Chapter 44
Responses to Commentators: Berinstein, Kovach, McDowell, Neta, Sethi, Smithies

Crispin Wright

I am most grateful to all my commentators for their searching and interesting reactions to "Perceptual Justification: Two Conceptions Compared". I cannot, in the space allotted to me here, react to all the issues they raise but I will comment briefly on what impress me as the most important for the purposes of the agenda of the PEER conference and the volume.

* * *

1. *On the bearing of the New Evil Demon Scenario on the relative merits of Internalist and Externalist conceptions of perceptual knowledge and justification.*

Is the New Evil Demon Scenario (NEDS) victim epistemically justified in the systematically false beliefs they form based on the character of what they take to be their perceptual experiences or do the latter merely provide them with *excuses* for what are actually unjustified beliefs? In the paper I suggest that the answer to the question, which is the better account,

> depends on the relative weighting assigned to values of epistemic product —truth and knowledge —on the one hand and values of doxastic management —rational integration and explanatory coherence — on the other,

and go on to suggest that since the opposition between broadly externalist and internalist conceptions of knowledge and justification in general already springs from and reflects such a disparity in relative weighting, the NEDS thought experiment is actually useless as a means for deciding which account has the better of the issue.

Declan Smithies responds that

C. Wright (✉)
University of Stirling, Stirling, UK
e-mail: cjgwright@gmail.com

M. Vuletić and O. Beck (eds.), *Empirical Reason and Sensory Experience*, Logic, Epistemology, and the Unity of Science 60,
https://doi.org/10.1007/978-3-031-52231-4_44

> On this view, the substantive question is whether epistemic properties that satisfy internalist criteria (e.g rationality) are more valuable than epistemic values that satisfy externalist criteria (e.g knowledge) or vice versa. But I don't think this can be right. To the extent that these epistemic values are commensurable at all, I don't see how an internalist can reasonably claim that rationality is more valuable than knowledge. And I don't see why they should. Given a plausible epistemic pluralism, an internalist need only claim that rationality is an important dimension of epistemic value. They need not claim that it's the only one, the most basic one, or the most valuable one. Its merely one among many important and mutually irreducible epistemic values.[1]

Smithies' reaction evinces a misunderstanding, for which I am wholly to blame, of the point I intended to which **Adam Kovach** also succumbs when he writes,

> In Crispin Wright's diagnosis, internalists and externalists about epistemic justification disagree about whether the NEDS victim has justified perceptual beliefs because they disagree more fundamentally about the relative importance of two kinds of epistemic values. The externalist gives greater weight to values of the epistemic product, such as truth or knowledge, but the internalist favors values of epistemic agency ("doxastic management"), such as rationality or coherent belief.[2]

Let me now try to express my intended point better. Values in general—maybe better: valuables— come in a variety of kinds. Some things are *intrinsically* valuable—valuable in virtue of certain of their intrinsic properties—and others are valuable in virtue of certain of their properties that they do not have intrinsically. A different distinction is that between things which are *finally* valuable—valuable for reasons other than a tendency toward the promotion or enablement of other values—and things that are instrumentally valuable, whose value derives exactly from their tendency to promote or enable the promotion of other values. The crucial distinction for the putative significance of NEDS is the latter—the distinction between final and instrumental *epistemic* value. My externalist holds that the only genuinely final epistemic value is knowledge, conceived in some way congenial to externalism—as true belief reliably formed, or safely formed, or sensitively formed, e.g.—and that values of doxastic management (coherence, rationality, explanatory integration) are merely instrumentally valuable and that only to the extent that they promote or enable the achievement of knowledge. For such an externalist, once the achievement of perceptual knowledge is foreclosed as in NEDS, there is strictly no point in managing belief in a coherent, rational fashion, certainly no sense in which one *ought* so to organize one's beliefs and is thereby justified in doing so. My internalist[3] holds, to the contrary that, however knowledge should be conceived and while knowledge

[1] This volume, p. 314.

[2] This volume, p. 317.

[3] Why "my internalist"? I am aware that internalism about perceptual justification is traditionally identified with the idea that the justifiers are ultimately just the internal phenomenal states of the subject. But this, I suggest, is not consistent with the idea that knowledge is the sole final epistemic value. The former idea entails that the NEDS subject and their phenomenal duplicate in the GOOD case have the same justification for their beliefs, so if either is believing as they ought, both are. But conceiving of knowledge as the *sole final* epistemic value surely entails that, when knowledge is impossible, there is actually nothing that the subject ought to think. The tension can be resolved only by accrediting rationally coherent belief management with value in both circumstances. And

is undoubtedly a significant—maybe even a supreme—epistemic value, rational, coherent belief management is also a final epistemic value in its own right and can thus sustain justification and entail epistemic 'oughts' even in a far-fetched scenario where true belief, and hence knowledge, are frustrated by the very doing of what one epistemically ought.

The issue between externalism and internalism so conceived is thus not which if either is the dominant value, as Smithies and Kovach[4] take it, but whether rational coherence is a final or merely an instrumental value. It is to be compared to the issue between a consequentialist who holds that fairness in the distribution of goods is valuable only insofar as it promotes general well being, and one who holds that fairness is a good even when it has no tendency to enhance general well-being. Such a position involves no commitment to, though it may go along with, a particular relative weighting of general well-being promotion and fairness.

2. *Is there a Problem with the Philosophical Methodology implicit in the invocation of the New Evil Demon Scenario?*

Sofia Berinstein thinks there is, but I am not sure I have managed to understand her concern. She writes,

> The articulation of the skeptical scenario doesn't aid us in deciding whether justification arises from the dyadic relationship or the intrinsic properties. For instance, in NEDS, since it is contentious if the perceiver is justified or not, there is no settling of the question of the source of the justificatory charge. Those who think that it came from the intrinsic properties of the mental state will tend to believe that the perceiver is justified and those who think that it came from the relationship to the worldly environment will tend to believe that it is not. Since whether the perceiver is justified or unjustified is not like whether a lamp is on or off—something that we can simply see and agree upon—the skeptical scenario doesn't arbitrate effectively. Rather, it encourages one to think that *either* justification depends on the perceptual mental state *or* on the world-involving dyadic states—to commit to internalism over externalism or vice versa.[5]

So far, that is, more or less, the point I was making—that (bearing in mind the clarification I have offered above in response to Smithies and Kovach) NEDS serves merely to elicit competing conceptions of epistemic value. But now Berinstein continues,

> I want to raise the concern that these value priorities are not demanded by our inquiry into perception as a natural phenomenon. If doxastic product and doxastic process primarily come into contest with each other because of artificial skeptical scenarios, then it may not be necessary to frame these values as being in competition. The value of thought experiments like NEDS might not then be worth the cost, methodologically. Because they focus our inquiry on disagreements over how to describe instances in which pre-doxastic perceptual states are severed from their consequent perceptual beliefs, skeptical scenarios like NEDS abstract away from the connection we seek to understand.[6]

---

obviously this value cannot be instrumental value in circumstances where rationally coherent belief management doesn't have any instrumental value.

[4] Kovach expresses some distance from my internalist's conception of the matter. I'll consider his reservations further below. (§3).

[5] This volume, p. 322.

[6] This volume, p. 322.

I am not sure I want to agree that philosophical enquiry into the nature of perceptual justification is or should be primarily concerned with "perception as a natural phenomenon". That is, while it is a perfectly good question

> … how it is that our perceptual experiences can generate perceptual beliefs and how these beliefs can be suitable for us to conduct empirical inquiry and empirical reasoning on their basis …[7]

it doesn't impress as the same question as that which I took to be central to our concerns: viz. the question how must or may we conceive of perceptual experience in order to have a smooth account of how such experience justifies the beliefs that it triggers, indeed in the best case how it renders them knowledgeable. When the latter is the central question, then—or so I want to argue—consideration of the values whose pursuit underwrites justification (where that consists in doing what, in the light of those values, one ought), has to take centre stage. And, when there is an apparent plurality of relevant such values, the construction of scenarios in which they are differentially subserved or obstructed has surely to be of relevance to the question of the way we conceive of their relative importance, character and relationships. A familiar example from the ethical sphere is, once again, the testing of our reactions to scenarios in which utility is maximised at the expense of distributive fairness and conversely. NEDS, it seems to me, is at least a perfectly valid thought experiment when conceived as standing to teach you something about your own epistemic values.

That said, however, Berinstein is incontestably right that, as we ordinarily think of these matters, process values in doxastic management are interlinked with product values—with the achievement of knowledge, truth and understanding. Manage your beliefs badly—ignore questions of mutual coherence, fitting of evidence, freedom from bias and prejudice, etc.—at your peril. Unless you are very lucky, you are thereby likely, so we think, to wind up bumping against the world. However, that is only to say that, as we believe, the pursuit of rational coherence does have instrumental value for the pursuit of knowledge—or at least represents our best chance of getting knowledge. That is perfectly consistent with the reaction of those who share Stewart Cohen's reaction to NEDS. One value can be *both* instrumental in the attainment of another and final in its own right. Indeed arguably something like that relationship obtains de facto between fairness and utility.

So, I see no reason to doubt that we can assign to the NEDS scenario and its ilk a legitimate role in helping us to identify the values we actually have: if you want to know what are your values in some domain of activity, and whether they are final, or instrumental, intrinsic of extrinsic, or whatever, then list some plausible candidates and explore your reactions to scenarios in which they are differentially realised and/or frustrated. I do not find anything in Berinstein's discussion to cast doubt on the validity or interest of that project.

However, all that said, there is arguably a separate *metanormative* project, on which these points do not encroach: the project of trying to determine whether we have *correct* or at least defensible values in some domain. I can imagine a hard-headed

[7] This volume, p. 319.

reliabilist arguing that, whatever about our actual values of belief management, it is (i) questionable to what extent we do actually *manage* our beliefs in the first place—whether the notion of epistemic responsibility has any real, generally applicable operational content—and (ii) open to question whether "what really matters" is not simply the attainment of the truth. It should be clear that the sympathies of my paper are with a relatively full-blooded notion of epistemic responsibility that abrades with that. But I cannot embark on discussion of the metanormative issues here.

3. *Does the Revealed Evil Demon scenario undermine the Internalist reaction to NEDS?*

Adam Kovach is unsympathetic to the internalist suggestion that NEDS brings to our attention (that we have) a final epistemic value of rational coherence. He constructs a new scenario—the *Revealed Evil Demon Scenario* (REDS)—that, reporting on his own reaction, may be taken to suggest that rational coherence has no such status.

The Revealed Evil Demon tells us that there is a significant likelihood that from hereon he will regularly cause us to suffer false perceptual experiences but that we will not know when. We believe him. Kovach puts the question:

> Would you still aim at coherence between your beliefs and your experiences in this case? Would you still care about looking things over carefully? Aside from aesthetic appreciation, I don't see much point in that. When it comes to perception, I think I value truth conductivity [sic] over coherence between experience and perceptual beliefs. I think I care about *that* kind of coherence only to the extent that I believe it takes me toward truth.[8]

This certainly looks like a useful case to consider. For of course *one* explanation of the impression that, in forming perceptual beliefs as usual, the NEDS subject is justified—is doing what they ought—is, not that they thereby successfully pursue a final value of rational coherence, but rather that they unsuccessfully but coherently pursue knowledge and that the value that excites the internalist intuition about the case is not actually a final value of doxastic management, but a derivative—instrumental or enabling—value that we attach to unsuccessful but competent and well intentioned attempts at a final value of knowledge. This value would be an instance of the value that in general we (some of us) place on "giving it your best shot", *provided* you can reasonably, if erroneously, take it that by doing so, you *do* improve your chances of success, even if the circumstances are such that doing so is not actually going to improve your chances of success. And the case for offering an account along these lines of the reaction to NEDS that allows that the hapless subject is in some sense doing doxastically well will be reinforced should it prove that our reaction to the REDS scenario is different; that is, if with Kovach we do indeed feel that the pursuit of rational coherence is, for the REDS subject, who knows of the Demon and its plan, without any epistemic value or purpose.[9]

However, I don't think REDS achieves that. Put yourself in the subject's shoes. Once you accept that there is a significant likelihood of deception in any particular

[8] This volume, p. 318.

[9] I am not sure why Kovach didn't simply add to the NEDS scenario the detail that the Demon again declares his hand. It seems the NEDS scenario, so tweaked, would equally well serve his purpose.

case, is it any longer rational to form *any* external world beliefs—any 'takings to be true'—as such at all? The question the scenario is meant to address is whether a final value of rational coherence operates over the management of a system of *belief*. An example where it is questionable whether it is rational to go in for believings in the first place can teach us nothing about that.[10]

4. *On the Anti-sceptical Significance of Disjunctivism*

Three commentators—**John McDowell**, **Ram Neta** and **Umrao Sethi**—are sympathetic to what has come to be known as Disjunctivism about perceptual justification. I take the root thought here to be metaphysical: a rejection of the idea that an episode of normal perception and a phenomenologically perfect but hallucinatory counterfeit of it are, as far as the state of mind of the subject is concerned, states of exactly the same type. Rather the former is *constitutively* a relational, world-involving state while the latter is a purely phenomenal state.

Now we can of course conjure such a distinction by stipulation, as when, for example, we fix it that nothing counts as a state of nettle rash unless actually caused by abrasion of the skin by the little barbs on the leaves of nettles, so that an otherwise perfect physiological duplicate of the characteristic pathology of the skin that nettles cause but which is actually otherwise caused, is, necessarily, not nettle rash. There is absolutely nothing amiss with such a stipulation. But equally it is of little importance if e.g. the practical significance of that pathology of the skin—for instance, its appropriate dermatological treatment—is exactly the same however it is caused. Likewise the metaphysical disjunctivist's thesis is of little significance if the epistemological consequences of the two state-types thereby distinguished are exactly the same. So an interesting disjunctivism has to include the idea that this is not so. Henceforward by "disjunctivism" and its cognates I will mean: the *interesting* versions.

The two questions I want to focus on here are:

(i) What does it take to underwrite a position according to which the epistemological significance of the two state types is indeed importantly different? And:

(ii) How exactly, if at all, does accrediting the two state types with differing epistemological significance make for a satisfying response to scepticism about perceptual knowledge/justification?

The correct answers, I believe, are respectively

(i) Justificational externalism

and

(ii) It doesn't.

McDowell and I have exchanged conflicting ideas about these matters on more than one previous occasion[11] and I must forbear here to rehearse the details of those

[10] There are interesting further issues raised by REDS. It is clear that in the REDS predicament you should do *something*, ... or die! But I cannot pursue these here.

[11] For example in: McDowell (1982), Wright (2002), McDowell (2006), Wright (2006).

exchanges. Nevertheless it's good to have a further opportunity to try to explain my reservations about the view.

It will be helpful to begin with a version of Moore's 'Proof' but in a disjunctivist setting. So suppose Moore had reasoned as follows:

Disjunctive Moore:

(i) Either (the GOOD case) I am perceiving my hand in front of my face or (the BAD case) I am in a state that phenomenologically seems exactly like GOOD but in which no genuine perception is taking place.
(ii) I am perceiving my hand in front of my face.
(iii) There is an external material world (since a hand is a material object existing in space).

The obvious first-pass rejoinder to this version of the proof when—whatever the actual historical Moore may have intended for his original version—it is presented as a response to scepticism, is that it is unclear how the transition is to be sanctioned from (i) to (ii). If, that is to say, the hedged disjunctive premise, (i), is a full articulation of Moore's epistemic situation when, as it seems to him, he holds up his hand and sees it before him, then what justifies the transition to (ii), with its presupposition that he is in GOOD?

I take it that a disjunctivist will want to reply that this version of the 'proof' in effect starts in the wrong place: that it undersells Moore's starting epistemic situation and that the correct starting point is rather, in effect, (ii) itself.

The question is how this claim is to be made out. One way of doing that is simply directly to invoke an externalist notion of perceptual justification. Moore, that is to say, would be justified in claim (ii) by the very fact—assuming it is a fact—that he is indeed perceiving the external hand in question.

It can be difficult to distinguish McDowell's response from that, but it is clear that a difference has to be intended. For he writes

> On an internalist conception, the justification for a perceptual belief is an experience that is, as Wright puts it, a state purely of the subject.
>
> Wright presupposes that internalism cannot accommodate a conception of perceptual experience as "a *world-involving* dyadic state, involving real relations between the subject and external circumstances". States purely of the subject, as Wright understands that specification, do not include perceptual experiences conceived as relating their subject to environmental circumstances in such a way that the subject could not be in such a state if the environmental circumstances did not obtain.
>
> This has the effect that an attractive conception of how perceptual experience figures in the epistemology of perceptual knowledge, and, derivatively, in the justification for perceptual beliefs, goes unconsidered. Wright does not argue against it; it simply has no place in his mapping of possible positions.
>
> According to the conception I mean, a state in which a subject perceives that things are a certain way in her environment is *both* a state purely of the subject *and* a state that the subject could not be in if things were not that way in her environment.[12]

[12] This volume, p. 325.

The intent of these remarks is, pretty clearly, to canvass a conception of the perceivable aspects of the external world which views them as potential ingredients of states of consciousness, as things of which a subject is directly aware and thereby as figuring in of states of mind that properly count as *internal*. For McDowell, accordingly, internalism about perceptual justification can properly avail itself of a conception of perceptual awareness that already embraces aspects of the external world. When a subject is in such a state of awareness of the worldly circumstance that P, they enjoy a *factive* warrant for the claim that P but, in contrast with the externalist version of the account, that they enjoy such a warrant is itself something of which they are self-consciously, non-inferentially aware—a circumstance not external to the mind but something of which they are directly conscious. Asked how he knows there is a hand in view, Moore can now reply—by way, as McDowell wants us to conceive it, of an authoritative report of *his state of mind*—that "I can see it".

McDowell describes this idea as "attractive". I think the question is whether it is ultimately intelligible.

I noted above that an externalist conception of perceptual justification can straightforwardly underwrite the differential epistemic significance of GOOD and BAD experience more or less as disjunctivism intends. But, again, McDowell is clear that externalist justification isn't what he is after. And I agree that it leaves us with a shortfall. Elsewhere, I have characterised the shortfall by saying that what we want—what externalism doesn't give us—is to *be in position rationally to claim* that we are in GOOD. To be sure, an externalist can understand that operator—"being in position to claim that P"—as likewise a matter of, possibly opaque to the subject, external situation—and I cannot here embark on the task of enforcing the reading I want it to have.[13] Fortunately, I think McDowell would be happy to grant that reading. Where, as I interpret matters, we differ concerns what puts us in position rationally to claim we are in GOOD. My proposal, developed in a number of other places,[14] is to involve the resource of a conception of epistemic entitlement that allows for the rational acceptance of certain propositions on no evidential basis. McDowell, by contrast, seems to be thinking of one's situation in GOOD as, in a recent terminology, *luminous*—that is, as a condition, which, when one is in it, is effectively recognizable as such. Thus it is not merely that one is endowed, by perception, with a direct awareness of situations in the external material world. More—a point that it is crucial not to confuse with that—one is also endowed, as something of which one is in position to be directly conscious, with an awareness that one is in that way perceptually directly aware of material worldly situations.

So, at any rate, is how I interpret this remark:

> Being in a position to know one perceives that P is part of what it is to be in a state in which one perceives that P. The subject knows she perceives that P by virtue of *self-consciously perceiving that* P. [My italics] That is obviously not a ground on which she might claim to

[13] For such an attempt, see my (2008).

[14] See in particular Wright (2004a, b, 2014).

> know that she perceives that P, but it is the answer to the question how it can be that she knows she perceives that P.[15]

When, a little later, he anticipates the reservation that

> In various ways, it can seem to someone that she perceives something to be the case when the seeming is a mere seeming.
>
> Does this threaten the idea that a subject who perceives that P is enabled to know that she perceives that P by the fact that she self-consciously perceives that P?

his answer is of course to reject the notion that there is in fact any such threat. I take it then that the core thesis of McDowellian disjunctivism is that one's perceiving that P in GOOD is something of which one can simply be conscious. Obviously there can be nothing of that form of which we are conscious when in BAD. So, purely in terms of states of which we can be conscious, our epistemic situations in GOOD and BAD respectively crucially differ. There is no need to 'go external' to secure the distinction. On this account, premise (i) in *Disjunctive Moore* is indeed, unnecessarily modest from an internalist standpoint.

So, here is the most immediately natural objection to this proposal. The objector protests:

> But the experience of the subject in BAD is *stipulated* as phenomenologically indistinguishable from the experience of the subject in GOOD. If that stipulation is allowed, then if you were in BAD, you wouldn't be able to tell from the character of your experience that you were. So how can being in GOOD possibly be something of which you are simply conscious? What can intimate to you that you are not in that rogue BAD state?

In discussion,[16] Disjunctivists often seem to be satisfied with the following rejoinder to this: from the fact that, should not-P be the case, I would not or might not be able to tell as much, it by no means follows that should P be the case, I might not be able to recognise *that*.

The *form* of this rejoinder is perfectly apt in two kinds of situation: first, when, under not-P conditions, I would not have the same cognitive capacities as I enjoy when P; and, second, when a verification of P, whenever true, can be accomplished by a method which cannot or cannot be guaranteed to deliver a verification of not-P, if it is true.[17]

---

[15] This volume, p. 326.

[16] I seem to remember McDowell himself making this move at the PEER conference.

[17] In chapter two of his *Philosophical Explanations* (Harvard 1981), Robert Nozick famously proposed an account of knowledge that required that knowledge should be *sensitive* to the fact known: in particular, that had the fact that P not obtained, the knowing subject (using the same method etc./…etc.…) would not have come to believe that P. When the suggestion that, if perceiving normally in Good, a subject can be directly conscious that they are, is coupled with the admission that were they in BAD, they would be unable to tell, the result is thus an egregious violation of sensitivity. Nozick's account has of course been roundly criticized over the ensuing years, but there is no doubt in my mind that it hits off one aspect of our concept of knowledge sufficiently well to ensure that special explanation is needed of the putative persistence of knowledge in any case where sensitivity fails.

As an example of the first kind, let P be "I am perfectly sober". If it is true, I will be cognitively lucid, able to remember that I have not had any alcohol recently or not enough to be even slightly intoxicated, etc., so will be in position to recognise P's truth. But if I am drunk, I may be drunk enough to be deluded into thinking I am perfectly sober.

As an example of the second kind, let P be the proposition that some particular formula, S, is a theorem of 1st order classical logic. Then there is a partial decision procedure for P: set a Turing machine to recursively enumerate the theorems. If S is one of them, it will show up in a finite time. But if S is not a theorem, we know by the undecidability of 1st order classical logic that the fact may never show in finite time.

Neither template seems apt to deliver what McDowell needs. There need be no suggestion that being in BAD need involve impairment of any of the cognitive systems involved in normal perception in the way that the fact of drunkenness might be masked by the very fact of being drunk. BAD can be stipulated to be a scenario wherein only the *distal* causes of one's experience are interfered with. And the theorem-hood analogy seems quite inapt: what is the procedure that detects, when one is in GOOD, that one is, but may fail to detect, when one is in a BAD but perfect phenomenological counterfeit of GOOD, that one is in BAD? It cannot have anything to do with the phenomenology of ones experience, since that is common ground, by hypothesis. But what else might a procedure that was up to the task have to go on?

The critical thrust of the foregoing is the demand for a needed but missing explanation. To be sure, a disjunctivist can reply that the relevant 'procedure' is simply to attend to whether one is "self-consciously perceiving" worldly situations—by hypothesis, a state of consciousness that is unavailable in BAD but according to disjunctivism available in GOOD. But the question at issue is whether that one is perceiving that P is indeed properly included in the range of psychological states which, when one is in them, are states of which one is indeed expectably directly conscious. So the reply is merely tantamount to a bald reaffirmation that this is the case. It does nothing to dislodge the opposing idea that perfect phenomenological duplication encompasses all aspects of consciousness of which a subject can be directly aware.

McDowell's would-be internalist disjunctivism confronts a simple dilemma. Its root idea is that the having of GOOD experience is something of which one can simply be self-consciously aware. But now if it is allowed that BAD may be so subtly contrived that if in BAD, there may be no way of telling that one is—and it is with that concession that the whole dialectic about Cartesian scepticism gets started—then either GOOD experience has some detectable feature, F, which goes missing when one is in BAD, or it does not. Suppose it does. Then the account of what it is to be in BAD has to explain what F is and why when in BAD, it is not salient to one that one's experience lacks F. No account of what F might be or why its *absence* can fail to be salient in BAD has been offered.

If on the other hand there is no such detectable feature, then the idea that one can know, just by having GOOD experience, that one is indeed in GOOD becomes magical.

However all that may be, significantly many contemporary epistemologists seem to be pretty much on board with McDowell's disjunctivism. **Ram Neta** is one such. In his present comment he advances a thought that, if sustained, would begin to address the charge of magic:

> Just as the fact that I have such-and-such perceptual experiences puts me in a position to know by reflection alone that I am having those experiences, so too, the fact that I have such-and-such *factive* [my italics] perceptual experiences puts me in a position to know, by reflection alone, that I am having those factive perceptual experiences.[18]

On this suggestion, it is thus *reflection* that in GOOD clues one in to the factivity of one's experience. It is *reflection* that can advise one, not (just) that one's experience is in all respects as if P, but that one is experiencing that P is the case.

This strikes me, however, as no doubt it will many others, as indeed a remarkably implausible claim if "reflection" is understood to connote some form of *pure thought*. What kind of faculty is properly described as "reflection" that can draw a distinction that requires a difference in causal provenance? Neta does not in his present contribution respond to that question[19] but he anticipates an objection which makes it clear that he is indeed thinking of the relevant notion of reflection as an a priori capacity:

> Such a view may seem incredible: *how*, we might wonder, could reflection alone disclose to me that my perceptual experiences are factive?

—as well we might!—

> Indeed, if it could do so, then couldn't I gain non-empirical knowledge of my perceptible surroundings themselves, by deduction from the knowledge that reflection alone provides for me?
>
> The answer to this last question is clearly no, and that's because, as Wright himself first pointed out years ago,[20] the closure of knowledge under known entailment doesn't imply that knowledge *transmits* across such entailments. While knowing that I see a cat on the mat may require me to know that there is a cat on the mat, it doesn't follow that I can acquire the latter knowledge by inference from the former. Closure leaves open that some of the things that I can know by reflection alone are things that I can know only if I *independently* know some other things empirically.[21]

However, it seems to me very doubtful whether the phenomenon of transmission failure is of service to Neta to deflect the worry of 'easy knowledge' that he anticipates. Here is not the place to embark on an investigation of the details of the whys and wherefores of it, but the core phenomenon, I suggest, is of cases were a rational subject's antecedent *open-mindedness* about the conclusion, C, of a valid

---

[18] This volume, p. 330. In their (2007) Neta and Duncan Pritchard argue that this is McDowell's view. Pritchard has of course subsequently endorsed and defended the view in depth in his *Epistemological Disjunctivism.* (Oxford 2012).

[19] Neta defends the suggestion in his (2019a, b).

[20] Neta cites my (2000). For an updated discussion, see Wright (2023).

[21] This volume, p. 330.

argument—that is, their correctly taking it that they have so far no grounds either to accept or to reject C—creates a context in which they rationally ought *also* to be open minded about the force of the particular credentials offered for one or more of its premises. Such an argument cannot be rationally cogent for such a subject. However, this template doesn't look a good fit for Neta's range of cases. Its applicability to those cases would require that anterior open-mindedness about the presence of a cat on the mat (perhaps because one has not yet entered the room wherein the mat in placed) would undermine the capacity of reflection to impart the knowledge that, on entering the room, one is seeing a cat on the mat. But the whole point of the manoeuvre with refection, I take it, is that it is to work in such a way that it is to be simply apparent on reflection, when one is, that one is seeing a cat on the mat, that that is to be a situation of which reflection can bestow a conscious awareness, independently of any anterior open-mindedness about the presence of the cat.

If transmission failures are indeed one and all, in one way or another, of the kind just adumbrated, transmission-failure is not a phenomenon that can be put to the service of bolstering Neta's claim against the objection he anticipates. So we are no further forward. Our question was: how is it, supposedly, manifest, when one is in GOOD, that one is, whereas when in BAD, a phenomenologically perfect duplicate state, that fact is not manifest? Neta, like Disjunctivists generally, doesn't explain how refection works to make one aware when in GOOD that one genuinely perceives but cannot deliver knowledge, when one is in BAD, that one does not. Moreover his suggestion remains, so far, vulnerable to the Easy Knowledge difficulty that he anticipates.

5. *A Transcendental Argument for the reliability of our Total Evidence*

However Neta follows up with an argument that would finesse the need for the play with reflection. It is an argument that

> our entitlement to regard the whole of our sensory engagement with the external world as, by and large, genuine[22]

can be corroborated by pure reason. We take as a premise what Neta terms the *Principle of Total Evidence*:

> (1) It is rational for us to proportion our confidence in a proposition to the degree to which that proposition is supported by our total evidence. (PTE)

However, this seems plausible:

> (2) It is rational for us to proportion our confidence in a proposition to the degree to which that proposition is supported by our total evidence *only if* it is rational for us to be confident that our total evidence is not generally misleading.

Hence

> (3) It is rational for us to be confident that our total evidence is not generally misleading — (from 1 and 2).

[22] This volume, p. 330.

So, by the PTE again,

> (4) Our total evidence supports the proposition that our total evidence is not generally misleading — (from 1 and 3).

This ingenious argument is deductively valid. Is it cogent?

There is a tension that emerges like this. What if one accepts the idea that "metaphysically heavyweight" hinges are, for reasons brought out by sceptical arguments (e.g. the I-II-III argument of Wright (2004a, b)), beyond evidential support? In that case, the PTE entails that we should place no confidence in those propositions. And as a result, since confidence in them is needed for rational confidence in the plethora of local, quotidian propositions that they underwrite, the PTE will commit us to an attitude of no confidence in the latter.

So far from furnishing us, then, with an a priori route to justified confidence in the non-misleading character of our total evidence, Neta's argument is tantamount to the assumption that every proposition in which we are rationally confident is one which is supported by evidence. That is in effect the assumption that the various sceptical paradoxes have mistaken conclusions. It is of course congenial to think so. But the argument presupposes it in its premise, rather than justifying it. It accomplishes nothing towards identifying the mistakes.

6. *Disjunctivism as the Solution to a Content Problem?*

Umrao Sethi argues that Disjunctivism solves a problem about the content of experience that besets both orthodox internalist and orthodox externalist approaches to perceptual justification. She glosses her Disjunctivism as the idea that

> (c)ontrary to Wright's presupposition, one can hold an externalist account of experience while nonetheless accepting an internalist account of justification. On this kind of view, one holds that the character and content of a perceptual experience are fixed by facts external to the skull, but also that the subject ought to align her beliefs with the evidence that is available to her. I'll argue that metaphysical disjunctivism—a particular brand of externalism about experience—is best understood as motivated by an internalist conception of justification. I will argue that this combination of commitments places disjunctivism in the unique position of simultaneously accommodating the epistemic values of both rationality and knowledge.[23]

Sethi's argument is two horned. First she questions whether any purely phenomenological conception of experience like that featured in the NEDS scenario—a conception whereby the role of the external world is reduced to one of brute external causality, so that the very same kinds of experience are available to the NEDS subject as to a normal perceiver—can succeed in explaining how experience enjoys any specific content nor therefore how it can justify at all. To substantiate this claim, Sethi draws on the familiar thought (call it the "Putnam point") of Chapter 2 in *Reason Truth and History* that

> … an item cannot intrinsically represent anything.[24] A line drawing in the sand cannot, just in virtue of its intrinsic characteristics, constitute a drawing of Churchill. And mental

[23] This volume, p. 334.

[24] Putnam (1981).

> items are no different from physical items here. A mental state cannot, just in virtue of its intrinsic properties, represent anything outside of itself. So, on an internalist view of perceptual experience, it can never so much as seem to the subject as if the world is some way.[25]

The consequence, Sethi contends, is that experience, conceived as by internalism, cannot justify at all.

The second horn of Sethi's argument reacts against the standard externalist response to the Putnam point that what is needed in order for experience to carry potentially justificatory content is exactly causality:

> A pattern in the sand comes to represent Churchill only if there is some causal chain that links the production of the pattern back to the individual himself. Similarly, a perceptual experience comes to represent a bowl of fruit on the table only if it is typically caused by bowls of fruit. The content of an experience, on this standard variety of externalism, is fixed by the causal etiology of the experience.[26]

But this, Sethi contends, entails that, absent antecedent knowledge of the ways in which external states of affairs have been causally active in one's experience, experience can carry no potentially justifying content "from the point of view of the experiencing subject", and so

> This brand of externalism about experience, then, cannot respect the epistemic value of rationality, according to which a subject must form beliefs solely on the basis of the evidence that is accessible to her (because no such evidence is accessible). At best then, it can aim for *knowledge* in place of rationality and argue that the subject's perceptual beliefs are nonetheless justified because they are formed on the basis of experiences that reliably lead the subject to true beliefs and knowledge.[27]

It's notable that both horns rest on major contestable assumptions. The first presupposes that the justificatory power of experience rests on its carrying content—in order to justify, it is being supposed, experience has to deliver a message, as it were. Standing opposed to that assumption are, of course, the various accounts that hold that the justificatory potential of experience is accountable in terms its non-contentual properties when taken in conjunction with a background theory or 'view' of the world.

It is notable that a further major assumption is also needed to underwrite the second horn of Sethi's argument: the assumption at work, for example, in the so-called McKinsey paradox, that if the content of a thinker's thought depends on facts about the external causation of their use of the symbols which that thought configures, then the thinker can knowledgably identify their thought contents only if they know those facts. The issues here have of course been much discussed. I will not pursue them here.

Sethi's recommended solution to the bind which she takes it that the two horns present is, exactly, the kind of conception of perceptual experience recommended by

---

[25] This volume, p. 334.

[26] This volume, ibid.

[27] This volume, p. 335.

Disjunctivism: let the justifying perceptual experience involve direct awareness of the external state of affairs that confers truth on the proposition it justifies. Now there can be a rational response to the experience consisting in the formation of a belief with a content appropriate to the manifest content of the experience that at the same time can be knowledgeable since made true by the very state of affairs of which the experience constitutes a direct awareness.

This way of arguing for the conception of perceptual experience and its justificatory role that Disjunctivism recommends would, naturally, be compromised by the worries mooted above in discussion of McDowell if those are well founded. But bracketing those concerns, I have two points to make in response.

First, it's worth pointing out that the question raised by the NEDS scenario need not presuppose any conception of perceptual content that is jeopardised by the Putnam point. Any sensible version of the view that perceptual experiences justify in virtue of carrying content has to allow that hallucinations, to take the extreme case, can have content. Whatever the correct account of the conditions necessary for that possibility, it must take one of two forms. First, it may involve that the relevant conditions are indeed consistent with the subject's always having been in the NEDS scenario—so that the Putnam point is simply wrong. In that case, Sethi's argument is blocked. Second, it may involve that the relevant conditions are indeed not consistent with the subject's always having been in the NEDS scenario. In that case, imagine them realised at some (sufficiently extended, if need be) earlier period but then cancelled as the New Evil Demon sets about its prank. The question still arises whether and if so why the subject is still justified in following their ordinary perceptual belief routines by the experiences they have after the NEDS scenario overtakes them, when those routines become completely unreliable. So the dialectical role of the NEDS thought experiment—its canvassing of the suggestion that epistemic justification can persist in the absence of reliability, or safety, or sensitivity, or whatever other externalist condition—need not be compromised by the Putnam point.

However I think there is a more basic point. Suppose that we grant that the Putnam point is sound and hence that in the 'all time' version of the NEDS scenario, the subject has no genuinely contentful experience, with the result that—on conceptions of the justificatory power of perceptual experience that require its contentfulness—they are indeed in no position to justify anything on the basis of the contents of their experience. Content-externalist views come with the possibility of certain kinds of *illusion* of content. Most of us would, for example, accept an externalist view of singular *demonstrative* content. When we use demonstratives, real external worldly semantic relations are required before we can succeed in thinking anything with a determinate truth condition. Suppose a Victorian lepidopterist on an Amazonian expedition, looking for rare and exotic species. Unmindful of the slightly strange-tasting but rather delicious stew of the local fungi she lunched on, and seemingly catching a sight of a huge butterfly, brilliantly patterned in orange, purple and aquamarine, she thinks to themself, "That is a hitherto undocumented species and will make me famous if I can catch and chloroform it, and bring it undamaged back to Cambridge." Rising from her camping stool, she reaches for her net. She does not think anything true. But nor is there any *determinately false* singular thought which

is exactly what she thinks—the nearest determinately false thought will be an existential thought: that there is a brilliantly coloured butterfly over there… etc. In so far as the content she intentionally thinks purports to be a *demonstrative* content, it is an illusory content. If we nevertheless regard her resulting action—reaching for the net and rising from her stool—as, in context, rationalised by her thought process, we have to allow that the subject is in a contentually illusory state that nevertheless makes rational sense of their response. She is doing what she rationally ought to do, given what she thinks and wants. And if that is right, then, large though the issue is, we have to allow that contentually illusory states can contribute to the rationalization of practical responses.

So why not epistemic responses too? And in that case, the Putnam point cannot straightforwardly sustain the load that Sethi's argument puts on it.[28]

References

McDowell, J. 1982. Criteria, defeasibility, and knowledge. *Proceedings of the British Academy* 68: 455–479.

McDowell, J. 2006. The disjunctive conception of experience as material for a transcendental argument. In *Disjunctivism: Perception, Action, Knowledge*, ed. F. Macpherson, and A. Haddock, 376-389. Oxford University Press.

Neta, R. 2019a. The basing relation. *Philosophical Review* 128 (2): 179–217.

Neta, R. 2019b. Disjunctivism and credence. In *New essays in epistemological disjunctivism*, ed. C. Doyle, J. Milburn, and D. Pritchard, 241–258. Routledge.

Neta, R., and D. Pritchard. 2007. McDowell and the new evil genius. *Philosophy and Phenomenological Research* 74 (2): 381–396.

Pritchard, D. 2012. *Epistemological disjunctivism.* Oxford University Press.

Putnam, H. 1981. Brains in a vat. In *Reason, truth, and history*, ed. H. Putnam, 1–21. Cambridge University Press.

Wright, C. 2002. (Anti-)sceptics: Simple and subtle: G. E. Moore and John McDowell. *Philosophy and Phenomenological Research* 65 (2): 330–348.

Wright, C. 2004a. Wittgensteinian certainties. In *Wittgenstein and scepticism*, ed. D. McManus, 22–55. Routledge.

Wright, C. 2004b. Warrant for nothing (and foundations for free?). *Aristotelian Society Supplementary* 78 (1): 167–212.

Wright, C. 2006. Comment on John McDowell's 'The disjunctive conception of experience as material for a transcendental argument.' In *Disjunctivism: Perception, action, knowledge*, ed. A. Haddock and F. Macpherson, 390–404. Oxford University Press.

Wright, C. 2014. On epistemic entitlement (II): Welfare state epistemology. In *Scepticism and perceptual justification*, ed. D. Dodd and E. Zardini, 213–247. Oxford University Press.

Wright, C. 2023. Closure and transmission again. In *New perspectives on epistemic closure*, ed. D. Pritchard and M. Nope, 163–189. Routledge.

[28] Special thanks to Paul Boghossian, Adrian Haddock and Alan Millar for very perceptive and helpful comments on an earlier draft of these Replies.