightly attacked by McDowell in "Knowledge and the Internal," and also in "Knowledge by Hearsay," in B. K. Matilal and A. Chakrabarti (eds.), *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony* (Dordrecht: Kluwer, 1993), reprinted in *Meaning, Knowledge, and Reality*, pp. 414-43. But McDowell makes a strange inferential move that, in the terminology of this paper, could be put thus: given that justification is *not* a luck-free zone, given that the world must make us a favor for us even to be justified in our beliefs, then we don't really need the notion of warrant, for our being justified can, when the world *does* cooperate, just be our having knowledge. If this indeed is McDowell's line of reasoning, then it rests on the surely false presupposition that there is only one kind of epistemically relevant luck—for it presupposes that the luck involved in being justified is the same as the luck involved in being warranted.

¹⁹ This passage (which, as well as the next, was excised from the version included in *Meaning, Knowledge, and Reality*) appears on p. 209 of the version of "Criteria, Defeasibility, and Knowledge" included in Jonathan Dancy (ed.), *Perceptual Knowledge* (New York: Oxford University Press, 1988), pp. 209-19.

Consider a situation in which someone supposedly has such knowledge, as compared with a situation in which someone experiences the same independently ascertainable circumstance but, since the state of affairs for which it is a 'criterion' does not obtain (the defeasible support is defeated), he lacks the knowledge in question: can the blankly external obtaining of the fact, in the first case, make it intelligible that the subject knows, though he is supposedly indistinguishable, in the reach of his experience, from the subject who does not know?²⁰

Note, first, that the accusation of being friendly to skepticism is not a fair accusation. In the first passage quoted, McDowell seems to imply that, because the relation between how things look and how they are is less than perfectly reliable, "for all one knows" things might be different from how they look. There is a crucial ambiguity here, which has to do with the scope of the epistemic possibility signaled by "for all one knows." If the scope is, as suggested, that things sometimes are different from how they look, then it is surely true but it doesn't entail any serious kind of skepticism. If, on the other hand, McDowell is saying that, *whenever* one believes that *p* on the basis that it looks to one as if *p*, "for all one knows" *p* is false, then that might entail a strong kind of skepticism, but it is false. For, when I perceptually know that *p*, I believe that *p* on the basis that it looks to me as if *p*; and in this case the falsity of *p* is obviously *not* compatible with everything I know.

There is a similar ambiguity in the second passage. It is not clear why McDowell assumes that the obvious answer to his question in that passage is "No." The vital notion now is that of two subjects being indistinguishable "in the reach of [their] experience." If that means that, e.g., one of them is seeing a tomato if and only if the other one is also seeing a tomato, then the answer to McDowell's question might well be "No," but it is a "No" that the HCF theorist will happily accept. If, on the other hand, that two subjects are indistinguishable in the reach of their experience just means that it looks to both of them as if *p*, then it is by no means clear that one of them knows that *p* if and only if the other does. Once again, the distinction between justification and warrant is crucial here. Both of those subjects would be at least *prima facie justified* in believing that *p*, though whether they know depends on whether the proposition that *p* is warranted for them, and it is perfectly possible for the proposition to be warranted for one of them and not warranted for the other.

A defender of McDowell could say, at this point, that there is still something akin to skepticism in the idea that, when I have perceptual knowledge that *p*, there are possible epistemic agents for whom the *phenomenological* aspects of their experience are just like mine, and yet they don't know that *p*. This is, indeed, true: according to the kind of HCF theories that I am defend-

²⁰ McDowell, "Criteria, Defeasibility, and Knowledge," p. 213 of the version mentioned in note 19.

ing, knowledge doesn't supervene on phenomenological states alone. Whether this is akin to skepticism is, I think, a debatable matter. In any case, McDowell can't accept this kind of help, for even his disjunctivism is a theory according to which the fact that *p* is making itself perceptually manifest to me and the fact that I am having a mere appearance as if *p* can be, as remarked before, phenomenologically indistinguishable for me:

The most obvious attraction [of HCF theories] is the phenomenological argument: the occurrence of deceptive cases experientially indistinguishable from non-deceptive cases. But this is easily accommodated by the essentially disjunctive conception of appearances that constituted the alternative. The alternative conception can allow what is given to experience in the two sorts of cases to be the same *in so far as* it is an appearance that things are thus and so...²¹

McDowell wants his disjunctivism to be compatible with the phenomenological indistinguishability of facts making themselves perceptually manifest and mere appearances. If so, he can hardly argue against HCF theories on the basis that they (also) entail that indistinguishability.

We have seen that HCF theories are not committed to the truth of (4); on the contrary, they have readily available a good explanation, in terms of the distinction between warrant and justification, of its falsehood. McDowell sometimes protests that an internalist notion of justification will issue in skepticism and thus make thought about the world seem impossible. But neither is internalism forced on HCF theories (even if they distinguish, as I have urged they should, between justification and warrant) nor is it clear that internalism is especially vulnerable to skeptical considerations.²²

²¹ McDowell, "Criteria, Defeasibility, and Knowledge," p. 389.

²² Many thanks to Manuel Comesaña, Michael Pace, Carolina Sartorio, Ernest Sosa, James van Cleve, and an anonymous referee for this journal for very helpful comments on previous drafts.