Synthese (2024) 203:166
https://doi.org/10.1007/s11229-024-04587-y

ORIGINAL PAPER

Factivity, hallucination, and justification

Jack C. Lyons[1]

Received: 3 May 2023 / Accepted: 31 March 2024 / Published online: 10 May 2024

Abstract
Veridically perceiving puts us in a better epistemic position than, say, hallucinating does, at least in that veridical perception affords knowledge of our environment in a way that hallucination does not. But is there any *further* epistemic advantage? Some authors have recently argued that veridical perception provides a superior epistemic benefit over hallucination not just concerning knowledge, but concerning justification as well. This contrasts with a traditional view according to which experience provides justification irrespective of whether it's veridical or hallucinatory. I think both views are mistaken. Although this traditional view should be rejected in favor of one on which *some* hallucinations are epistemically worse than veridical perceptions (and some are not), I don't believe there is good reason to think that the mere fact of hallucination—or factivity more generally—has any consequences for justification. Susanna Schellenberg has endorsed both the traditional and the factive views (for different elements or kinds of perceptual justification), and I critique her views in detail, though I also draw out more general epistemological lessons about factivity and evidence.

Keywords Perception · Hallucination · Disjunctivism · Epistemological disjunctivism · Evidence · Reliabilism

Surely there's at least one respect in which the veridical perceiver enjoys an epistemic advantage over someone who's undergoing a qualitatively identical hallucination: the perceiver, though not the hallucinator, is in a position to know something about the external world in virtue of that perceptual event. This much is uncontroversial. But is this the *only* epistemic advantage, or does veridical perception confer some further benefit over hallucination, beyond this? In particular, could the veridical perceiver also thereby have a higher degree or superior grade of epistemic justification than the hallucinator?

✉ Jack C. Lyons
jack.lyons@glasgow.ac.uk

[1] University of Glasgow, Glasgow, UK

A couple of decades ago, few epistemologists would have seriously considered the possibility that there might be a *justificatory* difference between a convincing hallucination and veridical perception. In recent years, however, several authors have endorsed a view that's sometimes known as "epistemological disjunctivism" (in reference to some similarities between it and the (metaphysical) disjunctivist view associated with naive realism; I won't be addressing the metaphysical view here at all), which holds that the perceptual beliefs of hallucinating subjects thereby have less justification (perhaps none at all) than those of their phenomenally identical veridically perceiving counterparts. Like the epistemological disjunctivist, I reject the traditionally dominant view that hallucination cannot make a justificatory difference; but not for the disjunctivist's reasons (Lyons, 2016b, 2019). I agree with the traditionalist that factivity per se has no consequences for justification, so the mere fact that I'm hallucinating doesn't undermine or degrade my justification. Hallucinations might nevertheless sometimes differ from veridical perceptions in other, non-introspectable but epistemically relevant respects. This has broader implications for our understanding of the relation between factivity and justification, and broader yet for our understanding of the relation between evidence and justification.

In a series of recent work, Susanna Schellenberg (2013, 2014, 2016, 2018, forthcoming) has defended both of the views I want to reject. She divides perceptual justification into what she calls "phenomenal evidence" and "factive evidence"; she holds the traditional view to be correct about the former and the disjunctivist view to be correct about the latter. I will address her work in detail, but with an eye to larger epistemological issues, especially the justificatory significance of factivity and the role of evidence in epistemology. I think she's mistaken about these, but I think a lot of other epistemologists are as well, and I want to draw out the larger lessons in later sections of the paper.

Some terminological matters before I begin. I'll go with the standard stipulation: a hallucination is an episode that's phenomenally or introspectively just like a veridical perception, except that nothing is actually perceived in virtue of that episode; illusion involves the perception of a real object but a *mis*perception of one or more of its properties.[1] (On a view according to which perceptual states are essentially representational and involve a demonstrative or directly referring element, hallucinations involve the failure of this element to refer, while illusion involves erroneous predication or attribution of properties to the object referred to.) All ensuing talk about hallucination will intend "perfectly matching" hallucination, hallucination that is introspectively indistinguishable from a veridical perception. Without this assumption, claims about justificational parity between hallucinating and genuinely perceiving aren't especially plausible and denials of parity aren't especially interesting. I will use 'perceptual belief' in a way that is neutral between good (veridical perception) and bad (hallucination) cases, even though there's a sense in which it's only in the good case that the beliefs count as genuinely *perceptual*. Similarly, I'll refer to both kinds of experience as "perceptual experience".

The kinds of case that I'm interested in are cases of ordinary, spontaneous and fairly unreflective belief. We are capable of adopting a hyper-reflective, quasi-Cartesian

[1] For a more complex view of hallucination and illusion, see Macpherson and Batty (2016).

stance in which we attend to our experience, form an explicit introspective belief about it and from that draw an inference about the external environment. This is not the kind of case I'm interested in here, for in this kind of case I think even the disjunctivist will concede that the justificatory status of the conclusion is the same, in good and bad cases, at least to whatever extent the "perceptual" justification derives from this inference.

I'm aware that 'disjunctivism' is perhaps not the best term for the view in question, especially since some versions of the view hold that veridical perception offers additional kinds of justification over hallucination.[2] Nevertheless, even these views hold that the kind of justification had in hallucination and the kind exclusive to veridical perception are distinct, disjoint, kinds of justification. Anyway, I take this to be fairly standard terminology.

The central question here is whether the hallucinating subject is in as good a justificatory position as the (introspectively identical) veridically perceiving subject. It is common to simply rephrase this question as the question whether a hallucination *experience* offers the same (or as good) *evidence* as does a veridical perceptual experience. I will be resisting this rephrasing. As I'll explain shortly (in Sect. 1), these emphases on experience and on evidence would narrow our focus in undesirable ways. Additionally, for reasons to be seen below—some connected to perception and hallucination, but some much more general—we should not be thinking of justification in terms of evidence, but in terms of the process by which the belief is formed or maintained. I'll explain this in due course, in Sect. 4 below.

Finally, it *almost* goes without saying that I will be assuming throughout that justification is not coextensive with knowledge, in particular that it's possible to have a justified belief that isn't known. This was probably obvious from the framing of the initial question but still worth making explicit. I'm also assuming (slightly) further that it's possible to have justified but false beliefs. Though a few authors have recently proposed identifying knowledge and justification (though perhaps with a specialized sense of 'justification' in mind), I don't know anyone who allows that Gettiered beliefs can be justified, but only if they're true. Probably someone does or soon will, but I'm ignoring them here.

1 Traditionalist, disjunctivist, and mixed theories

Once more, our central question is this: is the hallucinating subject in as good a *justificatory* position as the phenomenally identical but veridically perceiving subject?

The traditional answer to this question is staunchly affirmative: the fact that one is hallucinating makes no difference to the agent's justification, unless, of course, the agent is aware of the fact or it makes a difference to their experience. Again, until quite recently, no epistemologist would have seriously considered any different answer. After all, justification is subjective in a way that knowledge is not, and everything

[2] Schellenberg, who I am classing as a disjunctivist, explicitly disavows the label. But it is clear that she means something different by the term than I do. She thinks, for example, that rejecting access internalism is sufficient for rejecting epistemological disjunctivism (2018, p. 183); I would assume that disjunctivism comes in non-access-internalist forms.

seems to the hallucinator exactly as it does to the veridical perceiver. Even if we don't think that justificatory status is fixed by introspectable features—and many of us do not—it still seems fairly intuitive that what the hallucinating subject ought to believe—provided of course that the hallucination is perfectly resembling veridical perception and the subject has no reason to suspect anything is amiss—is what their senses are telling them.

The rise of externalist theories of justification in the late 1970s (e.g., Goldman, 1979; Sosa, 1980) undermined a central rationale for the traditional view, by rejecting the claim that justification is fixed by introspectable features. And yet, the first epistemologists to reject the traditional view were not reliabilists, but epistemological disjunctivists. That is, rejection of the traditional view was grounded not in the difference in utilized process, capacity, skill, virtue, or the like, between the good and bad case; but in an alleged difference between the evidence, ground, reason, "criterion," or something similar, between the good and bad case—namely that the agent has a sound reason (/ground/evidence) in the good case but one that is in some way defective in the bad case. (I will return to this important distinction between, as I call them, process epistemology and indicator epistemology below, in Sect. 4) The locus classicus of this disjunctivist response to the traditional view is McDowell (1982), though I won't try to unpack his view in any detail, partly because McDowell is notoriously obscure, but also because his explicit concern is with knowledge rather than justification. The view in its general form holds that an agent is in a better justificatory position in the good case than in the bad case, simply in virtue of the factivity difference between the two cases. Seeing that p provides more or better justification for p than merely seeming to see that p, precisely because the former guarantees the truth of p and the latter does not. Several authors (e.g., Byrne, 2014; Hellie, 2011; Littlejohn, 2012; Millar, 2008; Pritchard, 2012; Schellenberg, 2013; Williamson, 2000) endorse some version or some close relative of this view.

One option for the factivist is to say that there's no justification whatsoever in the bad (i.e., hallucinatory) case (Byrne, 2014, 2016). Another is to say that there's still justification in the bad case but that it's somehow inferior to the justification had in the good case. "Inferior" justification in the bad case could then be understood in one of two ways: either simply as the view that there's a single quantity—justification—and there's less of it in the bad case and more in the good case; or as the view that there are two distinct *kinds* or *dimensions* of justification: an inferior or less important kind had in at least the bad case, and a superior kind had only in the good case. It's not always easy to tell who is taking which option, but Millar (2019) seems to take the distinct kinds approach.[3] Williamson (2000) and Schellenberg (2013, 2018, forthcoming), seem to take the single quantity approach; in their very different ways, each holds that the agent in the good case has all the evidence had by the agent in the bad case but also some additional evidence, and this additional evidence seems to be of a common

[3] Millar (2011, 2019) distinguishes "reasonableness" from "well-foundedness" and seems to take a sufficiently dim view of the former—which is the only epistemically redeeming feature had in the bad case—that there's not much difference between his view and one where there's simply no justification in the bad case: "reasonableness is but a shadow of what we are really after when we wish our beliefs to be well founded and in that sense justified" (2011, p. 345).

currency with that shared evidence, such that it makes sense to add them up, as one normally does with bits of evidence.[4]

These discussions sometimes presuppose a view I call "experientialism" (Lyons, 2009), which holds that perceptual beliefs derive their justification from perceptual experiences, which serve as evidence for those beliefs, in something like (but also perhaps something unlike) the way premises serve as evidence for conclusions. With that presupposition in place, we could state traditionalism and factivism provisionally as follows:

> **Trad$_{exp}$:** the justificatory force of an experience is determined only by introspectable features of the experience: there is no justificatory difference between veridical perception and perfectly matching hallucination.
> **Disj$_{exp}$:** the justificatory force of an experience is determined partly by the factivity of the experience: veridical experiences provide better justification than hallucinations.

These provisional, experientialist formulations are convenient, in that they clearly highlight a crucial contrast. But I don't want to assume experientialism here. This is for three reasons. First, I think experientialism is false (Lyons, 2009, 2016a). Even if I'm wrong about this, it's a problematic view for which serious arguments are rarely if ever given, against which significant prima facie objections exist; and it certainly should not be assumed as an obvious and innocuous starting point (Lyons, 2020). Second, not all proponents of the traditional or disjunctive views themselves embrace experientialism. Classical coherentists (BonJour, 1985; Davidson, 1986; Lehrer, 1990), for example, would all endorse the traditionalist claim that factivity makes no justificatory difference, but they probably wouldn't spell that out in terms of evidence and certainly not in terms of (nondoxastic) experience, as coherence on their views is only a relation among beliefs. And although the disjunctivist *can* think of good perceptual beliefs as distinct from and based on veridical experiences, another option is to think of the agent's seeing that p as *constituting* the agent's knowledge (and thus belief) that p, rather than serving as anything like a premise for it (Williamson, 2000). An experientialist framework won't illuminate such a view. Third, there's no room in an experientialist framing to articulate my own preferred answer to our central question about hallucination and justification, which is neither traditionalist nor disjunctivist, but a third, "mixed" view, according to which neither factivity nor intrinsic character of experience matter per se for justification; instead, what matters is whether the perceptual belief has the proper etiology (details to follow shortly). Some proponents of a mixed view might mention experiences when unpacking "proper" etiology, and some won't; but for none of them will the nature of the experience *exhaust* the proper etiology. In fact, in the last section of this paper I will say considerably more about the very general process framework for epistemology, which at the same time both redirects our focus away from experiential evidence and allows for the possibility of the mixed view that's neither traditionalist nor disjunctivist.

[4] In some sense, obviously, Schellenberg's "factive evidence" is a different kind of evidence than "phenomenal evidence". Similarly, however, perceptual evidence is a different kind than testimonial evidence in some sense; but they're also of the same kind in the more important sense that they are summative.

As mentioned, the earliest rejections of the traditionalist view came from disjunctivists. A very different way of rejecting the traditional view emerges if we think of justification as the result of exercising a certain kind of skill, or capacity, or virtue, or of utilizing a certain kind of cognitive *process*. Just to have a concrete example of this general approach: a (simplified) process reliabilism holds that a belief is justified just in case it is the output of a cognitive process that has a propensity to produce a high ratio of true to false beliefs (in, let's say, the environment in which the agent is employing that process). On such a view, and on the assumption that hallucination at least sometimes involves a different cognitive process than veridical perception (some hallucinations have endogenous origins, of a sort that veridical perceptual beliefs do not, for example), there will sometimes be justificatory differences between hallucination and veridical perception, even when introspectively identical. But this difference will be due to *general* features of the process/capacity/etc. employed, not to veridicality (or not) or referential success (or not) on this particular employment. To illustrate, we might think that Macbeth-style hallucinations (which originate endogenously, as the result of madness) have less justification than the errors in perceptual judgment that result from exogenous, brain-in-a-vat-style manipulation, even though the two types are equally non-factive.[5] The latter, in fact, might involve exactly the same cognitive processes as a normal case of veridical perception (Lyons, 2019), which are generally reliable, even though getting things wrong on this particular occasion. A victim of a one-off demonic intervention thus might be using a reliable process, though that same process would be unreliable were the agent to be persistently thus deceived. Thus, we might find justificatory differences among hallucinations (contra disjunctivism), and also between (some) hallucinations and (some) matching veridical perceptions (contra traditionalism).

And we might trace these justificatory differences to the fact that the relevant good-making feature, whatever we might take that to be—e.g., reliability (Goldman, 1979; Lyons, 2009), rationality (Siegel, 2016), proper function (Burge, 2003; Graham, 2012), is absent in the endogenous cases but present in normal veridical perception and at least potentially, in the exogenous hallucination cases. Alternatively we might think that certain epistemic vices are operative in the Macbeth case but not the brain-in-a-vat case, etc..[6] The operative differences between the "better" and "worse" hallucinations are obviously not a difference in factivity or successful reference, and they won't always include a difference in introspectable features.

So we have three views about the justificatory status of hallucination compared to veridical perception: a traditionalist view, a disjunctivist view, and a mixed view. Dropping the experientialist presupposition, temporarily adopted above, we can formulate the options as follows:

[5] For my own part, I'm not naturally inclined to think of exogenous, demon- or neuroscientist-induced, "hallucinations" as hallucinations, really, even though they clearly satisfy the standard definitions and are clearly nonfactive in the sense relevant to disjunctivism. When Andy Clark (2015) promotes the predictive coding view as the radical view that perception is "controlled hallucination," the rhetorical punch has to derive from the term being evocative of Macbeth, not of brains in vats.

[6] These examples show that one needn't be a process reliabilist to endorse this non-traditional, but non-factivity-based view about the justificatory status of hallucination. To the best of my knowledge, the first articulation of this mixed view in any detail is in Lyons 2016b (in improved form in Lyons 2019), although Burge (2003) has a footnote (n24, pp. 537–8) that could be read as stating or anticipating something similar.

Trad: The prima facie justificatory status of a perceptual belief is determined only by introspectable features[7]; a veridically perceiving agent and an introspectively identical hallucinating agent have exactly the same prima facie perceptual justification.

Disj: The prima facie justificatory status of a perceptual belief is determined in part by the factivity of the relevant states; everything else equal, a veridically perceiving agent has better prima facie perceptual justification than an introspectively identical hallucinating agent.

Mixed: The prima facie justificatory status of a perceptual belief is determined by the etiology of the belief (in a way that is not exhausted by which, if any, experiences are implicated in its production); some hallucinating agents will thus have better prima facie perceptual justification than their introspectively identical counterparts.

My main focus in this paper is on the disjunctivist view and why we should reject it. Although I won't try to defend the mixed view here, it should be clear that the arguments leveled against disjunctivism are not arguments that favor or presuppose a traditionalist view.

2 Reasons for disjunctivism

So why would anyone endorse epistemological disjunctivism? I guess someone might find it intuitive that every hallucinator has less justification than an otherwise matching veridical perceiver; but I suspect that disjunctivism gets most of its credibility from theoretical considerations, rather than intuitive pull. There are three main, not necessarily exclusive strands of argumentation. I can only discuss one of them in detail here.

The first traces back to McDowell (1982), who argues that the possibility of perceptual knowledge entails the possibility of a special kind of conclusive, indefeasible reason, possessed in the good case but not in the bad. A second traces back to Williamson (2000), who equates evidence with knowledge. Some followers have endorsed a slightly more inclusive conception of evidence, according to which only factive mental states can serve as evidence; and hallucinations are paradigmatically nonfactive. The third strand goes back to Millar (2008) but is developed in different directions and more detail in Schellenberg (2013, 2014, 2016, 2018, forthcoming), and it holds that although all perceptual judgment involves the exercise of a capacity, it is only in the good case that the capacity is being exercised in a non-defective way.

I will limit my attention to this third strand, and in particular, to Schellenberg's arguments and articulation of the view. Most of the first subsection (Sect. 2.1) is aimed directly at Schellenberg and doesn't apply to other views; much of the second

[7] Again, 'perceptual belief' is taken to be neutral between good and bad cases. The restriction to prima facie justification is there because I want to allow a traditionalist to say, for example, that whether you have defeaters right now isn't always introspectively available to you. Many traditionalists won't say that, but I want to allow it, because it's an unrelated issue and I don't want it to serve as a distraction.

half of this section (in Sect. 2.2) will apply to Millar's and similar views, though I won't try to press the particulars. In the subsequent section (Sect. 3), I'll offer general objections to disjunctivism of any form and however motivated, and the rest of the paper is concerned with the more general view.

The basic idea in Schellenberg is that every perceptual event involves the exercise of a capacity that's systematically linked to the truths it would pick out if everything were going well—hence providing some degree of justification. If everything *is* going well—i.e., the capacity is employed in a good case—there's an *ideal* link, hence a greater degree of justification. Putting this general idea in terms more at home in Millar (2011, 2019) than in Schellenberg: hallucination involves a *bad* exercise of a *good* capacity. Because the cognitive capacity being employed is an otherwise and generally good one, the result is some positive degree of justification. But when the particular exercise is a bad one, it must be a lower degree of justification than would obtain were it a *good* exercise of a good capacity. This has some intuitive appeal, but of course, we'll need to know more about the good and the bad, or returning to Schellenberg, the ideal and the merely systematic linkage, to evaluate the proposal.

Schellenberg claims that perceptual justification in general derives from the fact that a percept is the outcome of the exercise of a capacity whose function is to pick out distal features of the environment. A perceptual experience as of a red square, for example, is the result of the exercise of a capacity the function of which is to pick out red squares—that is (and here's how I'll be using 'pick out':) both to mean *red square* (i.e., to have red squares as the satisfaction or veridicality condition) and to be tokened veridically *of* some red square.[8] The content, and therefore the phenomenology, of an experience is determined by which capacities were exercised on that occasion. Capacities on this view are not necessarily skills or virtues; they needn't be reliable or indeed even ever successful in their representational aspirations. Schellenberg thinks—for reasons I'll try to understand shortly—that because an experience is thus grounded in capacities, it's always prima facie rational and thus in some sense epistemically proper to heed the experience, i.e., to believe in conformity with it, e.g., that there's a red square. When we do so, we have what Schellenberg calls "phenomenal evidence". This is present in both the good and bad cases. The good and bad cases are also importantly different, however. The bad case involves a defective exercise of this capacity, while the good case involves a successful exercise; thus we have an additional form of evidence in the good case that's absent in the bad case. She calls this "factive evidence" (e.g., Schellenberg, 2013, 2016, 2018, forthcoming). Thus the traditionalist is right to think there's perceptual justification in the good and bad cases, but the disjunctivist is right to think there's *more* justification in the good case, even when absolutely everything else is the same.

This, as I say, is the basic idea: the exercise of a capacity generates a certain amount of normativity; but a greater normative status accrues to successful exercises than to defective exercises.

[8] She uses 'single out' instead of 'pick out' here, although her term has, to my ear, discriminatory connotations that she doesn't intend. Both are used as success terms in the sense that the function is not merely to *mean* "red square" but to indicate red squares by firing or being tokened (only) *in the presence of* a red square. In hallucination, the perceptual state fails to "single out" (or to "pick out") a red square, though it continues to *mean* "red square".

2.1 Schellenberg on "phenomenal evidence"

So how does the first part of this work? That is, how does the fact that something is an exercise of a capacity imply that it's in any sense worthy of our belief? If you think of capacities as skills, or reliable tendencies to represent accurately, there's no mystery here, for if that's what capacities are, then if your percept that p is the result of the exercise of a capacity, it is thereby objectively likely that p is true. This is Millar's understanding of capacities, but it is patently *not* Schellenberg's understanding; she explicitly rejects reliability as in any way explaining the justificatory power of experience (e.g., 2018, p. 177). For Schellenberg, to have a certain perceptual capacity is simply to be such that you could possibly pick out, say, F, if you were confronted with F in the right conditions.[9] Capacity talk *looks* initially promising for her, but I think this is because capacity talk has such strong reliabilist associations and because reliability so plausibly (though admittedly, controversially) has something centrally to do with justification. Once those implicatures are cancelled, there's no obvious route from capacity to justification.

Her official argument for phenomenal evidence invokes a "systematic linkage" between perceptual states and truth—or at least, the truths that those perceptual states represent in the good case (Schellenberg, 2013, 2016, 2018, forthcoming). She says, "if phenomenal states are systematically linked to what they are of in the good case in the sense specified, then it is rational to heed their testimony" (2018, p. 177). Again, this sounds quite promising initially, but less so once we unpack what she means by 'systematic link'. Clearly there are many kinds of systematic linkages and clearly many of them don't make heeding rational. Every representation is in some sense systematically linked to what it represents (and also in some sense linked to its negation), but we don't want to conclude that every representation generates justification and evidence (and certainly not justification for believing its negation!). Unjustified beliefs are systematically linked to their contents, but it's not rational to heed unjustified beliefs. Rather, it must be some specific *kind* of systematic link that generates the justificatory force. She says that "[p]henomenal states are systematically linked to what they are of in the good case *in the sense that* the perceptual capacities employed in the bad case are explanatorily and metaphysically parasitic on their employment in the good case" (2018, p. 175, italics added; cf. pp. 174–9; 2013, p. 714; 2016, p. 880; forthcoming; etc.). The metaphysical and explanatory priority just noted consist in the fact that any analysis of a bad-case exercise of a capacity must appeal to the good-case exercise; what makes this a perceptual experience as of a red square is that it is an exercise of the same capacity that would successfully represent a red square as such in the good case.

But even granting all of this, it is far from obvious how this suggests any positive epistemic status in either the good case or the bad. Let C be a capacity that might possibly, occasionally, represent something correctly when conditions are favorable. On the fluke occasion that it does represent correctly (the good case), how does this

[9] In fact, her possession conditions for capacities, especially as articulated in Schellenberg (forthcoming), imply that normal humans likely *currently* have electroreception capacities—maybe even clairvoyance capacities—because *if* we had the relevant sense organs, we'd have the relevant representational abilities (see Lyons forthcoming for details). This is certainly not a reliabilist or skill-based understanding of capacity.

imply that the agent who employed the capacity has any degree of justification? And even if C does produce justification in *some* conditions of employment, how does the fact that *other* conditions of employment can only be understood in terms of those good ones argue that these other conditions are therefore also conducive to justification? It's one thing to argue that if *any* use of a capacity yields (prima facie) justification, then *every* use of that capacity does; but this clearly isn't her argument or she'd have said so. Instead her argument seems to be that if E-employments of a given capacity yield prima facie justification, then if F-employments can only be understood by reference to E-employments, then F-employments of that capacity also yield prima facie justification. I'm not seeing any reason at all to think this is true, or indeed any reason to think that epistemic value accrues in virtue of being metaphysically *secondary* to something successful. It's certainly not *generally* the case that whenever one property, capacity, disposition, or whatever is explanatorily and metaphysically parasitic on a good property, capacity, etc., the parasitic or secondary one is also therefore good. Failure can typically only be understood by reference to success. The shortcomings of my junker car cannot be understood or analyzed except by reference to what a good car would be like, but this surely doesn't imply that my junker is in any sense good, despite the systematic linkage.

When discussing the metaphysical priority of the good case, Schellenberg frequently also mentions (teleo-)functions. It's not just that the capacities do represent in the good case; they *function* to represent what they succeed, in the good case, in representing. She invokes functions to underwrite her claims about metaphysical priority,[10] but teleofunctions are sometimes thought to be sources of normativity in their own right, so it's worth exploring the possibility that these functions might provide the needed link between capacities and justification. (To be fair, I don't think this is the role Schellenberg intends functions to play, but I discuss it here in the interest of thoroughness.) I see three problems with this approach, however: (a) these functions don't obviously generate the right kind of normativity, (b) they don't attach to the right state, and (c) appeal to functions wouldn't get the desired outcome anyway.

As for (a), any thing that does what it's designed to do is in some way good, but it doesn't yet follow that it's *epistemically* good. Belief-relevant capacities—think wishful thinking mechanisms or hypersensitive predator detectors—might be functioning perfectly well while still generating epistemic *dis*value. In any case, an argument from teleofunction to epistemic value would need to be spelled out.

[10] In fact, I see a significant tension between the role that functions seem to play in her view and the claim about the metaphysical priority of the good case. Suppose that *what makes* something an F-detector is that its function is to pick out F-instances. In the good case it does so; in the bad case it "says" F even though there's nothing F nearby. To the extent that function is determining contents, function is determining success-conditions and thus determining which cases are good ones and which bad; and it's *both good and bad cases* that are parasitic on functions, with neither *case* nontrivially prior to or parasitic on the other. If we had some way of individuating capacities antecedently to or independently of the contents involved (I don't think Schellenberg does; see Lyons (forthcoming)), then we could use the performance of these capacities in the good cases to fix contents and thus make sense of (mis)representational content in the bad cases. But then functions wouldn't be playing any significant role. Either the metaphysical and explanatory primacy can be understood independently of functions, in which case functions are idle; or functions are doing real work for the theory, in which case, good and bad cases are on the same level and claims about primacy are undermined.

More importantly (b), the normativity wouldn't attach to the state that Schellenberg wants it attached to. If perceptual experience has a function, then that function could underwrite the ascription of evaluative assessments (along some—possibly epistemic—dimension) *to those experiences*. But Schellenberg's point is not that some of those experiences are proper (and some, presumably, improper); it's that a *belief* that heeds, i.e., conforms to, *any* perceptual experience is thereby proper. Maybe it's plausible that perceptual belief is justified when it's based on an epistemically good experience (but not when it's based on a bad one),[11] but that's quite clearly not her view, especially since, if there's any sense to be made of "good" and "bad" experiences, this is going to figure into factive evidence (which differs between the good case and the bad case), while the current concern is phenomenal evidence (which is the same in the good and bad case). Why, then, should we care whether perceptual *experiences* have functions? What we want to know is whether perceptual *beliefs* have functions. If she wanted functions to explain why the belief is proper, she would need a view where it's beliefs, or *doxastic* capacities, and not just experiences, that have functions. But that's a view much more like Burge (2003, 2020) or Graham (2012, 2024, etc.) or Bergman (2006) than like Schellenberg.

Finally, and relatedly (c), it doesn't look like the function she describes could do anything at all to explain what the agent is doing right or well even in the bad case. It's not *having* a function that generates positive evaluative status; all hearts have a function, but this doesn't mean that all hearts are good. Rather, it's *fulfilling* a function or functioning *properly* that generates positive evaluative status. But (a point we'll return to below), the function of perceptual capacities is unfulfilled in the bad case. And she's given us no reason—nor does one leap to mind—why every exercise of a perceptual capacity (even all of the failures) is an instance of proper function. Is it somehow impossible for perceptual capacities to malfunction?

Thus it remains completely mysterious why, on a view like Schellenberg's, there's anything positive to say about the agent in the bad case. But saying something positive about the agent in the bad case was precisely what the construct of phenomenal evidence was supposed to enable.

Why, then, is it in any sense epistemically proper to "heed" our percepts? I think the only idea remaining is that it's rational to heed any state that is supposed to say that p only when p (i.e., the function is: don't represent p as obtaining unless p really obtains). But this just cannot be right. If any states aim at truth in this way, beliefs do, in the sense that every belief is supposed to say that p only if p. But if it's rational to heed all such states, then it's rational to heed the unjustified ones. That is, we should believe even the unjustified beliefs. I don't see much difference between this and the claim that all beliefs are justified.[12]

[11] Some (e.g., Siegel 2016) do hold such a view. I think there are good reasons to resist this (Lyons 2016c).

[12] Schellenberg addresses an objection similar to this (2018, p. 178; cf 2013, pp. 720–1; 2014, p. 97, etc.) (although, surprisingly, she seems to think it's irrational to treat *any* beliefs as evidence; I am only objecting to treating *unjustified* beliefs as evidence). Her response seems to involve a rejection of the functional principle that it's ok to heed anything that's supposed to only yield truths and a re-centering on the principle that justification results from explanatory and metaphysical primacy of the good case. I confess I'm unable to see how this answers the challenge. If some representational contents are grounded in perception (and she concedes that any belief about particulars probably is), then won't *any* application

Above, I suggested that Schellenberg's view was almost but not quite the view that hallucination involved the bad exercise of a good capacity. I tried to make sense of goodness of capacity in terms of systematic linkage and then in terms of functions, but concluded that neither works in her framework. In the end, the suggested view couldn't have been quite hers anyway, because she doesn't want a distinction between good and bad capacities, but only between good and bad exercises. Whenever any perceptual capacity is employed—even if it's a diseased, malfunctioning capacity, one that never actually gets things right, is employed irrationally, etc.—it generates phenomenal evidence, on Schellenberg's view. On this count, she stands and falls with the traditionalist.

To a reader with some kind of reliabilist or proper functionalist sympathies, her language—dense with terms like 'capacity', 'systematic link', and 'function'[13]—makes it look like she's a friend. She's not. The first two terms are explicitly stripped of any connection to reliability, objective probability, and the like; the third is never explicitly linked to normative status, and given where and how it appears in her theory, looks like it couldn't be.

In a similar vein, the term 'phenomenal evidence' makes her view look more similar to standard internalist epistemologies than it really is and presumably allows her view to benefit from whatever intuitive plausibility those standard views enjoy. She's happy to encourage this association: "[i]n contrast to, say, orthodox versions of reliabilism, [her view] makes room for the cognitive and epistemic role that conscious mental states play in our lives" (2018, p. 3). "To a first approximation, we can understand phenomenal evidence as determined by how our environment sensorily seems to us when we are experiencing" (2018, p. 167). Even though I'm a fairly orthodox reliabilist, I recognize the intuitive pull of the idea that what you should perceptually believe is fixed—or at least influenced—by your conscious experiences. Reliabilism denies this, however, on the very good grounds that justification is fundamentally a matter of believing what's objectively likely to be true; so to the extent that experiences do not conduce to reliability, they are not to be heeded; to the extent that they are conducive to reliability, they are to be heeded, but only in virtue of the reliability boost, not because of anything to do with phenomenology, or conscious accessibility, or the rest.

Schellenberg's view is *exactly* like orthodox reliabilism in this latter respect. What gives perceptual states their epistemic potency is that they're "systematically linked" to what they're of in the good case (i.e., that these states result from a capacity that can only be understood by reference to its underwriting an ability to "discriminate and single out" (i.e., to segregate from the background and refer to) particulars in good cases). But none of this has anything to do with conscious experience. Zombies and blindsighters can employ capacities to discriminate and single out Fs both in the presence and in the absence of Fs. The same could hold of clairvoyants if we describe them in the right way. And these employments can exhibit the same systematic linkages as ours. There is thus nothing obviously *phenomenal* about "phenomenal evidence".

Footnote 12 continued
of those contents—including as constituents of unjustified beliefs—be explanatorily and metaphysically parasitic on the good case? Primacy and parasitism are transitive, after all. If so, any tokening of such a content would yield phenomenal evidence, which is clearly implausible.

[13] I haven't mentioned it here, but 'discrimination' gets heavy rotation in all these works as well.

Nor, for that matter, is there anything obviously *evidential* about it, since (or so we can suppose) blindsighters and zombies are simply *trusting* in their perceptual beliefs, rather than *basing* those beliefs on anything—they're certainly not basing them on conscious experiences. On such a picture, it's hard to see how *evidence* might fit in. And yet, again, the systematic linkage, which is doing all the theoretical work for Schellenberg here, is the same for them as it is for us. Indeed, to downplay these facts, she has to staple a conscious experience clause onto her formulation of the phenomenal evidence principle:

> If S perceives α or suffers an illusion or a hallucination as of α (while not suffering from blindsight or any other form of unconscious perception), then S is in a phenomenal state that provides phenomenal evidence for the presence of α. (2018, p. 17; cf 2013, p, 704)

But an agent who *is* suffering from blindsight is still exercising if not the very same capacity, a capacity that shares the only features relevant to its justificatory power, namely the "systematic linkage" to the good case. If "phenomenal evidence" is some real thing that normal perceivers get from perception, zombies and blindsighters get it too, although for that reason it's better not to *call* it "phenomenal evidence." Thus, the view just quoted could be rendered more simply (and more generally) as:

If S perceives α or suffers an illusion or a hallucination as of α, then S has some prima facie justification for the presence of α.

Any argument Schellenberg has for her own, internalist-friendly principle, is an equally good argument for this principle, though this one is not internalist-friendly.

Summarizing, Schellenberg holds that any belief where we take experience at face value is going to have a positive prima facie epistemic status of a particular sort. The official argument for this was that it follows from the fact that the erroneous experiences are the result of a capacity that can't be understood except by reference to the possible cases where it gets things right. I was unable to see how this established any kind of epistemic potency, let alone across all employments. She might have an argument up her sleeve that trades on functions to represent, but it's unlikely that functions can do the needed epistemic work, given the role she's already assigned them. Nor should anyone think that internalist considerations speak in favor of the view, since the internalist-friendly elements—phenomenality and evidence—aren't really features of the *operative* parts of the view (the stuff about capacities). For these reasons, we should not be moved by Schellenberg's traditionalist argument for thinking that anything has gone epistemically right in bad cases.

2.2 Schellenberg on "factive evidence"

But what about her disjunctivist argument that things have gone comparatively better in the good case, insofar as they've gone worse in the bad case? That would be enough to support some kind of epistemological disjunctivism. And it might be one that is of interest to other disjunctivists.

The central idea here is that hallucination involves a bad, a "defective" (e.g., 2018, p. 183) exercise of a (good) capacity.[14] Because it's defective in the bad case but nondefective in the good case, the latter enjoys enjoys a distinct type of justification (involving "factive evidence") and a greater overall quantity of evidence.[15] There's obviously an initial plausibility to the idea that a defective, botched, or otherwise poor exercise of a certain skill or capacity yields a lower degree of justification than a good exercise. But ultimately the plausibility will depend on how and in what sense the exercise is defective.[16]

First of all, any exercise of a perceptual capacity in the bad case will be ipso facto defective in the uninteresting sense that it's outputting a falsehood (setting aside the controversial idea of "veridical hallucination"). Clearly, however, if this is the only sense in which the exercise is defective, it gives us no reason to endorse the disjunctivist claim that there's less justification in the bad case, since we're happy (as I stipulated early on) with the idea of justified falsehoods. We might have very strong inductive evidence for p, even though p is false; inferring p in that case is not a defective exercise of an inferential capacity, at least not in a sense relevant to justification, not in the sense that hasty generalization is, for example. It is unsuccessful, but not defective. Our concern here is with justification, not knowledge; if defective employment is supposed to entail reduced justification, it will have to be defective in some *further* sense than merely generating a false belief.

With that out of the way, what reason might there be to think that bad-case employments of perceptual capacities are relevantly defective, defective in this further way? I find three arguments in Schellenberg, two explicit and one implicit. One explicit argument expands on the idea of a systematic link; the other invokes gappy contents; the implicit argument appeals again to functions.

The first argument states that whereas there's a systematic link between perceptual states and worldly facts in all cases, there's an *ideal* link in the good cases (2018, p. 181; cf 2013, p. 735; 2016, p. 881). Thus, while perceptual experience always yields justification, it yields even more justification in the good cases. One problem

[14] It's actually more Millar (2008, 2011, 2019) and McDowell (2011, 2018) who think of the *exercise* of the capacity as defective; Schellenberg (2013, 2016, 2018) always describes the *contents* as defective. I'll use the former phrasing, as it seems a more natural view.

[15] Schellenberg is quite explicit that factive and phenomenal evidence are coins of a single currency in the sense that adding factive to phenomenal evidence gives you *more* of the relevant quantity, rather than a separate, different kind of thing. This isn't the only option for disjunctivists, but I'll follow Schellenberg in supposing that the putative justificatory difference between the good and bad cases is a difference in amount, not in kind. She does also sometimes use 'factive evidence' in the different, trivial sense of evidence-that-happens-to-be-true or maybe evidence-that-is-not-misleading (see, e.g., 2018, p. 182). She even argues that because (veridical) perception is factive, the evidence that it provides inherits factive status from that (2013, 2018). This clearly only speaks to factive evidence in a trivial sense of the term, not the substantive sense of *additional* evidence that accrues in the good case, over what's present in the bad case. This latter, substantive sense—a sense that entails *more*— is what she normally means by 'factive evidence' and is the only thing I'll ever mean by the term. For a nuanced discussion of various senses of factivity, see Graham and Pederson (2020).

[16] In Schellenberg (2013, forthcoming), the capacity is said to be employed "baselessly", but that's just a technical term introduced to denote cases of reference failure, not to indicate any antecedently understood defect in terms of which reference failure is explained.

with this we've already seen, which is that there's no reason to believe that the systematic link—at least of the sort Schellenberg says is doing the work—really does generate any epistemic force, thus no reason to think it generates *even more* force in the "ideal" case. Second, although the idea of an ideal, a "stronger" (2013, p. 735; 2016, p. 881), a more-than-merely-systematic link makes perfectly good sense if we're understanding systematic linkage in broadly reliabilist ways, it becomes quite obscure when we remember that the link in question, between the perceptual state and a state of affairs that would make it true, is simply that the capacities actually employed are explanatorily and metaphysically parasitic on the employment of those capacities in the good case. This isn't something that obviously does come in degrees or admit of superlative, ideal, varieties. The argument seems to be that the good-case employment gets a justificatory boost, because *other* employments are metaphysically and explanatory parasitic on it. As with the argument for phenomenal evidence, I find myself mystified as to why any epistemic consequences should ensue, if this is what's meant by 'systematic link'.

There's a more promising argument in the vicinity, I think, although it can't be hers and is still problematic anyway. Consider the following:

> It's good to give money to the poor. Because of that, it's good to *try* (but fail) to give money to the poor (though not as good, of course, as actually succeeding). Thus, it's plausibly at least sometimes the case that, if Φ-ing is valuable in some way, then "attempted" Φ-ing, or "botched" Φ-ing, is similarly valuable, though in lesser degree. Maybe perception is analogous. Maybe veridical perception is epistemically good, and because of this, non-veridical perceptual belief formation is a derivative, but lesser, epistemic good.

I think this is an improvement, but it doesn't salvage an argument for factive evidence. First of all, this surely isn't really Schellenberg's view; it's more Williamsonian than Schellenbergian.[17] Second, even if we find the conclusion compelling, it doesn't yet follow that this is a difference in justification; it might simply be that the distinction here marks the knowledge/justification distinction, rather than two grades of justification. We'd need further argument. Third, although there are surely instances where a failed, attempted Φ shares some of the (positive or negative) value of a successful Φ (charity and murder come to mind), there are surely counterinstances as well. Though there's something epistemically good about proving a theorem, it doesn't follow that there's anything epistemically good about failing to prove a theorem. Thus, we'd again need further argument, this time to show that perception is more like murder than like proving a theorem. Finally, an important part of Schellenberg's project at this stage is to explain *why* there's something epistemically good about *veridical* perception.

[17] If this were her view or similar to it, she'd likely present factive and phenomenal evidence in a different order; but she uniformly (2013, 2016, 2018, forthcoming) argues for phenomenal evidence *before* arguing for factive evidence. "The analysis of the epistemic role of phenomenal evidence in virtue of a notion of systematic linkage carries over to an analysis of the epistemic role of factive evidence" (2016, p. 881; cf forthcoming). In her (2014) she offers her standard argument for phenomenal evidence without even mentioning the existence of factive evidence. Clearly, despite a related explanatory and metaphysical dependency that runs in the opposite direction, she thinks that the best way to understand factive evidence is in terms of phenomenal evidence (as a souped-up version of the latter, rather than the latter being a watered-down version of factive evidence).

This argument simply presupposed that there's an epistemic good. (This is why I think the argument is more Williamsonian, for whom factive mental states explain the goodness of nonfactive states, not the other way around.) What she's trying to do here is to explain why there is factive evidence and why it provides a "rationality boost" (2016, p. 881) over phenomenal evidence.

Her second argument relies on her theory of perceptual content. Briefly, and approximately, she holds that perceptual representations have demonstrative contents along the lines of

That_{α} is F.

This is the *type* of content had by perceptual states. In the case of hallucination, however, there's no particular for the *token* demonstrative element to attach to, nothing for it to attribute F-ness to, nothing whose F-ness would make the ascription true or veridical. The token therefore has, she says, a "gappy" content, like

That__ is F.

States with gappy contents can't be true, and it's irrational to heed something that can't be true, so hallucinatory perceptual states can't provide evidence for beliefs (2013, p. 732; 2018, p. 183).

The first problem with this argument is that she is clearly committed to the claim that perceptual states with gappy contents do indeed provide evidence (namely, phenomenal evidence) even though they can't be true. Two ways out of this apparent problem present themselves. First, she could dial back the claim so as to state only that things that can't be true can't provide *factive* evidence, not that they can't provide evidence simpliciter. Unfortunately, of course, this principle has little, if any, independent plausibility, especially since we're still trying to understand what factive evidence even *is*. A second option is to claim that ordinary perceptual belief involves two distinct heedings. She writes, "[a] gappy token content does not provide evidence, since a gappy content cannot be true.… It is not rational to heed something that by its very nature could only be false or lacking a truth value. Therefore, gappy content does not provide factive evidence" (2018. p. 183). This can be squared with her view about phenomenal evidence if perceptual states have two natures: one corresponding to the content type and one to the token content. It's only the latter that's gappy, and the former *can* in some sense be true (or let us grant); so there's no conflict with the view about phenomenal evidence. Very roughly, the idea would be that when I'm duped by a hallucination, I'm heeding the hallucination qua representation as of a white cup and also heeding the hallucination qua hallucination. More exactly, I'm heeding both the content type—which it's rational to do—and heeding the content token—which it's not rational to do. I'm not sure that this is quite her view. There is little textual evidence that she thinks there are two heedings in every case of perceptual belief, other than the fact that this would render the above quotation consistent with what's clearly her position regarding phenomenal evidence.

But anyway, is it really not rational to heed (to base beliefs on) something that, by its nature, can't be true? It's sometimes rational to heed lies, even though a lie is, by its very nature, something that can't be true. Maybe she means it's irrational to heed something whose content precludes it from being true. However, it seems very

plausible that we can have justified false mathematical beliefs; and when we do, it can be rational to heed these beliefs, even though they can't possibly be true. If I have very good reason to think that 447 is prime, I can justifiably infer from this that there's a prime number between 440 and 450. The conclusion is prima facie justified even though the premise can't possibly be true. Closer to current concerns: someone on the street shouts, "That man stole my purse!" It might be rational to believe that someone stole her purse, even I can't tell which man she's referring to. And I might be fully justified, even if it turns out that the person shouting is failing for some reason to refer to *anyone*, although in such a case, her claim has a gappy content and thus can't possibly be true. We might try to assimilate this case to a case of phenomenal evidence without factive evidence, but this seems even more counterintuitive than the perceptual case.

Schellenberg could partly avoid some of these objections by claiming not that things that can't be true can't provide *any* evidence, but only that they can't provide *as much* evidence as things that can, or that they can't provide the same *kind* of evidence (in some nontrivial sense of 'kind of evidence'). However, although the proffered claim is weaker than the one she makes, it doesn't have nearly the same intuitive pull. It might seem initially plausible that a thing that can't be true can't give reason to think something else is true. Once we see that initial plausibility to be misleading, it's hard to see why such a thing would have to provide any *less* of a reason. The examples from the previous paragraph seem to undercut this principle as well. If I'm really justified in thinking 447 is prime, then I should believe there's a prime number between 440 and 450. This inference, of course, is weakened by the fact that I could easily figure out that 447 is not prime, despite my evidence that it is (if, for example, I happen to know that any number whose digits sum to a multiple of 3 is itself a multiple of 3); but it's not weakened merely by the (unknown) fact that the premise is false, or that it's necessarily false.

A possible third, implicit, argument is from functions again. Really crudely, the idea is this: perceptual capacities have a function, so we should listen to them. In the good case, they fulfill that function, so we should *really* listen to them. That is, they generate justification whenever employed, but they generate extra justification when successfully employed, all in virtue of their functions.

But just as I argued above that no normative implications follow immediately from an item's merely *having* a function, no such implications follow from its *fulfilling* that function. Rather, it's *functioning properly* that generates or tracks normative status. A creature that's bleeding to death will have a heart that is no longer able to fulfill its function, even though that heart continues to function properly. There's *nothing wrong with* the heart; it's not malfunctioning. But it's unable to perform its function because there's simply no blood to pump. So failure of an x to perform (/execute/satisfy/fulfill, etc.) its function does not imply any normatively negative assessment of x. Nor, of course, does successful performance by x imply proper function or therefore positive evaluation of x. A *bad* heart might nevertheless pump blood (even if sporadically, or slowly, or with poor pressure, etc.), but it's still bad in the relevant sense.

This argument is at most implicit. Schellenberg never explicitly defends or articulates an argument along these lines. Still, I think function talk is sufficiently salient in her work, and it's clearly supposed to have *some* normative consequence, so it's

worth pointing out that functions won't help in this way. Functions of the sort that she invokes—or even of the sorts that Burge, Bergmann, or Graham invoke—won't underwrite the disjunctivist claim that there's ipso facto more justification in the good case than in the bad case.

This leads back to the general point that I started this section with. Surely the disjunctivist is correct that there's something epistemically better about the good case than about the bad case, and surely there's *some* sense in which the perceptual capacity exercise in the bad case is a failure and hence in *some* sense defective. The question is whether it's the kind of failure that affects justification and not just truth or knowledge. Let's grant that, for a suitable range of exercises E: a successful E is in some sense epistemically better than an unsuccessful E; and that a more justified E is in some sense epistemically better than a less justified E. It clearly doesn't follow that a successful E is more justified than an unsuccessful E. Successful (i.e., good-case) exercises of perceptual capacities are surely epistemically better in the sense that they put one in a position to know, while bad-case exercises do not. What the disjunctivist needs to show is that *in addition* to this difference, there's also a difference in the quality or amount of justification that's had simply in virtue of the one case being good and the other bad. Schellenberg's arguments don't show this.

3 Arguing against disjunctivism

The discussion up to now has been concerned with fending off one particular argument for disjunctivism. Obviously, my attacking that argument does nothing to undermine other arguments for disjunctivism, but addressing them all would be far beyond the scope of this paper. Instead I'll now go on the offensive and argue that there are good reasons to reject disjunctivism.

First, it is worth reminding the reader that disjunctivism is counterintuitive in a way that's brought out by the famous New Evil Demon objection to reliabilism (Cohen, 1984; Lehrer & Cohen, 1983). This argument states that a brain in a vat or victim of a Cartesian demon will be forming all their beliefs in an unreliable way and thus, according to reliabilism, have no justification at all; although it's intuitively quite appealing that many beliefs—including perceptual beliefs—would be justified even in these conditions. Obviously, this argument is as much a problem for disjunctivism as it is for reliabilism. I think more, because while reliabilism gets this controversial result as a side effect of explanatory and (at least otherwise) plausible claims about the relation between justification and probable truth, this counterintuitive result is disjunctivism's main offering. I realize that disjunctivists are unlikely to be moved by the New Evil Demon argument, but its popularity indicates that it should be taken seriously, even if ultimately resisted. Note that it's not just internalists who like this argument; many externalists go to great lengths to ensure that their theory doesn't have the unattractive implication (e.g., Burge, 2003; Comesaña, 2002; Goldman, 1986, 1988; Graham, 2012; Henderson & Horgan, 2006; Sosa, 1991). Even many of us who are willing to insist that (some) demonworlders are thereby unjustified (Lyons, 2013, Lasonen-Aarnio in prep) feel the intuitive pull of the argument and try to find ways to make the result palatable. In the end, I think we need to insist that a justified belief is one

that is in some sense likely to be true; and I think the way to understand that sense is in terms of belief that's produced by a process that *generally* delivers true beliefs (even if it's not delivering true beliefs in this instance—or in this house, or today). This allows that some demonworlders (especially new and/or temporary residents) *are* forming beliefs using ways that are generally truth-conducive (even if not truth-conducive right this moment) and thus have perceptual justification. For the others, including the "unanchored" demonworlders (Lyons, 2013; Graham, 2024; Lasonen-Aarnio in prep), the lack of a truth-connection makes it plausible on reflection (certainly much less implausible) that they're not justified after all.

Non-reliabilist versions of disjunctivism don't have a similar way of taking the sting out of the New Evil Demon objection, and they don't even have the right sort of truth-connection benefit to offset this disadvantage. Reliabilist versions of disjunctivism can say all of what I've just suggested, but if you're some kind of reliabilist anyway, you should be going with a mixed view, not a disjunctivist view: when hallucination happens in a way that generally makes for unreliability, the beliefs are unjustified; when it happens in a way that's otherwise compatible with reliability, then the beliefs are justified, even if demon-produced. Whether the agent is hallucinating or not is, in itself, irrelevant.

Second, although it's standard to contrast veridical perception with hallucination, most of the sorts of arguments that are supposed to show reduced justification in the hallucination case seem to apply straightforwardly to cases of illusion as well. Illusion doesn't involve reference failure, but it involves an erroneous, thus nonfactive, perceptual state. Thus, illusion doesn't provide conclusive reasons, and if all evidence is factive it doesn't provide reasons (evidence) at all. Illusion seems to be explanatorily and metaphysically parasitic on veridical perception in much the way that hallucination is.

Illusions by their nature can't be true, even if they don't have gappy contents.[18] And so on. Thus, if we're going to hold that hallucination results in less justified perceptual beliefs, we ought to say the same about illusion.

The problem with this is that it's perhaps even more counterintuitive to say that all illusions are unjustified (or less justified) than to say that all hallucinations are. While hallucination involves demonic intervention or something having gone pretty radically wrong with the agent or the environment, illusion typically just involves quirky stimuli. Indeed, the standard take on perceptual illusion in the sciences is that it's the result of a generally reliable process being employed outside of normal operating parameters. Some perceptual illusions are very compelling, and, it seems, *rightly* compelling. Consider the following, from Roger Shepard (1990):

[18] On Schellenberg's view, illusions do have gappy contents, because the predicative element is also demonstrative/de re.

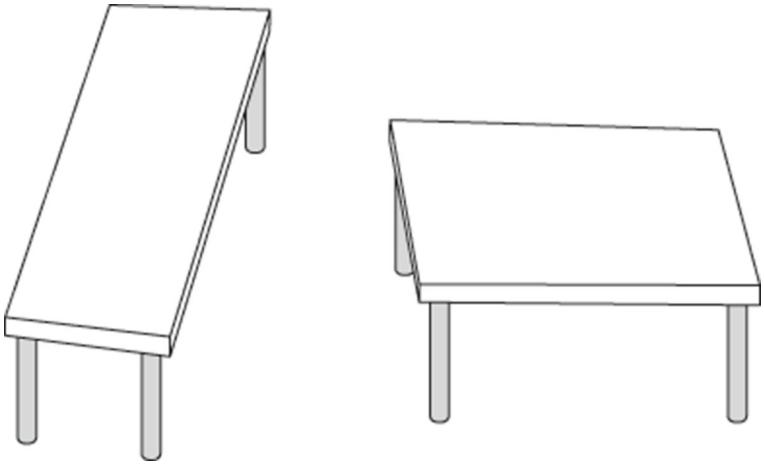

Someone who is unfamiliar with the illusion should count as justified—pretty strongly justified, I would think—in believing on the basis of visual perception that the parallelogram on the left is longer and narrower than the one on the right. It isn't. The disjunctivist has to disagree, and thus to insist, apparently, either that the perceiver should suspend belief—even though there's no reason to think anything is amiss—or perhaps to believe that the two parallelograms are the same size and shape—despite how they look and despite the agent's having no insight into the illusion. This strikes me as counterintuitive, indeed, considerably more counterintuitive than giving the unpopular response to the New Evil Demon. So even if disjunctivists are not worried about the New Evil Demon, they should be worried about illusion.

In fact, it is not obvious that the disjunctivist can allow for false but (fully) justified perceptual beliefs at all. If the false belief mismatches the percept, then it's not supported by that percept; it's not the belief that that percept supports. But if the belief matches the percept, then the percept was erroneous and thus can't confer full justification on a disjunctivist view. But surely false, fully justified perceptual belief is possible. The street performer might look exactly like a statue, the wax object exactly like real fruit, the animatronics robot just like a real turtle, etc. These aren't even the sorts of cases we normally label as illusion. They are, one should think, cases of justified belief.

Third, think about misleading evidence: evidence that speaks in favor of, and thereby justifies, something that turns out nevertheless to be false. Most epistemologists think that misleading evidence is still evidence. Indeed one way to frame the disagreement between disjunctivists and others is to say that while the rest of us think that the hallucinating agent has (at least sometimes) *misleading* evidence, the disjunctivist thinks the hallucinating agent doesn't have any evidence (or doesn't have any evidence of a certain, important sort), but merely *appears* to. The difference, of course, is that misleading evidence justifies and apparent evidence doesn't. Or, more exactly, the

justificatory power of merely apparent evidence is extremely, perhaps fully, attenuated; the justificatory power of misleading evidence is not thereby attenuated.[19]

Whether or not the disjunctivist says this about misleading and apparent evidence, the disjunctivist position is very much like a particular very strange and heterodox view about misleading and non-misleading evidence. Suppose one held that non-misleading evidence ipso facto provides a higher degree of justification than misleading evidence. Then, everything else being equal, if I have two conflicting pieces of evidence, but one of them happens (unbeknownst to me) to be misleading and the other not, I should believe in accordance with the one that's not. Suppose John tells me p and Jayne tells me not-p; I was properly completely agnostic about p before and know them to be equally reliable. The proposed principle implies that I'm justified in believing whichever of them happens to be telling me the truth. This strikes me as a reductio of the supposition that non-misleading evidence provides a higher degree of justification than misleading evidence.

But it gets worse, on this view, depending on how much higher a degree of justification non-misleading evidence is said to yield. I might know that Jayne is *considerably* more reliable than John. Nevertheless, the principle implies that I am still justified in believing John over Jayne, provided that p happens to be true, depending on the size of the non-misleading-evidence bonus (that is, depending on *how much more* justification is provided by non-misleading evidence than by misleading evidence).

Similarly with perceptual justification. Let the "default justification" associated with some percept token be the amount of prima facie justification that believing in accordance with that percept (i.e., believing that things are as the percept "says" they are) would receive from that percept—*if the percept is/were veridical*. That justification might be determined by intrinsic phenomenal properties of the percept, like vividness or clarity; by process reliability; by background beliefs; or by any combination of these and/or similar factors. Nondisjunctivist views, because they don't think that veridicality per se makes any justificatory difference, think that default justification exhausts the prima facie justificatory contribution of any percept. Disjunctivists think default justification exhausts the prima facie justificatory contribution only of veridical percepts; nonveridical percepts contribute *less* than the default justification.

Now consider two conflicting percepts, in different sense modalities, with equal associated degrees of default justification. You hear an event as occurring at location L1 but see it as occurring at L2, or a certain surface looks rough but feels smooth, etc. In all cases, tweak the parameters (vividness, lighting conditions, etc.) as needed for default parity between the competing sense modalities. Disjunctivism has to say that in such cases, you should believe whichever percept is veridical (if one is). That is, you're justified in discounting the erroneous percept and going with the veridical one, even though they have equal default justification, and you have no independent

[19] An anonymous referee mentions someone who has an intellectual seeming in connection with the gambler's fallacy: it *seems* to them like they have evidence for thinking that the next ball will land on black. They don't have any such evidence but only appear to. It does strike me as intuitive that such a person is more rational in committing the fallacy than someone who doesn't have such a seeming. So maybe seeming to have evidence for p just is a very attenuated kind of evidence for p. The more exact formulation in the text is intended to allow for neutrality on this question.

evidence whatsoever for thinking that the one you're trusting is the veridical one.[20] How strongly you are allowed to, perhaps required to, discount the erroneous percept will depend on how rebutting defeat is thought to work and how strong the "veridicality bonus" is said to be. On the plausible but simple view that you should believe the victor to a degree or with a credence proportionate to the justificational difference between them, and on the further assumption that the veridicality bonus is relatively slight (i.e., that a veridical percept only provides *slightly* more justification than an erroneous one), there might not be a great deal of discounting.

But for a disjunctivist who holds that there's no justification in the bad case, or vastly less, the veridicality bonus is large and the discounting will be severe. In connection with this view, consider two conflicting percepts where the erroneous one has significantly more default justification than the veridical one. You have a very clear, leisurely view of an object that looks to be (though it is in fact not) located at precisely L1. That object makes a faint noise, such that, were you going on only the sound, you'd have a pretty tentative (true) belief that it's at L2. On this form of disjunctivism, you should believe it's at L2. You're justified in discounting the clear, vivid, reliable visual percept in favor of the faint, fuzzy, unreliable auditory one, simply because (unbeknownst to you) the auditory percept happens to be correct. This, to me, is an unacceptable result, and it's more unacceptable the more serious the disjunctivism is, i.e., the bigger the justificatory gap it posits between the good and bad cases.[21]

4 Indicator epistemologies and process epistemologies

I've been arguing against disjunctivism, attempting both to undermine some of the arguments that support it (Sect. 2) and to offer direct arguments against the view itself (Sect. 3). I think disjunctivism is wrong. In this section I want to briefly diagnose the error, to explain how and why disjunctivism has gone wrong.

It's widely accepted that prima facie doxastic (henceforth, pfd) justification is determined by why the agent believes as they do.[22] There are two main ways of trying to understand this: in terms of indicators or in terms of cognitive processes. The first emphasizes the basis, or evidence, or ground, or reason on which a belief is based; the

[20] The disjunctivist might say—indeed, might be forced to say—that even if you didn't have any *independent* reason for trusting, say, touch over vision, the fact that you trusted the right one now gives you non-independent reason to do so. You end up with the justified belief that the surface is smooth, which is evidence for the claim that touch is getting this one right and vision is getting it wrong, which is evidence for thinking you should trust touch over vision in this case. Obviously this is an instance of the bootstrapping problem (Cohen 2002; Vogel 2008), and I think it only makes things worse for disjunctivism.

[21] Notice that this argument works whether the disjunctivist says that hallucination provides only apparent evidence, rather than genuine though misleading evidence; or says that hallucination provides evidence of an inferior sort; or says that evidence isn't even part of the justificatory story for perceptual belief.

[22] Propositional justification is not determined thusly, and since propositional justification can yield defeaters (e.g., if you thought about q for a second, you'd realize it's a counterexample), neither is ultima facie doxastic justification.

second emphasizes the process, or capacity, or skill, or virtue that's responsible for the belief.[23]

An *indicator epistemology* is one that holds that a belief is pfd justified just in case it's held on the basis of a good indicator for that belief. Some beliefs are *based on* some things—the paradigmatic example is inference, where a conclusion belief is based on a premise belief; though the concept is often liberalized a bit to allow beliefs to be based on nondoxastic experiences or extramental facts as well. An indicator epistemology holds that the pfd justification of a particular belief is determined by (a) the epistemic status of the indicator (the thing it's based on), and (b) the indication relation—i.e., the support relation—between the indicator and the justificandum belief. The generic view is neutral regarding what kinds of things can serve as indicators (e.g., propositions? experiences? beliefs?), the nature of the basing relation (e.g., to be based on e, must the belief causally depend on e?), the source of the status of the indicator (e.g., do all experiences provide justification, or only those with the right history? only the true ones?) and the support/indication relation (is it contingent and agent-relative, or does it follow necessarily from the contents?) etc. On this view, once you've specified a set of indicators and their epistemic status for an agent, you've settled the question what the agent is propositionally justified in believing, and you've gone a very long way to answering what the agent is doxastically justified in believing. Classical evidentialism (Feldman & Conee, 1985) is a paradigmatic indicator theory, but the class includes externalist varieties as well.[24]

A *process epistemology*, by contrast, holds that a belief is pfd justified just in case the whole relevant causal history of the belief is good. It differs from indicator epistemology in two crucial respects: (a) it doesn't require that a justified belief be *based on* anything, and (b) even for those that are, it's not enough to base it on good evidence; the transition from the indicator to the belief has to also pass muster.[25] The generic view is neutral about the *relevant* causal history (e.g., does it start with sensory transduction, or can it include distal objects?) and what kind of causal history is required for justification (e.g., is it a matter of reliability, of proper functioning,

[23] All of these terms—'evidence', 'reason', 'capacity', 'virtue', etc.—are are problematic and susceptible to being used in very different ways. Schellenberg's "capacities" underwrite an indicator theory, while Millar's (2019) "capacities" underwrite a process theory. That's why I present them in clusters of near-synonyms, so that the near synonymy can help clarify which meanings of these vexed terms are at issue.

[24] What I'm calling an indicator epistemology here might with some justice simply be called evidentialism. I avoid the latter terminology for two main reasons. The first is that the term 'evidentialism' is sufficiently associated (at least in my mind) with a Feldman/Conee-type internalism that I worry it would be misleading. The second is that the term 'evidence' is so polysemous, so contested, and increasingly so widely abused, that its use likely does more harm than good. Leaving out 'evidence' makes it possible to give a statement of indicator epistemology—it makes justification a function of what the belief *is based on*—without distractions.

[25] If an indicator view endorses a causal (/process) theory of the basing relation, it turns out to be a special case or degenerate version of process theory: one that holds that every (justified) belief is based on some indicator and also holds that this basing exhausts the relevant causal history between indicator and belief. (Thus, Comesaña (2010) presents a view he takes to be process reliabilism, though it turns out to be indicator reliabilism on my taxonomy, and Alston (1995) seems to not see an important difference between his own (1988) indicator reliabilism and process reliabilism (Goldman 1979, 1986).) In what follows I'll use 'process epistemology' a bit more narrowly and exclusively, to mean a process theory that *isn't also* an indicator theory.

of proceeding from virtue? etc.). Process reliabilism (Goldman, 1979; Lyons, 2009) is the paradigm process theory, but many other views count as well, including some forms of virtue responsibilism.

I mean for indicator epistemologies to be understood pretty narrowly and process epistemologies pretty broadly. A process epistemology, for example, can allow that our beliefs, perceptual or otherwise, are based on reasons—*provided that* the basing in understood in causal terms. Someone might endorse a view that looks a lot like a standard indicator theory but hold that pfd justification requires not just that the belief be *based* on good indicator but that it be *properly* based on a good indicator. Whether such a view really is an indicator epistemology will depend on how proper basing is to be cashed out. If proper basing is to be understood in terms of causal processes rather than in terms of the presence of good indicators, then the view is a process epistemology on the current taxonomy.

I'm not sure that everyone accepts one of these two views, although I don't know that anyone has articulated a third way. Traditional foundationalists (Chisholm 1977) and coherentists (BonJour, 1985) are probably best viewed as endorsing indicator theories, along with a non-causal account of the basing relation. Many authors are simply silent on the matter.

The idea that justification should have quite a lot to do with *reasons*, or evidence, is an intuitively appealing one, as is the narrower claim that perceptual justification should have a lot to do with reasons/evidence; and indicator theories are quite commonly endorsed. But they are also very problematic. The main difference between indicator and process theories is that the indicator theory doesn't care *how* you get from the evidence to the belief, so long as the belief "fits" the evidence in the relevant sense. Process theories do care; getting from p and 'p ⊃ q' to q by means of thoughtful, reflective, deliberate reasoning is epistemically better (more justification-conferring) than getting there by means of drunken, sloppy guesswork—*even though both inferences are instances of modus ponens* and equally valid. More generally, for nearly any e and h, we can imagine an epistemically bad way of getting from e to h, no matter how good the fit between e and h; and these bad ways do not result in justified belief. An indicator view cannot license this verdict, because the intervening process doesn't affect the fittingness relation, and an indicator epistemology holds that justification is fully determined by the fit between the belief and what it's based on. Process views have, in my opinion, the clear advantage in this respect.[26]

The problem is not that the indicator view makes too few distinctions; it's that it makes the wrong kind of distinctions. It evaluates beliefs solely on the basis of the indicators they're based on as opposed to the more general causal mechanism, of which this basing is, at most, only a part. A process theory can allow that some inferences are bad, even though they happen to be instances of modus ponens, but also conversely, that some inferences are good, even though they're instances of affirming

[26] The disadvantage is so clear that one might suspect that no one has really ever held an indicator view, thus construed. See Conee and Feldman (2004, p. 93) and Comesaña (2010, p. 384) for explicit statements. The view is implicit in Williamson (2000), insofar as justification (or blamelessness, as he now calls it) is to be actually analyzed in terms of evidence and thereby knowledge, as well as traditional foundationalism and coherentism, and effectively any theory that uses propositional justification as a proxy for doxastic justification.

the consequent (see Lyons, 2019, 2022 for more detail and argument), provided that they're the product of a *generally* good process, even though mistaken in this instance.

Similarly, compare process reliabilism and indicator reliabilism in fake barn country. Suppose that fake barn country is quite extensive, both spatially and temporally, and that, although very few things look like barns, the vast majority of things that do are facades. The environment is normal in all other respects. The indicator reliabilist will say that the perceiver's barn belief is unjustified, since the barn experience does not, in that environment, reliably indicate a real barn. But the process reliabilist will say that the belief is justified, since the processes involved are the very same processes as those involved in perception of houses, cars, etc. and are still highly reliable, despite the proportion of barn facades to real barns.[27] Thus, even if we agree that pfd justification is a matter of *why* the agent holds the belief, we can get different verdicts, depending on whether we unpack that 'why' in terms of indicator or in terms of process.

In this particular case, the process reliabilist gets the intuitively right answer: barn beliefs are justified even in fake barn country; indicator reliabilism gets it wrong. In demon worlds, they get the same, counterintuitive, answer (though see discussion at the beginning of Sect. 3 above). In cases of acute hallucination (in normal, non-demon-world, non-brain-in-vat, scenarios), the indicator reliabilist will generally say that the belief is justified, though this will depend on the environmental details and the details of the theory. What won't matter if it's an indicator theory is *anything else going on in the agent's head*. It won't matter whether the hallucination was endogenous or exogenously caused, whether it was the result of a 10 ms glitch or massive and widespread cerebral dysfunction, only whether that indicator (e.g., that perceptual experience) was a good (i.e., reliable) indicator of the truth of the belief. Surely, however, at least if we're not internalist traditionalists, these sorts of factors should matter. To my mind, such considerations strongly motivate even further the general theoretical position that we should endorse process epistemologies over indicator epistemologies.

Not all indicator theories are reliabilist, so not all offer the same verdict about fake barn country. But indicator reliabilism is an illustrative case, because it gives the wrong answer here for the same reason that indicator theories generally give the wrong answer about bad but correct or good but incorrect inference: that they put too much emphasis on the relation between indicator and belief and not enough on the larger, more inclusive, belief-forming process. For the process theory, it's the whole relevant causal chain leading up to belief that counts as *why* the agent holds the belief and that thus determines doxastic justification, not just one special link.[28] There are various ways for the process theorist to make sense of the "relevant" causal chain; no one, for example, thinks the whole causal chain all the way back to the big bang is relevant. The point is that the process theory has more to work with than a single, experiential link. An indicator theory, on the other hand, will always be at serious risk

[27] For more details, see Lyons (2016b, 2019, 2022, 2023).

[28] The complex of causes is generally less linear than the chain metaphor suggests, but this doesn't matter to the present concerns.

of overstating the importance of the indicator at the expense of the rest of the chain, since that's the only link the theory has to work with.

A certain kind of internalism requires an indicator theory. If the only factors that can determine (pfd) justification are introspectible, non-factive mental states of the agent, then perceptual justification will have to be determined by the beliefs and experiences of the agent. The rest of the causal history of the belief doesn't satisfy the internalist requirements. The result is a traditional (and traditionalist, in the sense of this paper) evidentialism, a paradigmatic indicator theory. For most evidentialists, the magic link is the perceptual experience, though for older and especially coherentist theories it was the perceptual belief itself that matters.

But if that particular internalism is abandoned—in favor of an externalist view, or a broader mentalist view that allows nonintrospectible mental states to play a role, for example—then there is no longer any reason to single out experiences as privileged links in the causal chain and the prime determinants of doxastic justification. Disjunctivism, though it rejects the traditionalist's internalism,[29] retains the indicator epistemology that follows from it.[30] The result is the same bad result that attaches to indicator theories more generally: the inability to make some of the epistemological distinctions we need to make (e.g. between "good" and bad cases of hallucination) and the enforced, artificial distinction between cases that demand similar epistemological treatment (e.g., between well-formed but false perceptual beliefs and well-formed true perceptual beliefs).

Perception typically involves a complex causal chain from distal object to perceptual belief. An indicator theory strongly privileges a very small number of links in this

[29] The disjunctivist might still be internalist in some other sense (McDowell 1982, 2011, 2018; Pritchard 2012).

[30] Williamson (2000, 2009) might appear to be a counterexample. He is certainly not an experientialist in the way that many disjunctivists are. Despite his insistence that the agent in the good case has different evidence than the agent in the bad case, it's unclear to me whether he's a disjunctivist in the present sense, that is whether he thinks the agent in the good case has has more or better *justification* than the agent in the bad case. (He explicitly denies being a disjunctivist in the metaphysical sense.) The answer to this question will depend on how he understands the status of the reflexive evidence involved in knowledge. For Williamson, all knowledge is reflexively justified in the sense that it provides evidence for itself. The agent in the good case perceptually knows that p, and that knowledge gives that agent evidence and thus justification for p; the agent in the bad case lacks this knowledge and hence lacks this evidence and justification. (Both agents also typically know that it perceptually seems to them that p, which provides additional evidence, but this doesn't distinguish the two cases.) If this reflexive evidence is meant to be substantive rather than trivial—i.e., to boost one's epistemic position over one who doesn't have such justification—then the position is indeed a disjunctivist one in the current sense. I think this take on reflexive evidence is (/would be) a very bad view, for the sorts of reasons laid out in Sect. 3 above, but also for reasons I'd like to take up elsewhere but which go beyond the scope of this paper. I'm uncertain whether this disjunctivist position is Williamson's view or whether he means for this reflexive evidence to be trivial and not producing any boost in justification, rationality, blamelessness or whatever evidence is supposed to boost. (One way to pursue this non-disjunctivist move though still in a Williamsonian vein can be found in Fratantonio 2021.) It's worth noting here that although Williamson says that the agent in the good case has different evidence, this only holds for an agent in the good case *who trusts their senses* and believes in accordance with them. One who suspends belief presumably fails thereby to know and thus loses that special reflexive evidence. Such an agent is thus not more propositionally justified in the good case than in the bad case.

In any case, Williamson seems to be committed to an indicator epistemology in general, and if he is a disjunctivist in the way just described, it's because the perceptual knowledge that p is serving as a good indicator for the perceptual belief that p.

chain as relevant to justification, dismissing the rest as relevant maybe to knowledge, and maybe to establishing a basing relation, but not otherwise directly relevant to justification. This, I have been arguing, is a mistake. And it's a mistake that's central to disjunctivism, as disjunctivism is simply the view that veridical indicators have a different justification-conferring status from hallucinatory indicators. If we abandon an indicator epistemology in favor of a process epistemology, as I think we should do for reasons independent of perception, we won't even be tempted by the disjunctivist position.

5 Conclusion

Everyone agrees that the veridically perceiving subject is in a better epistemic position than their hallucinating counterpart. Epistemological disjunctivists claim that this is not just a matter of knowing or being in a position to know, but of the perceiver having more or better *justification* than the hallucinator, simply in virtue of the fact that they're veridically perceiving. Traditional and mixed views reject this claim about justification.

I've argued here that we shouldn't follow the disjunctivist in this. One set of prominent arguments for disjunctivism is unconvincing on examination, and there are good reasons to reject the general view on its own merits. The idea that hallucination per se could make a justificatory difference is only plausible if we assume an indicator epistemology, which is problematic on independent and antecedently known grounds, in addition to the specific problems it faces in the form of disjunctivism about perceptual justification.

I haven't tried to argue for a mixed view in this paper, though I have in (Lyons, 2016b, 2019). For the current purposes, I'm happy to have concentrated on the negative argument, on the case against disjunctivism. However, if we endorse a process epistemology, neither traditionalism nor disjunctivism will look at all appealing; some kind of mixed view seems the natural result. The perceptual experience, on such a view, is just one of many links in the causal chain that determines the epistemic status of the belief. Thus it is unlikely that justification will be strictly a function of that one link, in either the way that the traditionalist or the disjunctivist would have it.[31]

Declarations

Conflict of interest No funds, grants, or other support was received, nor does the author have any other financial or non-financial interests to disclose.

[31] Earlier versions of this paper or some of the material in it were presented to Washington University, Indiana University, the University of Glasgow, and Ben Gurion University. Thanks to audiences there for helpful comments. Thanks also to Ori Beck, Derek Brown, Giada Fratantonio, Peter Graham, Allan Hazlett, Fiona Macpherson, Matthew McGrath, Susanna Schellenberg, and Timothy Williamson for helpful discussions.

References

Alston, W. P. (1988). An internalist externalism. *Synthese, 74*, 265–283.
Alston, W. P. (1995). How to think about reliability. *Philosophical Topics, 23*, 1–29.
Bergmann, M. (2006). *Justification without awareness: A defense of epistemic externalism*. Oxford University Press.
BonJour, L. (1985). *The structure of empirical knowledge*. Harvard UP.
Burge, T. (2003). Perceptual entitlement. *Philosophy and Phenomenological Research, 67*(3), 503–548.
Burge, T. (2020). Entitlement: The basis for empirical epistemic warrant. In P. Graham & N. Pedersen (Eds.), *Epistemic entitlement.* Oxford: OUP.
Byrne, A. (2014). Perception and evidence. *Philosophical Studies, 170*(1), 101–113.
Byrne, A. (2016). The epistemic significance of experience. *Philosophical Studies, 173*, 947–967.
Clark, A. (2015). Perceiving as predicting. In D. Stokes, M. Matthen, & S. Biggs (Eds.), *Perception and its modalities* (pp. 23–43). Oxford University Press.
Cohen, S. (1984). Justification and truth. *Philosophical Studies, 46*, 279–295.
Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research, 65*(2), 309–329.
Comesaña, J. (2002). The demon and the diagonal. *Philosophical Studies, 110*, 249–266.
Comesaña, J. (2010). Evidentialist reliabilism. *Nous, 44*, 571–600.
Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies, 48*, 15–34.
Davidson, D. (1986). A coherence theory of truth and knowledge. In E. Lepore (Ed.), *Truth and interpretation: Perspectives on the philosophy of Donald Davidson* (pp. 307–319). Blackwell.
Fratantonio, G. (2021). No infelicity for the infallibilist. *Theoria, 87*, 1307–1321.
Goldman, A. I. (1979). What is justified belief? In George Pappas, ed. *Justification and knowledge*. Dordrecht: Reidel.
Goldman, A. I. (1986). *Epistemology and cognition*. Harvard UP.
Goldman, A. I. (1988). Strong and weak justification. In J. E. Tomberlin, ed. *Philosophical perspectives, 2, Epistemology*.
Graham, P. J. (2012). Epistemic entitlement. *Noûs, 46*(3), 449–482.
Graham, P. J. (2024). The new evil demon problem at 40. *Philosophy and Phenomenological Research, 1*, 1–27.
Graham, P. J., & Pedersen, N. J. (2020). Dretske & McDowell on perceptual knowledge, conclusive reasons, and epistemological disjunctivism. *Philosophical Issues, 30*(1), 148–166.
Hellie, B. (2011). There it is. *Philosophical Issues, 21*, 110–164.
Henderson, D., & Horgan, T. (2006). Transglobal reliabilism. *Croatian Journal of Philosophy, 6*, 171–195.
Lasonen-Aarnio, M. (in prep). *Feasible dispositions: A normative framework*.
Lehrer, K. (1990). *Theory of knowledge*. Westview.
Lehrer, K., & Cohen, S. (1983). Justification, truth, and coherence. *Synthese, 55*, 191–207.
Littlejohn, C. (2012). *Justification and the truth-connection*. Cambridge University Press.
Lyons, J. C. (2009). *Perception and basic beliefs: Zombies, modules, and the problem of the external world*. Oxford University Press.
Lyons, J. C. (2013). Should reliabilists be worried about demon worlds? *Philosophy and Phenomenological Research, 86*(1), 1–40.
Lyons, J. (2016a). Experiential evidence? *Philosophical Studies, 173*(4), 1053–1079.
Lyons, J. (2016b). Goldman on evidence and reliability. In H. Kornblith & B. McLaughlin (Eds.), *Goldman and his critics.* Blackwell.
Lyons, J. C. (2016c). Inferentialism and cognitive penetration of perception. *Episteme, 13*(1), 1–28.
Lyons, J. C. (2019). Algorithm and parameters: Solving the generality problem for reliabilism. *Philosophical Review, 128*(4), 463–509.

Lyons, J. C. (2020). Two dogmas of empirical justification". *Philosophical Issues, 30*, 221–237.
Lyons, J. C. (2022). Cognitive Diversity and the Contingency of Evidence". *Synthese, 200*(3), 202.
Lyons, J. C. (2023). Methods, processes, and knowledge. In L. Oliveira (Ed.), *Externalism about knowledge.* Oxford University Press.
Lyons, J. C. (forthcoming). Capacities second (or never). In O. Beck & M. Vuletić, eds., *Empirical reason and sensory experience*.
Macpherson, F., & Batty, C. (2016). Redefining illusion and hallucination in light of new cases. *Philosophical Issues, 26*(1), 263–296.
McDowell, J. (1982). Criteria, defeasibility and knowledge. *Proceedings of the British Academy, 68*, 455–479.
McDowell, J. (2011). *Perception as a capacity for knowledge*. Marquette University Press.
McDowell, J. (2018). Perceptual experience and empirical rationality. *Analytic Philosophy, 59*(1), 89–98.
Millar, A. (2008). Perceptual-recognitional abilities and perceptual knowledge. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, action, knowledge.* Oxford University Press.
Millar, A. (2011). How visual perception yields reasons for belief. *Philosophical Issues, 21*(1), 332–351.
Millar, A. (2019). *Knowing by perceiving*. Oxford University Press.
Pritchard, D. (2012). *Epistemological disjunctivism*. Oxford University Press.
Schellenberg, S. (2013). Experience and evidence. *Mind, 122*(487), 699–747.
Schellenberg, S. (2014). The epistemic force of perceptual experience. *Philosophical Studies, 170*, 87–100.
Schellenberg, S. (2016). Phenomenal evidence and factive evidence. *Philosophical Studies, 173*, 875–896.
Schellenberg, S. (2018). *The unity of perception: Content, consciousness, evidence*. Oxford University Press.
Schellenberg, S. (forthcoming). What is explanatory fundamental in an analysis of perception? In O. Beck & M. Vuletić (Eds.), *Empirical reason and sensory experience*.
Shepard, R. N. (1990). *Mind sights: Original visual illusions, ambiguities, and other anomalies*. WH Freeman and Company.
Siegel, S. (2016). *The rationality of perception*. Oxford University Press.
Sosa, E. (1980). The raft and the pyramid: Coherence versus foundations in the theory of knowledge. *Midwest Studies in Philosophy, 5*, 3–25.
Sosa, E. (1991). *Knowledge in perspective: Selected essays in epistemology*. Cambridge University Press.
Vogel, J. (2008). Epistemic bootstrapping. *The Journal of Philosophy, 105*(9), 518–539.
Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
Williamson, T. (2009). Replies to critics. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge.* Oxford University Pressa.