

The Skeptic and the Naïve Realist

Heather Logue
University of Leeds¹

Consider the perceptual experience I have when I *see* that there is a yellow banana before me, and the indistinguishable experience I have when I have a *hallucination* as of a yellow banana. It's tempting to suppose that these experiences are fundamentally the same. Obviously, they're not exactly alike in *all* respects—for one thing, the first is a veridical perceptual experience in which I perceive something in my environment, and the other is a hallucination. But since the hallucination is indistinguishable from a veridical experience of a yellow banana, one might find it natural to conclude that they have the same fundamental psychological nature—that is, that the features the experiences have in common (e.g., the tendency to produce the belief that there's a yellow banana before me, their phenomenal character) are grounded in the same kinds of psychological facts.

Of course, the fact that I can't tell that the hallucination isn't a veridical experience of a yellow banana doesn't *entail* that it has the same fundamental psychological nature as the latter. Perhaps there's a difference that I cannot detect. So there is logical space for a theory of perceptual experience that denies that the fundamental psychological nature of my hallucination is the same as that of my veridical experience.

One view that appears to entail this claim is *Naïve Realism*, which holds that veridical perceptual experiences fundamentally consist in the subject perceiving physical entities in her environment. For example, Naïve Realism holds that my veridical perceptual experience of the yellow banana in front of me is fundamentally a matter of my perceiving the banana and its yellowness. It seems that the Naïve Realist cannot give the same account of the fundamental psychological nature of my hallucination, however, as the subject of a hallucination doesn't perceive anything in her environment.

Why would one endorse Naïve Realism? Some of the view's proponents have suggested that it undermines a certain kind of argument for skepticism about the external world. This motivation for Naïve Realism is the focus of this paper. In section 1, I will clarify Naïve Realism's commitments, outline the argument for skepticism it's supposed to undermine, and present a response to that argument. In section 2, I will present a case for the conclusion that Naïve Realism affords a better version of this response than its main rival. Finally, in section 3, I will argue that this case is flawed, and I will conclude that Naïve Realism isn't in a privileged position to respond to this argument for skepticism after all.

¹ An early draft of this paper was presented at a meeting of the Centre for Metaphysics and Mind at the University of Leeds; thanks to those in attendance for their helpful questions and comments. Special thanks to Alex Byrne for detailed comments on previous drafts of this paper.

1. Naïve Realism

Let the *good case* be a situation in which I see the banana on my desk, and it is the way it looks to me to be (yellow and crescent-shaped). Moreover, nothing “funny” is going on—e.g., it’s not the case that the lighting in my office is such that the yellow banana would look green if it weren’t for the exact wavelengths of light that happen to be coming through the window. In this case, I am having a veridical perceptual experience of my surroundings.

Let a case be *subjectively indiscriminable* from the good case if and only if the subject of the case cannot tell by introspection alone that she’s not in the good case. It will be helpful to single out two subjectively indiscriminable cases for illustrative purposes. First, there is the *illusory bad case*: one in which I see a *green* banana on my desk, but because of the lighting conditions in my office it looks *yellow*. Second, there is the *hallucinatory bad case*: one in which my brain has been envatted, and it’s being stimulated such that it seems to me that I see a yellow banana on a desk. In this case, I don’t perceive anything in my surroundings at all.

Above, I characterized Naïve Realism as claiming that a veridical perceptual experience fundamentally consists in the subject perceiving physical entities in her environment. To say that an experience *fundamentally* consists in something is to say that it is that in virtue of which it has all the other psychological properties it does. For example, my experience in the good case has the following psychological properties (among others): it is an experience as of a yellow banana, it naturally generates the belief that there is a yellow banana before me, and it has a certain phenomenal character. According to Naïve Realism as I have formulated it, the experience has all of these properties ultimately in virtue of its *consisting in my perceiving the banana and certain of its properties*. We can get a better sense of what this claim amounts to by contrasting it with the *Representational View*, Naïve Realism’s main rival. The Representational View holds that the experiences in all three cases fundamentally consist in my *representing* my environment as being a certain way (e.g., as containing a yellow, crescent-shaped banana). According to the Representational View, my experiences in *all three* cases have the psychological properties on the list ultimately in virtue of *being a certain kind of representational state*.

As I mentioned above, the experience in the hallucinatory bad case cannot fundamentally consist in my perceiving anything in my environment, simply because it doesn’t involve me perceiving anything in my environment. Thus, it seems that the fundamental psychological characterization of the experience in the hallucinatory case must be different than the fundamental psychological characterization of the experience in the good case. Borrowing some terminology from Byrne and Logue 2008, let’s call the claim that hallucinations and veridical experiences are of different fundamental psychological kinds *metaphysical disjunctivism*.² Arguably, Naïve Realism entails metaphysical

² In our 2008, Alex Byrne and I characterized metaphysical disjunctivism as claiming that the experiences in the good and bad cases have nothing mental in common *at all*—we argued that the denial that they are of the same fundamental kind is not sufficient for disjunctivism. I no longer endorse this characterization

disjunctivism.³ (As for the *illusory* case: since it *does* involve me perceiving things in my environment—I perceive the banana, it just isn't the color it looks to me to be—it is in principle open to the metaphysical disjunctivist to hold that the experience is of the same fundamental kind as the experience in the good case.)⁴

Naïve Realism isn't always formulated in the same way by its proponents, so we must pause to clarify how other formulations are related to the one I've proposed. For example, M.G.F. Martin sometimes characterizes Naïve Realism as claiming that “[s]ome of the objects of perception—the concrete individuals, their properties, the events these partake in—are constituents of the experience” (2004: 39) in the good case.⁵ An apparently different formulation is offered by William Fish, who mainly characterizes Naïve Realism as a thesis about the *phenomenal character* of the experience in the good case. He holds that “the phenomenal character of [my good case experience]...is the property of acquainting [me] with such-and-such a presentational character” (2009: 15, emphasis removed), where an experience's presentational character is just the properties one perceives when having it.⁶

Arguably, these formulations of Naïve Realism are *consequences* of the one I've proposed. As for the “constituency” formulation of Naïve Realism: if the experience in the good case fundamentally consists in a certain state of affairs (my bearing the perceptual relation to the banana and its yellowness), then the constituents of this state of affairs are presumably constituents of my experience. Indeed, it's not quite clear to how to make sense of the constituency claim other than understanding it as a consequence of the formulation I've offered. As for the “phenomenal character” formulation of Naïve Realism: the phenomenal character of my good case experience is one of its psychological features. The

of metaphysical disjunctivism. The proper characterization of metaphysical disjunctivism should only be as strong as is required to preserve Naïve Realism, and the weaker characterization offered in the main text meets this requirement.

³ Mark Johnston denies this claim. On his view, hallucinations involve perception of “sensible profiles”, which are (roughly) complexes of properties and relations that would be instantiated by things in one's environment if it were as it appeared to be. Johnston would say that my experiences in all three cases fundamentally consist in perception of the same sensible profile; it's just that in the illusory case it's partially uninstantiated by my environment and in the hallucinatory case it's totally uninstantiated by my environment. For our purposes here we can set this view aside—most Naïve Realists embrace metaphysical disjunctivism presumably because they are wary of Johnston's claim that we can perceive uninstantiated properties.

⁴ This choice point is reflected in the distinction between *V v IH* and *VI v H* (metaphysical) disjunctivism drawn in Byrne and Logue 2008. *VI v H* disjunctivism treats the experience in the illusory case as being of the same fundamental kind as the experience in the good case; see Brewer 2008 for a view along these lines. *V v IH* disjunctivism treats the experience in the illusory case as being of the same fundamental kind as the experience in the hallucinatory case; although Martin does not discuss illusions in great detail, it appears that he endorses such a view (see, e.g., Martin 2006: 361).

⁵ See also Campbell 2002: 117.

⁶ See also Brewer 2008: 171 and Campbell 2002: 114-5.

claim that this experience has the *phenomenal character* it does in virtue of the obtaining of the perceptual relation is a consequence of the claim that it has *all* its psychological features ultimately in virtue of the obtaining of the perceptual relation. Given that Naïve Realists intend to say something about the metaphysical structure of veridical experience in general and not just its phenomenal character, the formulation I've proposed is more accurate.

Now that we've clarified what Naïve Realism is claiming, let us turn to the question of why someone might want to endorse it. One motivation for the view put forward by some of its proponents is that it affords a response to a certain kind of argument for skepticism about the external world. The argument goes roughly as follows:

1. If my perceptual evidence puts me in a position to know that there is a yellow banana before me, it favors the proposition that there's a yellow banana before me over the proposition that I'm a brain in a vat enjoying a hallucination as of a yellow banana.
2. The perceptual evidence I have in the *bad* cases *doesn't* favor the proposition that there's a yellow banana before me over the proposition that I'm a brain in a vat enjoying a hallucination as of a yellow banana.
3. Thus, the perceptual evidence I have in the bad cases, if I were to have it in the presence of a banana, wouldn't put me in a position to know that there's a yellow banana before me.
4. The perceptual evidence I have in the good case is the same as the perceptual evidence I have in the bad case.
5. Thus, the perceptual evidence I have in the good case doesn't put me in a position to know that there's a yellow banana before me.

From here, it's a short step to the conclusion that I *can't* be in a position to know that there's a yellow banana before me. In the good case, I'm in the best informational position I can possibly be in with respect to whether there's a yellow banana before me—so if this position isn't a position *to know*, then it must be *impossible* for me to be in a position to know.⁷ Further, since we could run analogous arguments for any subject and any proposition about the external world, the skeptic concludes that we can't have any knowledge of the external world at all.⁸

Premise 1 is an instance of an “underdetermination” principle—the principle in play here is that one is in a position to know that *p* only if one's

⁷ This isn't quite right. I could be in a better informational position if, say, my belief was confirmed by the testimony of an *expert* banana-spotter. But of course, given that I don't know that there's a yellow banana before me in the good case as described, *no amount* of expert banana-spotter testimony is going to change this (since an analogous argument applies to the expert's perceptual evidence).

⁸ Cf. the “argument from underdetermination” as outlined in, e.g., Byrne 2004: 304-6 and Vogel 2004: 426-9.

evidence favors the proposition that p over all propositions one knows to be incompatible with it. This principle seems plausible, as it's not clear how one's evidence could put one in a position to know that p if it didn't give p some sort of advantage over the competition. Premise 2 is also plausible. The best perceptual evidence I can get in the bad cases would be something along the lines of the following: it *appears* to me that there's a yellow banana before me. This claim doesn't favor the proposition that there is a yellow banana before me over the proposition that I'm a brain in a vat being stimulated so as to have a hallucination as of a yellow banana. Premise 3 follows from 1 and 2, so if we grant premise 4—the claim that I have the same perceptual evidence in the good case as I have in the bad cases—the skeptic is off and running.

But arguably, premise 4 isn't as plausible as the other premises. One could reasonably suggest that the perceptual evidence afforded by my experience in the good case *isn't the same as* the evidence afforded by my experiences in the bad cases. In conjunction with the claim that the evidence I have in the good case is good enough to put me in a position to know that there's a yellow banana before me, we have the outlines of a promising response to the argument for skepticism just outlined. Borrowing some terminology (from Byrne and Logue 2008 and Snowdon 2005), let us call this response *epistemological disjunctivism*.⁹

While metaphysical disjunctivism holds that the experiences in the good and bad cases differ with respect to their metaphysical structure, epistemological disjunctivism holds that they differ with respect to the perceptual evidence they afford. On the face of it, metaphysical disjunctivism and epistemological disjunctivism are entirely independent claims. Metaphysical disjunctivism doesn't obviously entail epistemological disjunctivism—one could consistently hold that the experiences in the good and bad cases are fundamentally different, but that this difference doesn't make for a difference in the perceptual evidence afforded by the experiences. And epistemological disjunctivism doesn't obviously entail metaphysical disjunctivism—one could consistently hold that the experiences in the good and bad cases are fundamentally the same, but that something other than a difference in fundamental psychological nature makes for a difference in the perceptual evidence afforded by the experiences. Nevertheless, one might hold that metaphysical and epistemological disjunctivism are connected through *Naïve Realism*. In particular, one might think that metaphysical disjunctivism is required to secure Naïve Realism, which in turn yields the best form of epistemological disjunctivism. The task of the next section is to investigate the second part of this claim.

⁹ For the purposes of this paper, I will assume that epistemological disjunctivism is an effective response to this argument for skepticism. As we will see, the task of this paper is essentially to compare *different versions* of epistemological disjunctivism. Thus, a defense of this kind of anti-skeptical strategy from general objections that would apply to all versions of it would be beyond its scope.

2. Naïve Realism and epistemological disjunctivism

Some Naïve Realists have suggested that their theory provides the best version (or perhaps even the only viable version) of epistemological disjunctivism. William Fish expresses sympathy with this way of motivating Naïve Realism, although he thinks further investigation is required to determine whether it is sound (2009: 23-6). Matthew Kennedy endorses a motivation for Naïve Realism along these lines in a recent paper in passing (2010: 78-9), although his project there isn't to argue for it. Mark Johnston suggests that any view but a Naïve Realist one is doomed to skepticism:

...so long as our sensory starting points are taken to be items 'in our own minds,' be they qualia or narrowly supervening propositional contents, one can only agree with the skeptic's suggestion that the environment-characterizing contents of our immediate perceptual judgments are conjectural, and neglect a host of equally good alternatives left open by our entertaining such subjective mental items (2006: 286).

The idea seems to be that as long as the good case experience is a wholly "inner" state, a view like epistemological disjunctivism can't get off the ground: if the good case experience doesn't consist in a relation to my environment (just like the experiences in the bad cases), it's not clear how there could be a difference in perceptual evidence across the good and bad cases.¹⁰

If Naïve Realism provides the best or only viable version of epistemological disjunctivism, and if epistemological disjunctivism affords a defensible response to the argument for skepticism presented above, this would be a powerful consideration in favor of Naïve Realism. In this section, I will argue that although Naïve Realism's main rival can support a version of epistemological disjunctivism, there is a respect in which the Naïve Realist's version seems to be superior. In particular, a Naïve Realist seems to be in a better position to give an account of *what it is* about my experience in the good case that puts me in a position to know things about my environment. But first, we need to get clearer about what exactly the evidential difference across the good and bad cases might be. Then we will be able to determine which of the various candidates for the evidential difference can be accommodated by the rival theories of perceptual experience.

¹⁰ A word is in order about John McDowell's place in the dialectic. He is one of the most prominent proponents of epistemological disjunctivism, and deserves much of the credit for putting the view on everyone's radar. Although he champions epistemological disjunctivism, it's not clear whether he endorses metaphysical disjunctivism (see Byrne and Logue 2008: 65-8 for McDowell exegesis on this point). And if he isn't a metaphysical disjunctivist, he's almost certainly not a Naïve Realist. He might be a non-Naïve Realist epistemological disjunctivist, and if he is, he'd be sympathetic to the conclusion of this paper.

a. What is the evidential difference?

When you have a perceptual experience, it gives you at least some evidence for propositions about your environment. For example, in the good and bad cases, my experiences afford me evidence for the proposition that there's a yellow banana before me (in the bad cases, misleading evidence). But what evidence do the experiences provide, exactly? For our purposes, we're interested in the evidence that constitutes the evidential difference between the good and bad cases—the perceptual evidence I have in the good case and lack in the bad cases.¹¹

The first choice point concerns whether the relevant perceptual evidence is *psychological*, i.e., whether it concerns my experience, as opposed to just what I am experiencing. If the relevant perceptual evidence is psychological, we face another choice about whether it is *environment-entailing*, i.e., whether the psychological state at issue entails that there's a yellow banana before me. For example, if I perceive the banana on my desk and its yellowness, or perceive that the banana on my desk is yellow, my being in these psychological states entails that there is a yellow banana before me. By contrast, its appearing to me that there is a yellow banana before me does not entail that there is a yellow banana before me (I might be a brain in a vat, or I might be looking at a green banana).

With these choices clearly in view, we're now in a position to articulate what my "good case" perceptual evidence for the proposition that there's a yellow banana before me might be. First, if the relevant perceptual evidence is non-psychological, it might be something along the lines of the following:

(1) that there is a yellow banana before me

Of course, there's room to quibble about exactly which non-psychological propositions my experience serves up as evidence. Perhaps my experience yields *object-dependent* propositions as evidence (e.g., that *this banana* is yellow). For our purposes, we just need one plausible candidate—what I will have to say about this one will apply to all other non-psychological candidates.

Second, if the relevant perceptual evidence is psychological and non-environment-entailing, it would be something along the lines of the following:

(2) that it looks to me as if there is a yellow banana before me

Again, there are other candidates that fit this description; for example, a proponent of the Representational View might suggest that the relevant

¹¹ I will assume that perceptual evidence is *propositional* in form, as opposed to consisting of objects, events, or states of affairs. However, for every proposition put forward as perceptual evidence, there's a correlated state of affairs that would presumably serve as perceptual evidence if it isn't propositional (e.g., we can replace the proposition *that the banana is yellow* with the state of affairs *the banana's being yellow*, or replace the proposition *that I see the banana and its yellowness* with the state of affairs *my seeing the banana and its yellowness*). The arguments to come apply equally to the propositions and (with some minor tweaks) to their state of affairs counterparts.

perceptual evidence is the proposition that I *perceptually represent* that there is a yellow banana before me. Again, what I will have to say about (2) will apply to any other psychological, non-environment-entailing candidate.

Third, if the relevant perceptual evidence is psychological and environment-entailing, it might be something along the lines of the following:

(3) that I see that there is a yellow banana before me

Another candidate under this heading is the following:

(4) that I see the banana and its yellowness.¹²

Now, which of these four candidates constitutes the epistemological disjunctivist's difference in perceptual evidence across the good and bad cases? Let's begin with (2), the proposition that it looks to me as if there is a yellow banana before me. Something along the lines of (2) must constitute my perceptual evidence in the *bad* cases—otherwise, I wouldn't have any perceptual evidence for the claim that there is a yellow banana before me in those cases. But this just goes to show that (2) cannot constitute a *difference* in perceptual evidence across the good and bad cases. Since it looks to me as if there is a yellow banana before me in all three cases, my perceptual evidence in the *good* case must concern more than just non-environment entailing psychological states if epistemological disjunctivism is true.

What about (1), the proposition that there is a yellow banana before me? To avoid the criticism that felled (2), (1) has to be evidence that I have only in the good case. And since it perceptually appears to me that there is a yellow banana before me in all three cases, it might seem that (1) has no chance of satisfying this criterion.

One might suggest that evidence can't be *false*, and so insist that I don't have (1) as evidence in the illusory and hallucinatory cases (since it's false that there's a yellow banana before me in those cases). However, this won't do—to see this, let's consider another case that is subjectively indiscriminable from the good case. Let the *veridical illusory* bad case be one in which I see a yellow banana that looks yellow to me, but only as a result of the "canceling out" of two illusion-generating features of the situation—each of which *alone* would have caused the banana to look some color other than yellow. In this case, it's *true* that

¹² What does it mean to say that I *perceive the banana's yellowness*? This question opens a big can of worms. For example: it's plausible that I perceive the banana only if it plays a role in causing my experience. If perception of properties has the same structure as perception of individuals that instantiate them, then I perceive the banana's yellowness only if its yellowness plays a role in causing my experience. This means that the banana's yellowness is causally efficacious, which suggests that it is a trope rather than an abstract universal. So it appears that either perception of properties doesn't have the same structure as perception of individuals, or that the proper account of property perception entails that there are tropes. I explore this can of worms in detail in Logue ms.. In the context of this paper, I will set these issues aside, and assume that there is an acceptable Naïve Realist account of property perception.

there's a yellow banana before me, but intuitively I'm not in a position to know this. And if (1) is evidence I have in both the good and veridical illusory cases, it cannot be my good case perceptual evidence (i.e., that which puts me in a position to know that there is a yellow banana before me).

A better strategy is to follow Williamson (2000) in identifying one's evidence with one's *knowledge*. The idea is that since I know that there is a yellow banana before me in the good case, but don't know this in the bad cases, there is an evidential difference across the good and bad cases—the proposition that there is a yellow banana before me is part of my perceptual evidence only in the good case. However, one could resist this move on the following grounds: any account of what the evidential difference is should explain *what it is* about the good case that puts me in a position to know that there's a yellow banana before me. If we specify the evidential difference across the cases in terms of what I *know*, we're simply presupposing that I have this knowledge. One might have hoped for more; namely, an account of the evidential difference that specifies *the facts in virtue of which* I'm in a position to know that there is a yellow banana before me in the good case.

What about (3), the proposition that I see that there is a yellow banana before me? It faces the same objection just outlined. To say that I *see that* there's a yellow banana before me is just to say that I know (on the basis of vision) that there's a yellow banana before me. And to say that my perceptual evidence for the claim that there's a yellow banana before me is that I *know* that there's a yellow banana before me fails to illuminate what it is about the good case that puts me in a position to know this in the first place.

So we're left with (4), the proposition that I see the banana and its yellowness. Fortunately, (4) avoids the criticisms that have felled the other candidates. We aren't compelled to include it in my perceptual evidence in the bad cases, as with (2). Moreover, (4) *does* elucidate what it is about the good case that puts me in a position to know that there's a yellow banana before me—namely, that in the good case, and in the good case *alone*, I *see* the banana and its yellowness.

Of course, one might question whether the reason given for throwing out (1) and (3) is legitimate—one might deny that an account of what the evidential difference is must *explain* what it is about the good case that puts me in a position to know that there is a yellow banana before me. However, I will assume that this is a desideratum on an account of the evidential difference. For given this desideratum, the Naïve Realist can construct a case for the superiority of her version of epistemological disjunctivism. So let us grant this desideratum for the sake of argument, and consider the Naïve Realist's case.

b. Naïve Realism, the Representational View, and the evidential difference

Given that an account of the evidential difference between the good and bad cases must explain what it is about the good case that puts me in a position to know that there is a yellow banana before me, the best candidate for my good case perceptual evidence is the proposition that I *see* the banana and its yellowness. This appears to be the only candidate that could explain what puts me in a position to know that there is a yellow banana before me.

The Naïve Realist account of the good case experience can easily accommodate the claim that this proposition constitutes the evidential difference across the cases. According to the Naïve Realist, my experience in the good case fundamentally consists in my perceiving (more specifically, seeing) the banana and its yellowness. So the Naïve Realist can say that my good case experience puts me in a position to know that there is a yellow banana before me simply because it is fundamentally a seeing of the banana before me and its yellowness.

What about the Representational View? Can it accommodate the claim that the evidential difference is the proposition that I see the banana and its yellowness? Not straightforwardly. Sure enough, the Representationalist holds that I see the banana and its yellowness. But she also holds that the obtaining of the perceptual relation between me and the banana isn't part of the fundamental psychological nature *of the experience*—the latter is a matter of my representing my environment as being a certain way (something I do in all the cases). So, on the face of it, the Representationalist cannot point to some aspect of my *experience* that puts me in a position to know that there's a yellow banana before me.

Of course, the Representationalist can point to a feature of *the good case* that puts me in a position to know this (the fact that I see the banana and its yellowness). Although this feature of the good case isn't a feature of the *experience* in it, it is a state of affairs that presumably obtains partly in virtue of my having the experience. But this is a rather odd way of characterizing my epistemic situation. Arguably, part of our notion of a perceptual experience is that it's the source of perceptual evidence; but on the picture just suggested, the crucial bit of perceptual evidence that differentiates the good case from the bad ones isn't provided by the experience itself. More importantly: if the proposition I perceptually represent, or the fact that I perceptually represent it, doesn't constitute the evidential difference across the good and bad cases, then the fact that I perceptually represent my environment as being a certain way isn't the most fundamental psychological explanation of the good case experience's epistemological power. Thus, it seems that the Representationalist would have to concede some ground to the Naïve Realist in order to secure a form of epistemological disjunctivism that meets the desideratum at issue. So the desideratum under discussion should be refined. An account of the evidential difference should not only explain what it is about the good case that puts me in a position to know that there's a yellow banana before me—it should also explain how *the experience* in the good case puts me in this position.

I should emphasize that this criticism of the Representational View's anti-skeptical power is predicated on this desideratum. Rejecting it would open the door to Representationalist versions of epistemological disjunctivism that seem to be just as effective as a Naïve Realist epistemological disjunctivism. For example, the Representationalist could say that one's perceptual evidence consists of the propositions one perceptually represents *and* knows (as on the Williamsonian elaboration of (1) discussed in the previous subsection). On this version of epistemological disjunctivism, in the good case my perceptual evidence includes the proposition that there is a yellow banana before me; not so in the bad cases.

To summarize: Naïve Realism isn't the only route to epistemological disjunctivism—you can also get there via the Representational View. But it looks like the only versions of epistemological disjunctivism you can get from the Representational View are ones that give up on the quest for an account of *what it is* about the experience in the good case that puts me in a position to know things about my environment. We're assuming (for the sake of argument) that a version of epistemological disjunctivism that can perform this explanatory task would be the best possible result. Naïve Realism appears to be in a better position than its main rival to secure this result—but can it deliver? This question is the topic of the next section.

3. Is the Naïve Realist's epistemological disjunctivism superior?

Let us revisit the dialectic between the skeptic and the Naïve Realist. The skeptic claims that I have exactly the same evidence for claims about the external world in both the good and bad cases, and that this evidence supports the proposition that I'm a brain in a vat being stimulated so as to have an experience as of a yellow banana just as much as it supports the proposition that there's a yellow banana before me. Given that knowing something requires having evidence that supports it over alternatives known to be incompatible with it, the skeptic concludes that I'm not in a position to know that there's a yellow banana before me in the good case. It's a short step from here to skepticism about the external world.

The Naïve Realist replies by invoking epistemological disjunctivism—he claims that I have much better perceptual evidence for claims about my environment in the good case than I do in the bad case. In the good case, my evidence includes the proposition *that I see the banana and its yellowness*, and this evidence is supposed to be sufficient to put me in a position to know that there's a yellow banana before me. Moreover, unlike the Representationalist's version of epistemological disjunctivism, the good case perceptual evidence the Naïve Realist appeals to is alleged to *explain* how my good case experience puts me in a position to know that there's a yellow banana before me.

However, the Naïve Realist's reply is not yet complete. We require assurance that the proposition that I see the banana and its yellowness is really part of my evidence. For if the correct account of what my evidence is leads to the conclusion that this proposition *isn't* part of my evidence, then the Naïve Realist's anti-skeptical strategy doesn't even get off the ground. So we need to consider what it takes for a proposition to be part of my evidence.

As mentioned earlier, Williamson holds that one's evidence *just is* what one knows (2000: Ch. 9). On this view, all it takes for proposition *p* to be part of my evidence is for me to know it. If this is right, then verifying that *p* is part of my evidence requires an account of how I know it. But even if one doesn't want to follow Williamson in identifying one's evidence with one's knowledge, it still seems that an account of how I know that *p* plays a crucial role in verifying that *p* is part of my evidence. To see this, let's start by considering whether one can have *p* as evidence for something even if one bears no psychological relation to *p* whatsoever. This proposal seems implausible; presumably, part of what makes an item of evidence *mine* is that I bear some psychological relation to it. While a

proposition I've never so much as momentarily entertained might count as evidence for some other claim in an "impersonal" sense (concerned only with logical relations between propositions), it's not evidence that *I have*.

So it seems that a necessary condition on my having a proposition as evidence is that I bear some psychological relation to it. Which one? A natural suggestion is that the relation is *belief*. However, even if believing that *p* is necessary for having it as evidence, it isn't sufficient—one might have all sorts of irrational beliefs, and irrationally held beliefs don't count as evidence for anything. (If somebody believed that the Earth is flat, we wouldn't regard him as having evidence for the claim that the Earth isn't a sphere.)

So which psychological relation must one bear to *p* in order to have it as evidence, if not *knowledge*? Natural candidates are being *justified in believing* that *p*, and being in a *position to know* that *p*. I won't endorse any of these three possibilities here. For our purposes, the important point is that they are all closely connected to knowledge, connected in such a way that verifying that *p* is part of one's evidence requires at least a partial account of how one knows (or could come to know) that *p*.

An account of being in a position to know that *p* is parasitic on an account of how one could come to know that *p*—roughly speaking, you give the account of how one could come to know that *p*, and specify which necessary conditions for knowing aren't met but easily could be. Thus, if having *p* as evidence consists in being in a *position to know* that *p*, verifying that one has *p* as evidence requires an account of how one could come to know that *p*. As for being *justified in believing* that *p*: given that justified belief is necessary for knowledge, to give an account of justified belief that *p* is to give *part* of an account of how one knows that *p*. Hence, if having *p* as evidence consists in being justified in believing that *p*, verifying that one has *p* as evidence requires part of an account of how one knows that *p*.¹³

In short, on any plausible view about what one's evidence is, verifying that *p* is part of one's evidence requires at least a partial account of how one could come to know that *p*. The proposition we're interested is *that I see the banana and its yellowness*. So in order to verify that it's part of my evidence, we

¹³ What if justified belief isn't necessary for knowledge? Views that deny a justification condition on knowledge typically hold that evidence can *undermine* but not *ground* knowledge. For example, on a reliabilist theory of knowledge, if I know that *p*, it's because the process that produced my belief that *p* is a reliable one—I need not have evidence for *p*, and if I do, that's not the ultimate explanation of why I know that *p*. A reliabilist about knowledge might tweak her view in order to accommodate the intuition that I cannot know that *p* if my total evidence favors not-*p*. However, on such a view evidence only plays the role of undermining knowledge—it is not that in virtue of which one *has* knowledge. This kind of view is antithetical to epistemological disjunctivism—on the latter, I know that there is a yellow banana before me primarily in virtue of the fact that I have *evidence* that puts me in a position to know this (evidence I lack in the bad cases). Thus, a combination of Naïve Realist epistemological disjunctivism with a view that denies a justification condition on knowledge seems contradictory, and so I will set it aside.

need to investigate how I could come to *know* it. So let us turn to potential accounts of how one could come to know that one sees something.

According to Duncan Pritchard, the core of McDowell's anti-skeptical strategy is the claim that (contrary to philosophical orthodoxy) one can know by introspection that one is in an environment-entailing psychological state (Pritchard 2008: 290)—e.g., that one can know by introspection that one sees that *p* or that one sees *o*. If this is right, then I am in a position to know *by introspection* that I see the banana and its yellowness.¹⁴

However, this just pushes the bump around under the rug—how exactly can one come to know by introspection that one is in an environment-entailing psychological state? McDowell himself doesn't offer much in the way of an answer to this question. He writes:

What does entitle one to claim that one is perceiving that things are thus and so, when one is so entitled? The fact that one is perceiving that things are thus and so. That is a kind of fact whose obtaining our *self-consciously possessed perceptual capacities* enable us to recognize on suitable occasions, just as they enable us to recognize such facts as that there are red cubes in front of us, and all the more complex types of environmental facts that our powers to perceive things put at our disposal (2008: 387, emphasis mine).

This is a start, but much more needs to be said: *how*, exactly, do one's perceptual capacities afford knowledge that one is seeing? The claim that they are "self-consciously possessed" sheds no light on this issue.¹⁵

Alan Millar offers a somewhat more developed suggestion along similar lines:

¹⁴ A clarificatory remark regarding McDowell exegesis: on my interpretation of McDowell, it's not clear whether he endorses Naïve Realism (see note 10 above). If he doesn't, he shouldn't be read as marshaling this claim in support of the Naïve Realist's version of epistemological disjunctivism (as I am doing on the Naïve Realist's behalf here).

¹⁵ Johnston makes another suggestion that is subject to a similar criticism. He writes: "Among the things I can be aware of are my own mental acts, for example, my seeing the spoon on the table. Such mental acts are events, which like environmental events can be objects of awareness. If I attend to this act and attend to the property of being a seeing by me of a spoon on the table, and then predicate the property I have attended to, then my higher-order judgment that I am seeing the spoon on the table will have been made out of its present truthmakers [and in the absence of defeaters, will count as knowledge]" (2006: 287-8). However, we may ask: what exactly does one's awareness of one's own mental acts consist in? And in what sense can we "attend" to them? Johnston distances himself from the claim that there is *literally* an "inner sense" (2006: 285), but if that's not how we become aware of and attend to our mental states, then what alternative does he endorse? In the absence of an account of the "*mechanics* of self-knowledge" (Gertler 2000, emphasis mine), this suggestion about how one comes by knowledge that one is seeing isn't very illuminating.

...at least in typical situations when we know by looking that something is an F we are aware, that is know, that we have seen the thing to be an F. In those situations the visual experiences that, via the exercise of the relevant *perceptual recognitional ability*, enable us to know that something is an F, also enable us to tell that we see that thing to be an F. This latter knowledge is made possible by the exercise of a higher-order recognitional ability—an ability to tell whether or not one sees that an F is there (2008: 342, emphasis mine)

Millar is telling us a bit more than McDowell was. Knowing that one is seeing an F isn't merely grounded in one's capacity to *perceive* F's—it's grounded in one's ability to *recognize* F's by way of perception. (Perceptually recognizing an F involves more than just perceiving it; it requires having the concept of an F, and being disposed to apply that concept to things that have a certain appearance distinctive of F's.) However, we're still left wanting more, for we may ask: *how*, exactly, do our perceptual *recognitional* abilities afford knowledge of one's environment-entailing psychological states (like seeing an F)? Millar suggests that there's some connection between the ability to perceptually recognize F's and a "higher-order recognitional ability" to tell whether or not one is seeing an F. But what exactly is the connection? Does the former give rise to the latter? If so, how? And how tight is the connection? Is it possible to have the lower-order recognitional capacity without the higher-order one? These crucial questions are left unanswered.

Matthew Kennedy builds on Millar's suggestion by appealing to one's capacity to recognize something as being an *object of sight*. The idea is this: just as a subject can perceptually recognize something as being a banana, a subject (with the appropriate concepts) can perceptually recognize something as being an object of sight. Moreover, my perceptual recognition of the thing before me *as being a banana and being yellow*, when combined with perceptual recognition of it *as being an object of sight*, yields knowledge that *I see the banana and its yellowness* (Kennedy forthcoming).¹⁶ This seems plausible—if one has perceptually recognized something as being an F, as having property G, and as being an object of sight (i.e., something *that one sees*), then one has knowledge that one *sees* an F and its G-ness.¹⁷

This account of how we know that we see is a bit dubious once you unpack the details. To perceptually recognize something as being an F is to know

¹⁶ I'm assuming that perceptual recognition of properties is analogous to perceptual recognition of objects—just as perceptual recognition *of a banana* is to know that something is a banana on the basis of how it perceptually appears, perceptual recognition of something *as being yellow* is to know that it is yellow on the basis of how it perceptually appears.

¹⁷ Presumably, a similar story would be told about what *justifies my belief* that I'm seeing the banana and its yellowness: e.g., that my justification for that belief consists in my perceptually recognizing the thing before me as being a banana, as being yellow, and as being an object of my sight. So if one's perceptual evidence consists of propositions one is *justified in believing* (rather than knows or is in a position to know), one would appeal to this story in verifying that the proposition under discussion is part of my evidence.

that it is an F on the basis of how it perceptually appears. So, to perceptually recognize something as being an object of sight is to know that it is an object of sight on the basis of its appearance. However, it's not clear that one can know that something is an object of sight *on the basis of its appearance*. When one perceptually recognizes something as being an F, one does so by way of perceiving an appearance *distinctive of F's*. For example, when one perceptually recognizes something as being a banana, one does so by perceiving its yellowishness and crescent-shapedness. However, what is the appearance distinctive of *an object of sight*? The appearances of the things we see are quite varied—they can be greenish-brown, octopus-shaped, and gigantic; hot-pink and green, dress-shaped, and human-sized; and so on ad nauseam. Perhaps the appearance distinctive of an object of sight is quite general—just a matter having *some* color, shape, and size. But this can't be quite right, as some objects of sight are colorless, and some don't really have any particular shape or size at all (e.g., volumes whose boundaries are out of view, like the sky). It's just not clear that we can isolate any sort of appearance as being distinctive of objects of sight, and thus it isn't clear that we can *perceptually recognize* things as being objects of sight.

But let us turn to an objection that is more important in this dialectical context. Kennedy's account of my knowledge that I see the banana and its yellowness is essentially as follows: I perceptually recognize the thing before me as being a banana, as being yellow, and as being an object of sight. These acts of perceptual recognition are supposed to ground my knowledge that I see the banana and its yellowness. Notice, however, that to say that I perceptually recognize the thing before me as being a banana, as being yellow, and as being an object of sight is just to say that I *know* that the thing before me is a banana, yellow, and an object of sight (on the basis of how it appears). One can't recognize something as being an F (being G) and fail to know that it is an F (that it is G); recognition is supposed to be a *success* state (see Millar 2008: 333). Thus, I *know* that there is a yellow banana before me. Kennedy's account of my knowledge that I see the banana and its yellowness *presupposes* that I'm in a position to know that there is a yellow banana before me.

If the Naïve Realist appeals to this account of knowledge that one is seeing, he undermines his claim to a superior version of epistemological disjunctivism. Remember, we're assuming that an account of the evidential difference between the good and bad cases must explain how my good case experience puts me in a position to know that there is a yellow banana before me. As argued in section 2b, Naïve Realism is uniquely situated to deliver the only candidate for my good case perceptual evidence that appears to satisfy this desideratum, viz., the proposition that I see the banana and its yellowness. However, it now seems that this candidate, when fully elaborated with the help of Kennedy's account of knowledge that one is seeing, cannot satisfy the desideratum. If the proposition that I see the banana and its yellowness is part of my evidence because I *know* it, then I know that there is a yellow banana before me (because I have carried out the required acts of perceptual recognition). If this proposition is part of my evidence because I am in a *position to know* it, then I'm in a position to know that there's a yellow banana before me (because I'm in a position to carry out the required acts of perceptual recognition). If this proposition is part of my evidence because I'm *justified in believing* it, then I

know that there is a yellow banana before me (because I have carried out the required acts of perceptual recognition, which are presumably required for justification). No matter what it takes for a proposition to count as part of my evidence, the claim that I have the proposition that I'm seeing the banana and its yellowness as evidence *presupposes* that I'm in a position to know that there's a yellow banana before me. Thus, it cannot figure in an *explanation* of how my good case experience puts me in this position.

In short, the Naïve Realist's attempt to *explain* how my good case experience puts me in a position to know that there's a yellow banana before me ultimately rests on the claim that I'm in a position to know that there is a yellow banana before me. He attempts to explain how my good case experience puts me in a position to know that there is a yellow banana before me by claiming that it gives me perceptual evidence to the effect that I see the banana and its yellowness; but then he turns around and attempts to explain my possession of this evidence in terms of my being in a position to know that there is a yellow banana before me. If this perceptual evidence is supposed to explain how my good case experience puts me in a position to know that *p*, the explanation of why I have this evidence shouldn't ultimately rest on the fact that I am in a position to know that *p*.

In light of these considerations, the Naïve Realist's candidate for my good case perceptual evidence (the proposition that I see the banana and its yellowness) no longer seems any better than some of the other candidates considered in section 2a (in particular, the proposition that I see *that* the banana is yellow, or the proposition that the banana is yellow supplemented with the identification of one's evidence with one's knowledge). These candidates were rejected on the grounds that they do not explain what it is about my good case experience that puts me in a position to know that there is a yellow banana before me, because they presuppose that I am in a position to know that there is a yellow banana before me. But it turns out that the Naïve Realist's candidate doesn't offer an explanation of this fact either, because it also makes this presupposition. Thus, all of these candidates appear to be on equal footing.

The Naïve Realist might defend the superiority of his candidate by noting that his account of my good case evidence doesn't just *flatly assert* that I know that there's a yellow banana before me. If we say that my good case perceptual evidence is that I *see that* there is a yellow banana before me (i.e., that I *know* on the basis of vision that there is a yellow banana before me), or that my good case evidence is that there is a yellow banana before me, and that this is part of my perceptual evidence because I *know* it, we aren't saying anything remotely illuminating about what puts me in a position to know it. By contrast, at least the Naïve Realist has a bit of a story to tell. This story begins with the claim that my good case perceptual evidence is that I see the banana and its yellowness. Now, it's true that the story about why this is part of my perceptual evidence ultimately appeals to the fact that I'm in a position to know that there is a yellow banana before me. But perhaps the moral is simply that my knowledge that there is a yellow banana before me and my knowledge that I see the banana and its yellowness are inextricably intertwined. At least the Naïve Realist has given an account that elucidates the relationship between my knowledge of my environment and my knowledge that I perceive things in it.

However, we must remember that the reason given for preferring the Naïve Realist's candidate for my good case perceptual evidence in the first place was that it is supposed to *explain* how my good case experience puts me in a position to know that there is a yellow banana before me. And the story that the Naïve Realist can tell merely generates the *illusion* of explanation. When the Naïve Realist attempts to explain how my good case experience puts me in a position to know that there's a yellow banana before me by claiming that my good case perceptual evidence is the proposition that I see the banana and its yellowness, this claim seems explanatory only because it isn't *prima facie* obvious that it presupposes what it's supposed to explain. The fact that an account doesn't wear all of its commitments on its sleeve isn't a theoretical virtue. Given that all three candidates for my good case perceptual evidence under discussion presuppose that I'm in a position to know that there's a yellow banana before me, none of them seems to do a better job of *explaining* how my experience puts me in this position than the others.

Let us take stock. We began this section with the idea that a candidate for my good case perceptual evidence ought to explain how my perceptual experience puts me in a position to know that there is a yellow banana before me, and that Naïve Realism is uniquely situated to deliver the only candidate that appears to achieve this (*viz.*, the proposition that I see the banana and its yellowness). I then argued that establishing that this proposition is part of my perceptual evidence requires an account of how I could come to know it (regardless of whether we identify evidence with knowledge). Next, we considered several potential accounts of how I could come to know that I see the banana and its yellowness. The accounts suggested by McDowell and Millar shed little light on how one knows that one is seeing. Kennedy's elaboration of Millar's account is more illuminating, but it is poorly suited to the Naïve Realist's purposes in this dialectical context. The Naïve Realist needs an account of how I can come to know that I'm seeing the banana and its yellowness that can be used in an explanation of how my good case experience puts me in a position to know that there is a yellow banana before me. But since Kennedy's account assumes that I'm in a position to know that there is a yellow banana before me, it cannot be marshaled in an explanation of this fact.¹⁸

4. Conclusion

In the debate between Naïve Realists and Representationalists about perceptual experience, one point some Naïve Realists have suggested that they have in their column is that their view affords a distinctive response to skepticism about the external world. In particular, since Naïve Realism holds that the metaphysical structure of my experience in the good case literally includes things in my environment, it is a natural fit with epistemological disjunctivism—which holds

¹⁸ Alex Byrne offers an account of knowledge that one is seeing (see Byrne forthcoming). Like Kennedy's account, Byrne's account also presupposes that I am in a position to know that there is a yellow banana before me; so the same point would apply. Unfortunately, I don't have the space to elaborate Byrne's account here.

that I have much better evidence for propositions about my environment in the good case, indeed, evidence that puts me in a position to know them.

However, as we saw in section 2, epistemological disjunctivism can be combined with a Representationalist account of perceptual experience, which undermines Naïve Realism's claim to an anti-skeptical advantage. The Naïve Realist could attempt to resurrect the alleged advantage by offering a version of epistemological disjunctivism that explains how my good case experience puts me in a position to know things about my environment. I argued that, unlike the Representationalist, the Naïve Realist can easily say that my good case perceptual evidence for the proposition that there's a yellow banana before me is *that I see the banana and its yellowness*, and that this candidate is the only one that appears to be up to the explanatory task. Granting the assumption that we ought to give such an explanation if we can, the Naïve Realist appears to have one up on the Representationalist.

It's worth noting how precarious the Naïve Realist's case for having a better version of epistemological disjunctivism is. For it crucially depends on a couple of debatable claims: first, that an account of my good case perceptual evidence must explain how my experience puts me in a position to know things about my environment, and second, that the Representational View is in tension with the claim that my good case perceptual evidence is that I see the banana and its yellowness. I find both of these claims plausible, but I have not provided watertight arguments for them. Thus, it could be that the Naïve Realist's case falters when it comes to one or both of these claims.

However, even granting these claims for the sake of argument, it appears that the Naïve Realist's version of epistemological disjunctivism isn't any better than the versions afforded by its main rival. As I argued in section 3, the Naïve Realist's best shot at an explanation of how my good case experience puts me in a position to know that there's a yellow banana before me ultimately rests on the claim that I am in this position, and so it doesn't really explain how my experience put me in this position. Thus, Naïve Realism's epistemological disjunctivism appears to be on a par with the versions afforded by its rival, and so Naïve Realism has no privileged claim to any anti-skeptical advantages epistemological disjunctivism may yield.

References

- Brewer, B. 2008. How to account for illusion. *Disjunctivism: Perception, Action, Knowledge*, ed. A. Haddock and F. Macpherson. Oxford: Oxford University Press.
- Byrne, A. 2004. How hard are the sceptical paradoxes? *Nous* 38: 299-325.
- . forthcoming. Knowing what I see. *Introspection and Consciousness*, ed. D. Smithies and D. Stoljar. Oxford: Oxford University Press.
- Byrne, A., and H. Logue. 2008. Either/Or. *Disjunctivism: Perception, Action, Knowledge*, ed. A. Haddock and F. Macpherson. Oxford: Oxford University Press.
- Campbell, J. 2002. *Reference and Consciousness*. Oxford: Oxford University Press.
- Fish, W. 2009. *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.

- Gertler, B. 2000. The mechanics of self-knowledge. *Philosophical Topics* 28: 125-46.
- Johnston, M. 2006. Better than mere knowledge? The function of sensory awareness. *Perceptual Experience*, ed. T. Gendler and J. Hawthorne. Oxford: Oxford University Press.
- Kennedy, M. 2010. Naive Realism and experiential evidence. *Proceedings of the Aristotelian Society* 110: 77-109.
- . forthcoming. Naive Realism, privileged access, and epistemic safety. *Nous*.
- Logue, H. ms. Naive Realism's sophisticated account of veridical experience.
- Martin, M. G. F. 2004. The limits of self-awareness. *Philosophical Studies* 120: 37-89.
- . 2006. On being alienated. *Perceptual Experience*, ed. T. Gendler and J. Hawthorne. Oxford: Oxford University Press.
- McDowell, J. 2008. The disjunctive conception of experience as material for a transcendental argument. *Disjunctivism: Perception, Action, Knowledge*, ed. A. Haddock and F. Macpherson. Oxford: Oxford University Press.
- Millar, A. 2008. Perceptual-recognitional abilities and perceptual knowledge. *Disjunctivism: Perception, Action, Knowledge*, ed. A. Haddock and F. Macpherson. Oxford: Oxford University Press.
- Pritchard, D. 2008. McDowellian neo-Mooreanism. *Disjunctivism: Perception, Action, Knowledge*, ed. A. Haddock and F. Macpherson. Oxford: Oxford University Press.
- Snowdon, P. F. 2005. The formulation of disjunctivism: a response to Fish. *Proceedings of the Aristotelian Society* 105: 129-41.
- Vogel, J. 2004. Skeptical arguments. *Philosophical Issues* 14: 426-55.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.