

## Primitive Knowledge Disjunctivism

Berit Brogaard

September 30, 2008

I argue that McDowell-style disjunctivism, as the position is often cashed out, goes wrong because it takes the good epistemic standing of veridical perception to be grounded in “manifest” facts which do not necessarily satisfy any epistemic constraints. A better form of disjunctivism explains the difference between good and bad cases in terms of epistemic constraints that the states satisfy. This view allows us to preserve McDowell’s thesis that good cases make facts manifest, as long as manifest facts must satisfy epistemic constraints.

### I. Disjunctivism

Consider the following scenario. Mike is looking at a blue ball right in front of him, and sees that the ball is blue and is right in front of him. Mike has an experience as of there being a blue ball right in front of him, there is indeed a blue ball right in front of him, and the blue ball is appropriately related to his experience of the blue ball. This is a good case of perception: Mike’s experience is veridical. Consider a second scenario. Unbeknownst to Mike, Mike is given a drug that makes him hallucinate a blue ball right in front of him (in an otherwise exactly similar environment). Mike has an experience as of there being a blue ball right in front of him, but there is no ball right in front of him. It merely seems to Mike as if there is a blue ball right in front of him. This is a bad case of perception: Mike’s experience is hallucinatory.

What I will refer to as ‘metaphysical disjunctivism about perceptual experience’ (or just ‘disjunctivism’ for short) is the view that despite the fact that they can be subjectively indistinguishable, good and bad cases are of different fundamental kinds. For example, despite the fact that Mike cannot necessarily tell whether he is in a good or a bad case, his veridical

perception in the first case and his hallucination in the second case are of different fundamental kinds. This characterization, of course, calls for refinement. In particular more needs to be said about *fundamental kind*. There are several ways to cash out the concept. One is in terms of metaphysical structure. On this view, good and bad cases have different metaphysical structures. Veridical perceptions are relations to external objects. Hallucinations are relations to object-independent contents or intentional objects, or they are belief-like states (see Fish 2008). Something like this view has been defended by, among others, J. M. Hinton (1973), Paul Snowdon (1980-81), John McDowell (1982, 1986), Hilary Putnam (1999), M. G. F. Martin (2002), Charles Travis (2004), Harold Langsam (1997), Ram Neta (2008), and William Fish (forthcoming).

One motivation for metaphysical disjunctivism can be articulated in the form of a reductio on the premise that good and bad cases are metaphysically alike. Call it *the argument from hallucination*:

#### Argument from Hallucination

- 1) Good and bad cases are metaphysically alike
- 2) Bad cases are not relations to external objects
- 3) So, good cases are not relations to external objects

Unlike sense-data theorists who have traditionally embraced the conclusion, disjunctivists treat the argument as a reductio on the first premise. In the good case, Mike is appropriately related to the blue ball. Not so in the bad case. In the bad case Mike's experience is a mental state with an object-independent content. So, good and bad cases are not metaphysically alike.

Disjunctivism contrasts with what we might call *highest common factor views*.<sup>1</sup> Highest common factor theorists agree with disjunctivists that good and bad cases can be subjectively

---

<sup>1</sup> The terminology is Hinton's and has been taken up by McDowell and others.

indistinguishable. But they take the subjective indistinguishability of good and bad cases to be explainable in terms of a common factor or shared component. Representationalists, for example, take the phenomenal character, or ‘subjective feel’, of experience to determine the ‘phenomenal content’ of experience. On this view, if a veridical perception and a hallucination are subjectively indistinguishable, then they have the same phenomenal character and hence the same phenomenal (representational) content. According to weak representationalists, a veridical perception and a hallucination which are subjectively indistinguishable may have additional representational content besides phenomenal content, and this additional content may differ in the two cases. But that part of the representational content which is determined by the phenomenal character of experience will be the same in those cases in which the two states are subjectively indistinguishable to a given subject. And insofar as experiences are individuated in terms of their phenomenal content, good and bad cases are then of the same fundamental kind.

We can tentatively formulate metaphysical disjunctivism and the highest common factor views as follows:

Classical Disjunctivism: Good and bad cases are of different fundamental kinds

Highest Common Factor Views: Good and bad cases are of the same fundamental kind.

Hallucinations are paradigmatic examples of bad cases. Illusions are more tricky. Illusions are cases where one’s experience attributes mistaken properties to the object of experience. They can be classified in three different ways: As good cases, as bad cases, or as sometimes good cases and sometimes bad cases. The three distinct classifications of illusions give rise to three kinds of disjunctivism:

1) Illusions are bad cases, like hallucinations (Hinton, McDowell, Martin)

2) Illusions are good cases, like veridical perceptions (Langsam, Snowdon)

3) Illusions are sometimes good cases and sometimes bad cases (Fish)

Treating illusions as bad cases, like hallucinations, is motivated insofar as illusions are bad in some respect. While in illusory cases the subject stands in a causal relation to the object of her experience, her experience attributes mistaken properties to the object. Treating illusions as good cases is motivated insofar as illusions might be thought to give rise to knowledge. Even if my experience attributes redness to a white table, I might still be in a position to know that I am looking at a table, though I won't be in a position to know that the table is red (as it is not). Finally the mixed position is motivated by the thought that some illusions attribute very few mistaken properties to the object of experience, whereas others attribute very many mistaken properties to the object of experience. So, illusions of the former kind should be treated as good, whereas illusions of the latter kind should be treated as bad.

Metaphysical disjunctivism is a claim about the nature of experience and is to be contrasted with a much less controversial form of disjunctivism, viz. epistemological disjunctivism. The broad rationale for epistemological disjunctivism is that skepticism ensues unless good and bad cases are treated as having different epistemic standings. We can articulate the rationale for disjunctivism in the form of a reductio:

Argument via Skepticism

1) Good cases and bad cases are epistemically on a par

2) In bad cases you don't get knowledge

3) So, in good cases you don't get knowledge

If one were to accept the conclusion that good cases of perception never give rise to knowledge, one would be required to accept skepticism about most empirical claims. Since it can be agreed

upon by nearly all parties both that skepticism is a questionable thesis and that bad perceptual cases do not give rise to knowledge, rejecting the first premise should be uncontroversial. Good and bad cases then are not epistemically on a par. Rejecting the first premise leads to epistemological disjunctivism—a somewhat trivial thesis, which—following Byrne and Logue (2008)—we may articulate as follows:

#### Epistemological Disjunctivism

The perceptual evidence one has in the good case is better than the perceptual evidence one has in the bad case

In this paper I will focus primarily on metaphysical disjunctivism, in particular McDowell's version of it (below I will defend the claim that McDowell is a metaphysical disjunctivist against recent arguments to the contrary). I will argue that McDowell's position, as it is typically cashed out, goes wrong because it takes the good epistemic standing of veridical perception to be grounded in "manifest" facts which do not necessarily satisfy any epistemic constraints. I will then argue that a better form of disjunctivism explains the difference between good and bad cases in terms of epistemic constraints that the states satisfy, and that this allows us to preserve McDowell's thesis that good cases make facts manifest, as long as manifest facts must satisfy epistemic constraints.

## II. The Internality Argument for Metaphysical Disjunctivism

Let us begin with what I consider one of McDowell's main motivations in favor of disjunctivism, viz. the claim that disjunctivism ensues even if it is conceded to the internalist that the epistemic standing of one's experiential state is grounded in matters internal to the state itself. The consideration is expressed in the following excerpt from McDowell's influential piece "Criteria, Defeasibility and Knowledge":

The root idea [underlying arguments for highest-common-factor views] is that one's epistemic standing on some question cannot intelligibly be constituted, even in part, by matters blankly external to how it is with one subjectively. For how could matters beyond one's ken make any difference to one's epistemic standing? ... But the disjunctive conception of appearances shows a way to detach this 'internalist' intuition from the requirement of non-question-begging demonstration. 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, as it would be if the truth about that were exhausted by the highest common factor (1982: 26-27).

What is it for a fact to be made manifest to a subject? To a first approximation, a fact *p* is made manifest in an experience *E* just when *E* is a mental state that has *p* as one of its relata. Or as McDowell puts it, '[when a fact is made manifest] the fact itself is [the] object [of the experience]; so that its obtaining is not, after all, blankly external' (1982: 25).

In the passage just cited McDowell responds to the common internalist intuition that the good epistemic standing of good perceptual cases cannot be grounded in something blankly external to the agent's subjectivity. McDowell grants this intuition but then replies that (i) external facts *can* contribute to the experiencer's subjectivity by being constituents (or relata) of a mental state, and (ii) if the difference in epistemic standing between good and bad cases were not grounded in a difference between the states, then it would be grounded in something purely external to the states themselves. In arguing against his internalist opponent McDowell appeals to the intuitive pull of a principle which he takes to be endorsed also by his opponent. The principle suggested by McDowell, as I understand him, is this one:

#### Internality

The good epistemic standing of the good cases is not purely external to the state itself

Given this principle McDowell's internality consideration in favor of disjunctivism can be articulated as follows:

#### Internality Argument

- (1) If the difference in the epistemic standing between good and bad cases is not grounded in a difference between the states themselves, then what grounds the difference in epistemic standing between good and bad cases is purely external to the states themselves.
- (2) The difference in epistemic standing between good and bad cases is not purely external to the states themselves.
- (3) So, the difference in epistemic standing between good and bad cases is grounded in a difference between the states themselves.

Being a simple modus tollens this argument in support of disjunctivism is obviously valid. The question is whether the argument is at least prima facie sound and whether its conclusion supports disjunctivism. Let us look at each of the argument's premises in turn. The first premise should be uncontroversial. The difference in epistemic standing between good and bad cases is grounded either in something internal to the states or something external to the states. If it's not grounded in a difference internal to the states, then it's grounded in a difference external to the states. Despite McDowell's assumption that his opponent will endorse Internality (viz. the second premise), the second premise is not without its problems. Here is one motivation in its favor. Suppose that there could be a veridical perception and a hallucination which are subjectively indistinguishable and yet have different proximal causes. In that case one could argue that the difference in epistemic standing between the cases is at least partially grounded in a difference between the states themselves. So the difference in epistemic standing between good

and bad cases is not purely external to the state itself. But there is a simple reply to this line of argument. One could treat the proximal causes as external to the states themselves, in which case the good epistemic standing of good perceptual cases *would* be grounded in something external to the states themselves. Here is another sort of pressure towards granting the truth of the second premise. One could perhaps argue that if the difference in epistemic standing is grounded in something purely external to the states themselves, then the difference in epistemic standing is a matter of pure luck: it has no bearing on 'how [matters are] with one subjectively'. There is, however, no point in becoming embroiled in speculative argument as to what, exactly, McDowell thinks motivates the second premise. It will suffice for our purposes simply to take its truth for granted.

The question still remains, however, whether the internality argument supports disjunctivism. The argument does lend support to McDowell-style disjunctivism. For, ex hypothesi, the good epistemic standing of veridical perceptions is not grounded in something purely external to the state itself. A natural answer to what the good epistemic standing of veridical perceptions is grounded in is that it is grounded in the subject's relation to the object (or fact) of the experience. But if this is so, then it would seem that in the case of veridical perception the subject must be appropriately related to an object. In the case of a hallucination, on the other hand, the subject is not appropriately related to an object. So, there is then a fundamental difference between veridical perceptions and hallucinations. In the case of veridical perceptions, there is an appropriate relation to an object, in the case of hallucinations there is not. The difference in epistemic standing between the two kinds of states is fully explained by the obtaining of an appropriate relation between the subject and an object in the one case but not in the other. But if this is so, then it is hard to see what role the alleged highest common factor plays. The alleged highest common factor is not what grounds the difference in the epistemic standing of veridical perceptions and hallucinations. And the highest common factor is supposed to be distinct from the phenomenal character or 'subjective feel' of the experience. So, it is

tempting to conclude that veridical perceptions and hallucinations have no highest common factor, that they have nothing of relevance in common besides being subjectively indistinguishable.

The internality argument thus supports disjunctivism. But the argument's conclusion does not entail disjunctivism. For, there could be a difference between the states but also a common fundamental kind. E.g. perhaps veridical perceptions and hallucinations can share phenomenology but have different contents. The internality argument at least seems to make a good case for the thesis that veridical perceptions and hallucinations must differ in how the subject is related to the object of experience. Somehow veridical perceptions must be object-dependent in a way that hallucinations are not. However, this view, that veridical perceptions must be object-dependent in a way that hallucinations are not, could be combined with a highest common factor thesis. For example, one could hold that veridical perceptions and the corresponding hallucinations share phenomenal content (as predicted by weak representationalism) but that perceptual experiences have a further layer of content that is object-dependent in the case of veridical perceptions but gappy or non-object-dependent in the case of hallucinations (see e.g. Chalmers 2004). The good epistemic standing of veridical perceptions could then be taken to be grounded in the object-dependency of this further layer of content. So, the argument's conclusion does not entail disjunctivism, it merely keeps disjunctivism in the running along with multiple-layers-of-content views (and perhaps other highest common factor views). The argument's conclusion seems to rule out views of perception which make it possible for veridical perceptions and hallucinations to have exactly the same content. For if it is possible for a veridical perception and a hallucination to have exactly the same content, then it would seem that they should have the same epistemic standing. Below we will argue that the argument's conclusion does not in fact rule out such views. But before turning to these considerations let us briefly consider a recent claim to the effect that McDowell is not in fact a metaphysical disjunctivist but rather only an epistemological disjunctivist.

### III. McDowell as an Epistemological Disjunctivist?

Alex Byrne and Heather Logue (2008) classify McDowell as an epistemological (but not metaphysical) disjunctivist. But there are slim grounds indeed for this claim. In support of their claim they offer the following quotes from McDowell:

1982: it 'looks to one exactly as if things [are] a certain way' in the good and the illusory case.

2008: The central thesis of 'the disjunctive conception' is that 'the two sides of the disjunction differ in epistemic significance ... this difference in epistemic significance is of course consistent with all sorts of commonalities between the disjuncts. For instance, on both sides of the disjunction it appears to one that, say, there is a red cube in front of one'.

However, the implications of the two quotes are not far-reaching and hardly support the claim that McDowell is not a metaphysical disjunctivist. The problem with the Byrne/Logue line is this: it is perfectly compatible with metaphysical disjunctivism that things look or appear the same way to the subject in good and bad cases. The metaphysical disjunctivist could, for example, explain away common appearances in good and bad cases the same way that they explain away the subjective indistinguishability of good and bad cases (see Fish 2008).

Alternatively, the disjunctivist could argue that the commonality between the good and the bad cases disappears once a proper distinction has been drawn between mere appearances and facts made manifest -- move made by McDowell:

The 'highest common factor' conception has attractions for us that cannot be undone just by describing an alternative .... The most obvious attraction 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 constitutes the alternative. The alternative conception can allow what is given to experience in the two sorts of case to be the same *in so far as* it is an appearance that things are thus and so; that leaves it open that whereas in one kind of case what is given to experience is a mere appearance, in the other it is the fact itself made manifest. So the phenomenological argument is inconclusive (1982: 25).

Why do Byrne and Logue think the passages they cite support their claim that McDowell is an epistemological but not a metaphysical disjunctivist? Presumably because they think that there is common mental element between good and illusory cases and that this common mental element is a certain kind of appearance. They consider two cases in which Austin has an experience as of there being a red spherical tomato in front of him. The first case is a good case, the second an illusion. In both cases Austin is seeing the tomato. But seeing the tomato, they say, cannot be the common mental element between the two cases; it's not specific enough, as Austin would also be seeing a tomato if he had an appearance as of there being a green ovoid tomato in front of him. The common mental element, they say, is *the tomato looks red and spherical to Austin*. This, they claim, is a fundamental kind, which the good and the illusory case have in common. But if there is a fundamental kind, which good and illusory cases have in common, then metaphysical disjunctivism which treats illusions as bad cases is false. Or so the argument goes.

However, a defect of the argument is that the disjunctivist can agree that in both the good and the illusory case the tomato looks red and spherical to Austin. The disjunctivist can simply insist that 'looks'-reports of this kind report a commonality between the two cases but then add that upon further analysis, the commonality turns out not to be fundamental: in the first case the tomato looks red and spherical because a fact is made manifest, in the second case the tomato *merely* looks red and spherical – as we have already seen, a point urged vigorously by McDowell:

An appearance that such-and-such is the case can be *either* a mere appearance *or* the fact that such-and-such is the case making itself perceptually manifest to someone ... the object of experience in the deceptive case is a mere appearance. But we are not to accept that in the non-deceptive cases too the object of experience is a mere appearance, and hence something that falls short of the fact itself. On the contrary, the appearance that is presented to one in those cases is a matter of the fact itself being disclosed to the experiencer (1982: 23)

On McDowell's view, the tomato's looking red and spherical to Austin is not a common element between the good and bad case. In the illusory case the tomato *merely* looks red and spherical to Austin. In the bad case the tomato looks red and spherical to Austin because a red-and-spherical-tomato fact is made manifest in Austin's experience.

In fact, it is clear that McDowell does not take the fact-made-manifest thesis to be a purely epistemological position. He thinks that the fact-made-manifest thesis entails epistemological disjunctivism but conceives of it as distinct from metaphysical disjunctivism. Here is McDowell:

One can hardly countenance the idea of having a fact made manifest without supposing that that would make knowledge of the fact available to one (1982: 26)

The fact-made-manifest thesis entails epistemological disjunctivism but the two theses are not equivalent. One can 'hardly countenance' metaphysical disjunctivism without epistemological disjunctivism but the two positions are distinct. McDowell's version of disjunctivism is most clearly a kind of metaphysical disjunctivism, and as it is commonly understood the position exemplifies a version of disjunctivism which treats illusions as bad cases.

#### IV. An Argument against V v I/H

As McDowell-style disjunctivism is most commonly cashed out, the position is committed to the claim that all veridical (non-illusory, non-hallucinatory) perceptions have a better epistemic

standing than the corresponding hallucinations (see e.g. Byrne and Logue 2008 and Fish forthcoming). It is further thought that this claim motivates a treatment of good and bad cases as fundamentally different. To take this to be the case, however, would be too hasty. There is evidence that McDowell does not in fact hold this view. In a footnote in “Criteria, Defeasibility and Knowledge” he says:

An experience subjectively indistinguishable from that of being confronted with a tomato, even if it results from confrontation with a tomato, need not count as experiencing the presence of a tomato. Another case in which it may not count as that is one in which there are a lot of tomato facades about, indistinguishable from tomatoes when viewed from the front (1982: 26, fn 39)

The sentiment McDowell here expresses (albeit in passing) is that not *every* veridical (non-illusory, non-hallucinatory) perception of an object counts as seeing the object and hence that not every veridical perception counts as a good case. Presumably McDowell would say that in such cases veridical perception does not have a better epistemic standing than the corresponding bad cases. Nonetheless, McDowell-style disjunctivism is typically construed as drawing a sharp distinction between veridical perception and illusions/hallucinations. Following Byrne and Logue (2008), let us call this common construal of McDowell-style disjunctivism ‘V v I/H disjunctivism’. It will be the burden of the rest of this section to argue that V v I/H disjunctivism, and its closely related cousin V/I v H disjunctivism, are epistemically unmotivated.

The objection to V v I/H disjunctivism can be articulated as follows:

#### V v I/H Objection

- 1) Some veridical perceptions have no better epistemic standing than the corresponding hallucinations/illusions.

2) If 1), then a difference in kind between veridical perception and hallucination/illusion does not adequately account for their relative epistemic standings.

3) Hence, a difference in kind between veridical perception and hallucination/illusion does not adequately account for their relative epistemic standings.

The argument is obviously valid. Moreover, given V v I/H disjunctivism, the second premise is very compelling. Part of McDowell's motivation for disjunctivism is that when a fact is made manifest in experience, then the subject is in a position to know, whereas when a fact is not made manifest in experience, the subject is not in a position to know. As McDowell's position is most commonly cashed out, a fact is made manifest in experience in all cases of veridical perception. For example, if I stand in the right sort of relation to a red chair, then a red-chair-fact is made manifest in my experience. So, if there are cases in which a veridical perception has no better epistemic standing than the corresponding hallucination/illusion, then it is plausible that the fact-made-manifest thesis cannot account for the difference in epistemic standing between veridical perceptions and hallucinations/illusions. So, the second premise is compelling. The controversial premise is the first. I will offer two related arguments in favor of the first premise.

You are looking at a red wooden chair. Unbeknownst to you, someone slowly removes pieces of wood from the chair while simultaneously feeding you drugs which "reconstruct" the missing parts in your visual field. So, just as someone removes a piece of wood from the back of the chair, a drug is simultaneously administered, and it thus seems to you as if the chair is still intact. Of course, as we imagine the case, you are not aware of any of these changes. You are unaware that someone is taking apart the chair in front of you, and you are unaware that drugs are administered that "reconstruct" the missing parts of the chair in your visual field. Now, at the beginning of this scenario your experience is a veridical (non-hallucinatory and non-illusory) perception, it then becomes an illusion, and finally when all the pieces of the chair have been removed, your experience is a (falsidical) hallucination. On V v I/H disjunctivism, as illusive

experiences (normally) are falsidical, no fact is made manifest in such cases. But consider now an arbitrary instant just before the chair is dismantled and drugs are administered to you to “reconstruct” the missing parts. Since you believe that there is a red wooden intact chair in front of you throughout the dismantling process, you believe that there is a red wooden chair in front of you at this instant. Moreover, your belief is true. There is a red wooden chair in front of you. Your experience is a (non-hallucinatory and non-illusory) veridical perception. Nonetheless, your belief that there is a red wooden intact chair in front of you is not safe. For in some of the closest situations in which you have the same belief, there is not a red wooden intact chair in front of you. Since your belief is not safe, it does not count as knowledge. In fact, you are not in a position to know that there is a red wooden intact chair in front of you. We thus have an example of veridical perception which does not put you in a position to know and which therefore does not differ in epistemic standing from the corresponding illusion.

Note that it's not because you cannot subjectively discern the difference between your veridical experience and your illusion that you are not in a position to know in the good case. We can grant that knowledge is a purely external state. It still follows that your belief does not (and could not) amount to knowledge because it fails to be safe.

Of course, the disjunctivist could argue that safety is not required for knowledge. But this move would be idle. For, surely knowledge states must satisfy some anti-luck constraint. But no such constraint is satisfied in the envisaged case. If safety is not the right sort of anti-luck constraint, perhaps sensitivity is. *s*'s belief that *p* is sensitive just in case in all of the closest situations in which *p* is false, *s* does not believe that *p*. But your belief that there is a red wooden intact chair in front of you is not sensitive. Some of the closest situations in which there is no red wooden intact chair in front of you are the future situations in which, unbeknownst to you, small pieces of wood have been removed from the chair. But in these situations you believe that there is a red wooden intact chair in front of you. So, in some of the closest situations in which it is not

the case that there is a red wooden intact chair in front of you, you believe that there is one.

Hence, your belief is not sensitive.

It may perhaps be said that your veridical perception and your illusions have different epistemic standings because, even though your veridical perception does not put you in a position to know that there is a red wooden intact chair in front of you, it gives you justification for believing that there is one. But this cannot be right. For, there is no difference in your internal justification for your beliefs that there is a red wooden intact chair in front of you throughout the dismantling process. And if justification is construed as safety (a move made by Williamson 2000), then there is no difference in your justification before and after the onset of the dismantling process. So, there is no difference in epistemic standing between your veridical perception before the onset of the dismantling process and your illusions after the chair is partly dismantled.

In some discussions of disjunctivism illusions are treated as good cases, like veridical perceptions (see e.g. Langsam and Snowdon). Following Byrne and Logue (2008) call this 'V/I v H disjunctivism'. The main motivation for this treatment of illusions is that in the case of illusions, the subject is appropriately related to an object. Moreover, just like veridical perceptions illusions could reasonably be thought to give rise to knowledge. For example, if I am having an experience that attributes redness to a white chair I cannot come to know that there is a red chair in front of me but I can come to know that there is an object shaped in a certain way in front of me. We can still maintain a version of the fact-made-manifest thesis in this case. If I have an experience that attributes redness to a white chair, there is a fact made manifest to me, for example, the fact that there is an object shaped like a chair in front of me.

The chair example just considered undermines the main motivation for V v I/H disjunctivism. But V/I v H disjunctivism too falls prey to a version of the chair problem. Consider one of the last instants at which you are having an illusion as of there being a red wooden intact chair in front of you. At one of the very next instants your experience will be a

hallucination. On the view under consideration, illusions are supposed to put us in a position to know, whereas hallucinations are not. But in the envisaged scenario this is hardly so. Firstly, at one of the last instants during the dismantling process the cause of what you appear to be seeing is the drug. The drug makes the wooden stub that remains look like a much larger, differently shaped and differently colored object (viz, the chair). Secondly, in the closest situations in which you believe that there is an object in front of you, your belief is false. So, even your belief that there is an object in front of you fails to be safe. We thus have an example of an illusion which does not put you in a position to know anything and which therefore does not differ in epistemic standing from the corresponding hallucination.

The chair example admittedly does nothing to undermine versions of disjunctivism that classify illusions sometimes as good cases and sometimes as bad cases (see Fish forthcoming). Disjunctivists who hold this view could simply say that when only a few pieces of wood have been removed from the chair the perceptual case is still a good one, and this explains why the veridical hallucination and the illusions have the same epistemic standing. When, on the other hand, all but a stub of the chair has been removed the perceptual case is to be treated as a bad one, and this explains why the illusion and the hallucination have the same epistemic standing. However, there are other examples which provide a firmer basis on which to regard this alternative form of disjunctivism as inadequate. As we will see below, in some cases of veridical hallucination the subject stands in an appropriate relation to an object, and so some cases of veridical hallucination are misclassified even by mixed-case disjunctivism.

We can offer a further example in support of the first premise of the V v I/H objection. Suppose you are an invert in a normal perceptual community.<sup>2</sup> For you red things yield illusory green experiences, and vice versa. But one lucky time a red thing yields a red experience. In this case your experience is a veridical perception. But it does not put you in a position to know that

---

<sup>2</sup> The example was originally offered to me by David Chalmers as an objection to a version of the primitive knowledge thesis which I defend below.

the thing in front of you is physically red. For, you grew up in a normal perceptual community. So, typically when you have a phenomenally red experience you come to believe that you are looking at a physically green object. Since you do not know that your current phenomenally red experience is veridical, you come to believe that you are looking at a physically green object. But this time your belief is false. The object you are looking at is red, not green. Since your belief is false, it does not (and could not) count as knowledge. So, unlike your normal illusory red experiences, your current phenomenally red experience does not put you in a position to know that there is a red object in front of you. So we have an example of a veridical perception which has a worse epistemic standing than the corresponding illusion.

These two cases support the first premise of the V v I/H objection. And no doubt there are endless cases of this sort. But it follows via the V v I/H argument that the difference in kind between veridical perception and hallucination/illusion does not adequately explain their relative epistemic standing. If facts are made manifest in veridical perceptions, but not in illusions and hallucinations, then the fact-made-manifest thesis cannot explain the differences in epistemic standing. So, metaphysical disjunctivism is epistemically unmotivated.

The moral then is this. Some veridical perceptions are bad cases. One can hold onto the fact made manifest thesis only by denying that all veridical perceptions make facts manifest: Only those veridical perceptions which make facts manifest are good cases.

## V. Veridical Hallucination

I will now make a case for a new form of disjunctivism which I will call 'primitive knowledge disjunctivism'. On this view, good cases are primitive knowledge states, whereas bad ones are not.

I am going to remain neutral on the question of whether good perceptions have content or make facts manifest (in McDowell's sense). But let's start with content (for simplicity: structured content). Suppose Mike is looking at a blue ball right in front of him and sees that the ball is blue

and is right in front of him. What is the content or object of Mike's perceptual state? If perception is a mental state with a Russellian content which consists of physical objects and/or physical properties, then the content of Mike's perceptual state consists of the blue ball *o*, the surface spectral reflectance property blue and the property of being right in front of Mike. As this content is true in the envisaged circumstances, Mike's experience is veridical.

Consider a second scenario. In this scenario Mike undergoes a hallucination. There is no ball right in front of him. As there is no ball in front of him, the blue ball *o* does not enter into the content of his experience. The content of Mike's experience is gappy (see e.g. Braun 2005). It consists of a gap and the properties of being blue and being right in front of Mike. As gappy propositions are false, Mike's experience is falsidical, which is as it should be.

Consider a third scenario. As in the first scenario, Mike is looking at a blue ball right in front of him and sees that the ball is blue and is right in front of him. But after a few minutes Mike is given a palinopsia-inducing hallucinogenic drug. The drug causes his experience to persist after the corresponding stimuli have left and prevents his visual system from processing new visual information. As it happens, however, the blue ball is moved out of Mike's view just as the drug takes effect.

In this scenario, it is plausible that the content of Mike's experience continues to be a conglomeration of the blue ball *o*, the surface spectral reflectance property blue and the property of being right in front of Mike even after the drug takes effect. After all, advocates of the thesis that mental states have Russellian content will be happy to grant that one can have a belief directly about an object even when the object is not present. So, given a Russellian view of content, there shouldn't be any principled reason for denying that the blue ball can be a constituent of Mike's experience even after the drug takes effect. Of course, in order for one to have a mental state (e.g. a belief) directly about an object, one must be in some sort of causal contact with the object. But Mike is in causal contact with the blue ball even once the ball has been removed from the scene. It is thus plausible that the content of Mike's perceptual

experience consists of the blue ball *o*, the surface spectral reflectance property blue and the property of being right in front of Mike. As there is no ball within Mike's view, however, the content of Mike's experience is false. So, Mike's experience is falsidical. Again this is as it should be.

But consider now a fourth scenario. As in the third scenario Mike is given a palinopsia-inducing hallucinogenic drug which causes his experience to persist and which prevents his visual system from processing new visual information. But in this case the ball is left untouched. So if the content of Mike's experience consists of the blue ball *o*, the surface spectral reflectance property blue and the property of being right in front of Mike, which we can reasonably assume, then the content of Mike's experience is true. *o* does indeed instantiate the surface spectral reflectance property blue and is right in front of Mike. So, Mike is having a veridical hallucination.

Now, very few would be happy to grant that the content of perceptual experience is exhausted by its Russellian content. There are two considerations against this view. First, there is the problem of the inverted spectrum (see e.g. Shoemaker 1994, Chalmers 2004, 2006, Egan 2006). It seems possible for there to be a perceptual community in which red objects systematically give rise to green phenomenal experiences. When people in the inverted community perceive red objects, their experiences are phenomenally indistinguishable from the experiences we have when we are looking at green objects. Yet if the content of perception contains only physical properties (e.g. surface-spectral reflectance properties), then people who perceive red objects will have perceptions the content of which contains the surface-spectral reflectance property red, and people who perceive green objects will have perceptions the content of which contains the surface-spectral reflectance property green. So, if content determines phenomenal character, then inverted spectrum scenarios shouldn't be possible.

Second, it seems that Mike could have had a phenomenally indistinguishable experience if he had been looking at a perfect replica of the blue ball.<sup>3</sup> Yet on a Russellian conception of content, material objects enter into the content. So, Mike's experience of the blue ball and his experience of a perfect replica of the blue ball have different contents. Mike's experience of the blue ball contains the blue ball as one of its constituents, whereas Mike's experience of the perfect replica contains the perfect replica as one of its constituents. So, if the content of perceptual experience is exhausted by its Russellian content, and phenomenal content supervenes on phenomenal character, then we should expect Mike's experience of the blue ball and his experience of the perfect replica to be phenomenally distinguishable.

Of course, one can avoid the first problem if one holds that none of the content of perception is determined by the phenomenal properties of perception, and one can avoid the second problem if one holds that none of the content of perception supervenes on the phenomenal properties of perception. If, however, one takes *phenomenal* content to supervene on phenomenal properties and phenomenal properties to supervene on phenomenal content, as many thinkers do,<sup>4</sup> one cannot easily hold that all of the content of perception is Russellian.

To avoid these problems David Chalmers (2004, 2006) suggests that the content of perceptual experience has a Russellian and a Fregean component.<sup>5</sup> The Fregean component is the phenomenal content of the experience and is the same regardless of what the environment is like. The Fregean content yields a Russellian content in a particular environment. Roughly, the Fregean content of Mike's experience as of the blue ball of being right in front of him consists of the property of being the object that is causing the current experience, the property of being the property that normally causes phenomenally blue experiences, and the property of being right in

---

<sup>3</sup> Thompson (2006), Schroeder and Caplan (2007) and Tye (2007) discuss versions of this problem for strong representationalism.

<sup>4</sup> For discussion see Chalmers (2004).

<sup>5</sup> In his (2006) Chalmers argues that perceptual experiences also have edenic content, which consists of primitive non-physical properties. The edenic content of perceptual experience is (imperfectly) veridical just in case it matches the Russellian content, which it does if it is the sort of content normally caused by the properties and objects constituting the Russellian content.

front of the person having the current experience. If the blue ball *o* is right in front of Mike, the Fregean content yields a Russellian content which consists of the blue ball *o*, the surface spectral reflectance property blue and the property of being right in front of Mike. In normal cases of hallucination, no object is causing the experience. So, the Fregean content does not yield a complete Russellian content, and so the hallucination will fail to be veridical.

However, even the Fregean account of perceptual content leaves open the possibility of veridical hallucination. Suppose again that Mike is given a palinopsia-inducing hallucinogenic drug which causes his blue ball experience to persist and which prevents the visual system from processing new visual information. We might imagine that the Fregean content of Mike's experience remains the same for at least a few minutes following the administering of the drug. But Mike's experience was caused by the blue ball *o*, so arguably the Fregean property of being the object that caused Mike's experience yields *o* in our envisaged circumstances. But Mike's experience is then a veridical hallucination.

It may perhaps be argued that the blue ball is not an appropriate cause of Mike's experience, because there is no longer a causal chain that leads from the ball to Mike's experience. But as causation need not be instantaneous, it cannot be the time lag between the ball's reflectance of light and the experience taking place that stands in the way of causation. In fact, we can even stand in an appropriate causal connection to objects that are no longer present. For example, it is quite plausible that I am not hallucinating when I am star gazing. But if we can stand in appropriate causal relations to objects that are no longer present, then it is plausible that Mike can also stand in an appropriate causal relation to the blue ball in the envisaged circumstances. Of course, in Mike's case Mike's perceptual system mistakenly allows his past experience of the blue ball to persist when it ought to have allowed for the receipt of new visual information. But this just goes to show that Mike's experience is defective in some way, not that it fails to be properly caused by the object the experience is about.

One could attempt to avoid these problems by cashing out causation (proper) in terms of causal dependence. A causally depends on B iff if B hadn't occurred, A wouldn't have occurred. Let  $t_1$  be a time right before the hallucinatory drug takes effect, and let  $t_2$  be a time right after the drug takes effect. One could then say that Mike's experience at  $t_2$  causally depends on the blue ball being right in front of him at  $t_1$ , but that it does not causally depend on the blue ball's being right in front of him at times close to  $t_2$ . One could then explain the defectiveness of Mike's experience at  $t_2$  by saying that it fails to causally depend on the blue ball at times close to  $t_2$ . However, it is not difficult to conjure up cases in which a defective experience at a time  $t$  causally depends on the object the experience is about at times close to  $t$ . Suppose Alice is told to give Mike a drug that instantly neutralizes the effects of the hallucinatory drug just if the blue ball  $o$  is no longer right in front of Mike. Then in the closest worlds in which the blue ball  $o$  is no longer right in front of Mike, Mike does not have an experience as of  $o$  being right in front of him. So, Mike's experience at  $t_2$  causally depends on the blue ball at times close to  $t_2$ . Still, Mike's actual hallucination is clearly defective. Of course, in the envisaged scenario, it is tempting to say that the blue ball is a deviant cause of Mike's experience due to Alice's potential intervention. However, this objection is amiss. For, if one spells out causation (proper) in terms of causal dependence, then any time there is causal dependence, there is causation (proper).

Faced with these problems, one could admit that the Russellian and Fregean accounts give rise to the possibility of veridical hallucination but insist that both accounts are wrong on the grounds that the content of perceptual experience is not a structured proposition. One might, for example, insist that the content of perceptual experience is possible-worlds content. However, this will not impede the possibility of veridical hallucination. On the possible-worlds view, the content of Mike's experience is a set of possible worlds in which the blue ball,  $o$ , instantiates the surface spectral reflectance property blue and is right in front of Mike. This content is true just in case the actual world is a member of this set. Since the actual world is a world in which the blue ball,  $o$ , instantiates the surface spectral reflectance property blue and is right in front of Mike, the

actual world is a member of this set. Hence, Mike's experience is veridical. So, rejecting the hypothesis that the content of perception is a structured proposition does nothing to close off the possibility of veridical hallucination.

One might see the trouble the Mike case causes for highest-common-factor views as a reason to give up on such views and instead follow standard disjunctivists in taking cases of seeing to be relations to external objects or facts and hallucinations to be relations to object-independent contents. But this move is unlikely to help. For, standard versions of disjunctivism are perfectly consistent with the possibility of veridical hallucination. In fact, it is plausible that Mike stands in an appropriate acquaintance or causal relation to an external object or fact in both the good and the bad case. So, standard versions of disjunctivism then predict that Mike's veridical hallucination is a good case of perception. Hardly a satisfactory result.

But if veridical hallucination is a genuine possibility, which it is on standard views of perception, then veridicality cannot serve as a mark of good perceptual experience, and falsidicality cannot serve as a mark of bad perceptual experience. The question then arises what serves as the mark of good perceptual experience. As we have just seen, it won't do to say that in the case of good perceptual experience, the perceiver is non-deviantly causally connected with the object. For, it is quite plausible that Mike stands in an appropriate causal relation to the blue ball. Certainly, the drug does not *produce* the content of Mike's experience. The ball did. The upshot: *good cases of perception* and *veridical perception* are not necessarily co-extensive.

We can articulate the argument from veridical hallucination as follows:

#### Argument from Veridical Hallucination

- 1) Veridical hallucination is possible
- 2) Veridical hallucination cases are not good cases (epistemologically).
- 3) *Good cases of perception* and *veridical perception* are not necessarily co-extensive.

My solution to the problem of how to characterize good cases of perception is to say that good cases of perception are primitive knowledge states. Only some veridical experiences are primitive knowledge states. Only primitive knowledge states make facts manifest. In the next section we will spell out this proposal in further detail.

## VI. Primitive knowledge Disjunctivism

It is fair to say that there is no single analysis of knowledge that seems to be entirely free of counterexample. Nonetheless it is widely agreed that in order for someone to know (in the standard sense) his or her belief must be safe and formed via a reliable belief-forming method (see e.g. Sosa 1999).  $s$ 's belief that  $p$  is safe, roughly, if in the closest worlds in which  $s$  believes that  $p$ ,  $p$  is true.  $s$ 's belief that  $p$  is reliably formed, roughly, if the belief was formed via a method which tend to give rise to safe belief in the sort of environment in which the belief was formed. Some take reliability to just be safety (see e.g. Williamson 2000). But the two notions come apart. For example,  $s$  might form a belief to the effect that it's 20 degrees outside by looking at a thermometer. Looking at a thermometer when determining the temperature tends to give rise to safe beliefs in normal environments. So, if  $s$  forms a belief that it's 20 degrees outside by looking at a thermometer in a normal environment,  $s$ 's belief is reliably formed. If, however, the thermometer is defective, then  $s$ 's belief fails to be safe. It is not the case that the belief that it's 20 degrees outside is true in all neighboring worlds in which  $s$  believes that it's 20 degrees outside.

It is also possible to have a safe belief that  $p$  which is not reliably formed. The simplest case of this is one in which  $p$  is necessarily true. Suppose, for example, that  $s$  guesses correctly that  $17 + 89 = 106$ .  $s$ 's belief that  $p$  then fails to be reliably formed. Guessing does not tend to produce safe beliefs. Yet, as  $17 + 89$  is necessarily 106, it is true that  $17 + 89 = 106$  in all the closest worlds in which  $s$  believes that  $17 + 89 = 106$ . So,  $s$ 's belief is safe.

Safety and reliability are standardly formulated in terms of belief. The reason for this is that knowledge traditionally is thought to be a belief-like state. But it is not essential to the formulation of the safety and reliability constraints that they be formulated in terms of belief. Let it be granted, for argument's sake, that some knowledge states are perceptual states. We might then formulate safety and reliability as follows.

#### Safety

*s*'s perceptual state to the effect that *p* is safe just in case in all the closest worlds in which *s* has a perceptual experience as of *p*, *p*.

#### Reliability

*s*'s perceptual state is reliably formed just in case it is formed via a method that tends to give rise to safe perceptual states in the sort of environment in which the state is formed.

Safety and reliability are most likely not the only constraints on knowledge. But for any externalist constraint which turns out to be a constraint on standard knowledge, it is possible to formulate an analogous constraint in terms of perceptual states (the only exception being the requirement that knowledge states *are* belief states that satisfy certain constraints). But it seems reasonable, then, to take perceptual states that satisfy most to all of the constraints a state must satisfy to count as knowledge to be knowledge states.

Now, some will take knowledge to involve internalist constraints. For instance, one could say that *s* must have internal evidence for *p* in order for her to know that *p*. However, there need not be any conflict between conceptions of knowledge that take knowledge to require internal evidence and conceptions of knowledge that do not require this. One could, for example, follow Ernest Sosa (2007) in distinguishing between what he calls 'animal knowledge' and 'reflective knowledge'. For Sosa, animal knowledge is purely externalist and so will satisfy

constraints like safety and reliability. Reflective knowledge, on the other hand, is partially internalist in requiring that the knower have internal evidence, and so reflective knowledge can be had only by creatures who have the capacity for thought and who can engage in reflection. If good perceptual states are knowledge states, then they will be akin to animal knowledge states, but unlike Sosa's animal knowledge states, one could in principle be in them without being in a corresponding belief state. Let us call these sorts of perceptual states 'primitive knowledge states'.

The upshot then is this. Perceptual states can have a status analogous to knowledge, although they need not require belief.<sup>6</sup> For example, in so far as safety and reliability are conditions for belief to be knowledge, safety and reliable are also conditions for perceptions to be primitive knowledge.

Having outlined an account of perceptual states as primitive knowledge states, let us now return to the case of veridical hallucination. Prior to Mike's ingestion of the hallucinogenic drug which makes his blue ball experience persist and which prevents his visual system from processing new visual information, Mike's perceptual state satisfies reliability and safety. So, Mike's perceptual state is a candidate to be primitive knowledge. After Mike's intake of the hallucinogenic drug, on the other hand, Mike's perceptual state is no longer safe, or at least it need not be. Mike could easily have turned around or walked away when he was given the drug. So, in some of the closest worlds in which Mike has an experience as of the blue ball o being right in front of him, Mike turns around when he is given the drug, and so his experience is falsidical; there is no ball *right in front of him*. But if Mike's perceptual state fails to be safe, then it doesn't count as knowledge. Thus, primitive knowledge must satisfy constraints very similar to the constraints satisfied by standard knowledge. Good perceptual states are states that satisfy all of these constraints. Bad perceptual states are states that fail to satisfy some of the constraints.

---

<sup>6</sup> However, one could require that whenever *s* is in a good case, *s* is also in a corresponding primitive belief state.

A problem here arises. It may be that Mike could not easily have turned around or walked away when he was given the drug. Suppose Mike is tied down or that the drug makes it impossible for him to move or that he agreed to stand still. Then in all of the closest worlds in which Mike has an experience as of the blue ball *o* being right in front of him, *o* is blue and is right in front of him. But then Mike's hallucination is veridical in all the closest worlds in which he is in a perceptual state to the effect that *o* is blue and is right in front of him. So, Mike's hallucinatory perceptual state is not only veridical and reliably formed, it is also safe.

However, this alleged counterexample does not undermine our conjecture that good cases are primitive knowledge states. It only shows that safety and reliability are not necessary and sufficient for knowledge. In fact, the problems with taking safety and reliability to be necessary and sufficient for primitive knowledge are strikingly similar to the problems we encounter if we make safety and reliability necessary and sufficient constraints on standard knowledge.

Consider a variation on Alvin Goldman's barn case. Henry is driving in the country and stops in front of a barn. Unbeknownst to him, he is looking at one of few real barns in an area spawned with facsimiles. The facsimiles are so realistic that if he had stopped in front of any of them, he would have been tricked into thinking that he were looking at a real barn. The standard intuition is that Henry does not know that he is looking at a barn, because he could easily have had the same belief while looking at a facsimile. In the variation a guardian angel would blur Henry's vision if he were to look at a fake barn. So, he would fail to form the belief that there is a barn in front of him. In this case it is only natural to think that he does not have the knowledge that he is looking at a barn. However, his belief is reliable and safe. His belief is safe because his belief could not easily have been false. Moreover, his belief is reliable because beliefs formed in the same way tend to be safe. So, safety and reliability are satisfied. This sort of example shows that one can fail to have knowledge in spite of being in a mental state that satisfies safety and reliability. But then the fact that Mike's perceptual state may be defective in spite of satisfying

safety and reliability does not show that perceptual states should not be treated as primitive knowledge states.

Some think standard knowledge requires sensitivity (see e.g. Dretske 1970, Nozick 1981, DeRose 1995). When defined on belief, *s*'s belief that *p* is sensitive iff if *p* were false, then *s* wouldn't believe that *p*. When defined on perceptual states, sensitivity may be articulated as follows:

#### Sensitivity

*s*'s perceptual state to the effect that *p* is sensitive just in case *s* does not have a perceptual experience as of *p* in any close world in which not-*p*.

Sensitivity, of course, is just a special case of causal dependence. A mental state that *p* is sensitive iff it causally depends on *p*'s being true. Requiring that knowledge be sensitive will not help with the modified barn case. If Henry were looking at a fake barn, then the demon would blur his vision and he would not form the belief that he is looking at a barn. So, Henry's belief is sensitive in the envisaged scenario. However, requiring that knowledge be sensitive helps to explain why Mike's hallucinatory experience is defective. If Mike were tied down when given the drug, then the blue ball *o* would be right in front of him in all the close worlds in which he has an experience as of the blue ball *o* being right in front of him. However, in the closest worlds in which the blue ball *o* is no longer right in front of Mike, Mike would still have had an experience as of *o* being right in front of him. So, while Mike's experience is sensitive to the fact that *o* was right in front of him, it fails to be sensitive to the fact that *o* is right in front of him. Of course, as we have already seen, one can easily conjure up cases of sensitive veridical hallucinations. If Alice would have neutralized the effects of the hallucinatory drug had the blue ball not been right in front of Mike, then Mike's experience is defective but sensitive. So, one can be in a defective

mental state that is factive, safe, reliable and sensitive, which is just to say, factivity, safety, reliability and sensitivity are not jointly sufficient for knowledge.

It is plausible that knowledge does not admit of analysis, as argued by Williamson (2001: chap. 1, for related considerations see also Brewer 1999: chap. 4). For Williamson, knowledge is a basic mental state distinct from belief. If Williamson is right, then we might say that what distinguishes good perceptual states from bad ones is a brute fact about perceptual states. Good perceptual states will be basic knowledge states, whereas bad perceptual states will fail to be basic knowledge states.

Regardless of whether knowledge is analyzable or not, the view defended here is that good and bad perceptual states are of different fundamental (metaphysical) kinds. What does that mean? One thing it could mean is that good and bad cases differ structurally. This would be so if nonlucky veridical perceptions were relations to objects, whereas lucky veridical perceptions, illusions and hallucinations were relations to contents or intentional objects. But good and bad cases might also be of different fundamental (metaphysical) kinds if, following Williamson (2000), knowledge states and non-knowledge states are of different kinds, not because they differ in structure, but because they differ with respect to which epistemic constraints they satisfy. On this view, mental states are individuated partially on the basis of their epistemic standing.

As I already mentioned, I am going to remain neutral on the question of whether good perceptual experiences have content or make facts manifest. But a word on the former position is here in order. If good and bad perceptual states differ in kind simply in virtue of one being a primitive knowledge state and the other not, then there could be two perceptual experiences with the same content and the same phenomenology, even if one of them is a primitive knowledge state and the other is not. It might, in that case, seem natural to take the constraints satisfied by good perceptual states to be conditions on the state itself rather than as constituting a layer of content. But if one is already committed to the thesis that good perceptual states have content, one could also take the constraints satisfied by good perceptual states to form a layer of non-

phenomenal content. Of course, if knowledge does not admit of analysis, the extra layer of content will not be constituted by the external constraints satisfied by primitive knowledge states but will rather be something like ‘this experience is a primitive knowledge state’. Let us call this sort of content, regardless of whether it is constituted by external constraints or not, ‘knowledge content’.

I am broadly sympathetic to the second proposal—the proposal that a layer of knowledge content enters into the content of perceptual experience.<sup>7</sup> Several considerations count in favor of this proposal. First, if the external constraints satisfied by good perceptual states constitute a layer of non-phenomenal content or perceptual states have the content ‘this experience is a primitive knowledge state’, then we can treat good perceptual states as veridical and bad perceptual states as falsidical. Veridical states are then primitive knowledge states, whereas falsidical states are not. For example, in our drug-induced hallucination case Mike’s experience fails to be a primitive knowledge state. But the knowledge content of Mike’s experience is then false. So, Mike’s experience is falsidical.

Second, while the phenomenology of experience does not give us reason to assume that perceptual experience has knowledge content,<sup>8</sup> positing a layer of knowledge content is motivated by the inferential role of perceptual experience (see Chalmers [2006] for this sort of motivation for positing a Fregean content). If Mike discovers that his perceptual state does not satisfy the constraints it must satisfy for it to be a knowledge state, he will consider it defective. For example, if it is brought to his attention that he is under the influence of a drug that prevents continued visual input from his surroundings, he will doubt that his experience is veridical. The same holds if it is brought to his attention that the content of his experience could easily be false.

## VII. Virtues of Primitive knowledge Disjunctivism

---

<sup>7</sup> Assuming, of course, that we don’t treat good perceptual cases as having no content.

<sup>8</sup> Presumably the content reflected in the phenomenology is edenic. See Chalmers 2006.

The view just offered is a kind of disjunctivism because it takes good cases of perception and bad cases of perception (hallucinations, illusions and lucky veridical perceptions) to be of different fundamental (metaphysical) kinds. Nonlucky veridical perceptions are primitive knowledge states, lucky veridical perceptions, hallucinations and illusions are not. In short: when a perceptual state is a case of primitive knowledge, that is its fundamental kind. When a perceptual state is not a case of primitive knowledge, then it has a different fundamental kind. Nonlucky veridical perceptions are primitive knowledge states, lucky veridical perceptions, hallucinations and illusions are not.

In the dismantled chair case your veridical perception just before the chair is dismantled has the same epistemic standing as your illusion after the chair is partly dismantled. In the invert case Invert's (lucky) veridical red experience has a worse epistemic standing than Invert's (normal) illusory red experiences. Lucky cases of veridical perceptions fail to be good cases of perception for two reasons. First, lucky cases of veridical perception have the same epistemic standing as the corresponding hallucinations/illusions. Since hallucinations/illusions are bad cases of perception in virtue of being epistemically inferior and lucky cases of veridical perception have the same epistemic standing as the corresponding hallucinations, lucky cases of veridical perception are bad cases of perception too. Second, lucky cases of veridical perception are not good cases of perception because they do not play the action-guiding role which good cases of perception play. Consider Williamson's (2000) example of how knowledge might play a more fundamental action-guiding role than belief. Suppose we want to explain why the burglar kept searching for the diamonds in the house for hours. The best explanation for his continued search is that he knew the diamonds were there. Had the burglar merely believed that the diamonds were in the house, he might have given up before finding them. We might extend this sort of motivation for treating knowledge as fundamental in guiding action to the case at hand. If you have a lucky veridical perception, you are prone to making mistakes. If you are an invert in a normal perceptual community you probably have learned to respond to green objects on the basis

of your phenomenally red experiences the way that others respond to red objects on the basis of their phenomenally red experiences. If you want to identify a physically green object, you go for the object that gives rise to a red experience in you. If you want to identify a physically red object, you go for the object that gives rise to a green experience in you. That one day where you have a lucky red experience caused by a physically red object you will make mistakes. For example, if asked to categorize objects according to colors you will place the physically red things in the box labeled 'green things'. This is because you have learned to call things 'green' and treat them as green based on your phenomenally red experiences. Normally you make no mistakes. The defect in your visual system has been accommodated for. But on your lucky day where red yields phenomenal red, you make mistakes.

You're not prone to making mistakes of a similar kind in the dismantled-chair case. That is, your lucky veridical perception that there is a red wooden intact chair in front of you right before the chair is dismantled will support your actions just as well as non-lucky veridical perceptions. If you try to sit down on the chair, you won't land on your behind, and if you try to touch the chair there will be an object there for you to touch, and so on. Still, there is a sense in which you are prone to error. After all, what the case shows is that you could easily have been wrong. That is just what lack of safety amounts to. If a mental state fails to be safe, then its bearer could easily have been wrong. So, there is a sense in which you are prone to error. You may sit down on what appears to you to be a red wooden intact chair at a time  $t$  without landing on your behind but had you sat down on what appears to you to be a red wooden intact chair a few instants later, then you would have landed on your behind. Since lucky perceptual states make you prone to error, they are not perfectly good perceptual states. When you are in them, you don't make mistakes, but you could easily have done so.

A final issue: should we reject the internality argument which we attributed to McDowell at the outset of the paper? We should not. Granting the truth of the second premise (viz. Internality), the argument is sound. If there is a difference in the epistemic standing of good and

bad cases of perception, then that difference cannot be grounded in a common factor shared between good and bad cases of perception. I have argued that good and bad cases could plausibly have the same content. But the sameness of content then is not what grounds the difference in epistemic standing. What grounds the difference between good and bad cases is the difference in which external constraints the states satisfy. So, just as the internality argument predicts, the difference in epistemic standing is grounded in an element of perception which is not shared between the good and bad cases. The reason standard versions of disjunctivism go wrong is simply that they take the good epistemic standing of veridical perception to be grounded in manifest facts which do not need to satisfy any epistemic constraints. For example, McDowell-style disjunctivism, as this position is most commonly construed, predicts that because Invert has a veridical perception where a physically red object causes a phenomenally red experience, her state will not have a worse epistemic standing than her normal illusory states. But this is mistaken. Invert's normal illusory color experiences put her in a position to know the physical colors of things; Invert's lucky veridical perception does not. We can, of course, preserve a version of the fact-made-manifest thesis if there are epistemic constraints on making a fact manifest. On this view, not all veridical perceptions then make a fact manifest – they have a true content, but don't achieve the status of revealing a fact. For a fact to be made manifest certain epistemic constraints must be satisfied. But this view just is a version of the primitive knowledge thesis defended in this paper.

### VIII. Conclusion

I have argued that classical attempts to motivate disjunctivism epistemically run into trouble. The trouble issues from the underlying thesis that veridical perceptions and hallucinations/illusions are of different fundamental kinds in virtue of having different epistemic standings. The good epistemic standing of veridical perceptions is supposed to be grounded in a feature of veridical perception which veridical perceptions do not share in common with hallucinations/illusions, viz.

facts made manifest to the subject. The trouble with this thesis is that there are veridical (non-hallucinatory and non-illusory) perceptions which have no better epistemic standing than the corresponding (falsidical) hallucinations/illusions. But if veridical perceptions need not have a better epistemic standing than the corresponding hallucinations/illusions in spite of the fact that veridical perceptions make facts manifest, then the difference in epistemic standing between the two kinds of states is not adequately explained by the facts-made-manifest thesis. I have further offered arguments in favor of a new kind of primitive knowledge disjunctivism. On standard views of perception, which take perception to have content, there can be veridical hallucinations and good cases of perception with exactly the same content. On McDowell-style disjunctivism, as the position is most commonly cashed out, there can be veridical hallucinations and veridical perceptions which both stand in an appropriate relation to an external object or fact. But then *veridical case of perception* and *good case of perception* cannot be necessarily co-extensive notions. I have argued that what characterizes good non-lucky, non-illusory and non-hallucinatory cases of perception is that they are primitive knowledge states. Like standard knowledge states, such states satisfy anti-luck constraints such as safety and sensitivity. Unlike standard versions of disjunctivism, primitive knowledge disjunctivism can explain why some veridical (non-illusory and non-hallucinatory) perceptions have no better epistemic standing than the corresponding hallucinations/illusions. Some veridical perceptions have no better epistemic standing than the corresponding hallucinations/illusion because they fail to be primitive knowledge states. In short: Classical disjunctivism goes wrong because it takes the good epistemic standing of veridical perception to be grounded in “manifest” facts which do not necessarily satisfy any epistemic constraints. A better form of disjunctivism explains the difference between good and bad cases in terms of epistemic constraints that the states satisfy. This allows us to preserve McDowell’s thesis that good cases make facts manifest, as long as manifest facts must satisfy epistemic constraints.<sup>9</sup>

---

<sup>9</sup> The paper has benefited from comments provided by David Chalmers, Bill Fish, Frank Jackson, Joe

## References

- Block, N. 2007. *Consciousness, Function and Representation*, volume 1, MIT Press.
- Braun, D. 2005. "Empty Names, Fictional Names, Mythical Names." *Noûs* 39: 596-631.
- Brewer, B. 1999. *Perception and Reason*, Oxford: Clarendon Press.
- Byrne, A. and H. Logue. 2008. "Either/Or", in ed. A. Haddock and F. Macpherson, ed. *Disjunctivism: Perception, Action, Knowledge*, Oxford: Oxford University Press.
- Chalmers, D. 2004. "The Representational Character of Experience", *The Future for Philosophy*, ed. B. Leiter, Oxford: Oxford University Press: 153-81.
- Chalmers, D. 2006. "Perception and the Fall from Eden", in *Perceptual Experience*, ed. T. Gendler and J. Hawthorne, Oxford: Clarendon Press: 49-125.
- DeRose, K. 1995. "Solving the Skeptical Problem", *Philosophical Review* 104: 1-52.
- Dretske, F. 1970. Epistemic Operators, *Journal of Philosophy* 67: 1007-23.
- Egan, A. 2006. "Appearance Properties?", *Nous* 40: 495-521.
- Fish, W. 2008. "Disjunctivism, Indistinguishability, and the Nature of Hallucination", in F. Macpherson & A. Haddock, ed., *Disjunctivism: Perception, Action, Knowledge*, Oxford: Oxford University Press.
- Fish, W. Forthcoming. *Perception, Hallucination, and Illusions*, New York: Oxford University Press.
- Hinton, J.M. 1973. *Experiences*, Clarendon Press: Oxford.
- Johnston, M. 2004. "That Obscure Object of Hallucination", *Philosophical Studies* 120: 113-83.
- Langsam, H. 1997. "The Theory of Appearing Defended", *Philosophical Studies* 87: 33-59.
- Martin, M.G. F. 2002. "The Transparency of Experience", *Mind and Language* 17: 376-425.
- Martin, M. G. F. 2004. "The Limits of Self-Awareness", *Philosophical Studies*, 120: 37-89.

---

Salerno, Susanna Schellenberg, Sydney Shoemaker, Daniel Stoljar and audiences at the ANU, Auckland and Otago. Special thanks to David Chalmers for written comments on the paper.

- McDowell, J. 1982. "Criteria, Defeasibility, and Knowledge", *Proceedings of the British Academy* 68: 455-79, cited as reprinted in T. Baldwin and T. Smiley, ed, *Studies in the Philosophy of Logic and Knowledge*, Oxford University Press, 2005: 7-29.
- McDowell, J. 1986. "Singular Thought and the Extent of Inner Space", in P. Petit and J. McDowell, ed. *Subject, Thought, and Context*, Oxford: Clarendon Press.
- McDowell, J. 2008. "The Disjunctive Conception of Experience as Material for a Transcendental Argument", in F. Macpherson & A. Haddock, ed., *Disjunctivism: Perception, Action, Knowledge*, Oxford: Oxford University Press.
- Neta, R. 2008. "In Defense of Disjunctivism", in F. Macpherson and A. Haddock, eds. *Disjunctivism: Perception, Action, Knowledge*.
- Nozick, R. 1981. *Philosophical Explanations*, Cambridge, MA: Harvard University Press.
- Putnam, H. 1999. *The Threefold Cord*, Columbia University Press: New York.
- Schroeder, T. and B. Caplan. 2007. "On the Content of Experience", *Philosophy and Phenomenological Research* 75: 590-611.
- Shoemaker, S. 1994. "Phenomenal Character", *Nous* 29: 21-38.
- Snowdon, P. 1980/81. "Perception, Vision and Causation", *Proceedings of the Aristotelian Society* 81: 175-92.
- Sosa, E. 1999. "How to Defeat Opposition to Moore", *Noûs* 33, Supplement: Philosophical Perspectives, 13, *Epistemology*: 141-153.
- Sosa, E. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, New York: Oxford University Press.
- Thompson, B. 2006. "Color Constancy and Russellian Representationalism", *Australasian Journal of Philosophy* 84: 75-94.
- Travis, C. 2004. "The Silence of the Senses", *Mind* 113: 57-94.
- Tye, M. 2007. "Intentionalism and the Argument from No Common Content", *Philosophical Perspectives* 21: *Philosophy of Mind*, J. Hawthorne, ed.: 589-613.

Williamson, T. 2000. *Knowledge and its Limits*, Oxford: Oxford University Press.