CrossMark

Episteme, 13, 1 (2016) 1–28 © Cambridge University Press
doi:10.1017/epi.2015.60

INFERENTIALISM AND COGNITIVE PENETRATION OF PERCEPTION

JACK C. LYONS
jclyons@uark.edu

ABSTRACT

Cognitive penetration of perception is the idea that what we see (hear, taste, etc.) is influenced by such "cognitive" states as beliefs, expectations, and so on. A perceptual belief that results from cognitive penetration may be less justified than a non-penetrated one. Inferentialism is a kind of internalist view that tries to account for this by claiming that (a) some experiences are epistemically evaluable, on the basis of why the perceiver has that experience, and (b) the familiar canons of good inference provide the appropriate standards by which experiences are evaluated. I examine recent defenses of inferentialism by Susanna Siegel, Peter Markie, and Matthew McGrath and argue that the prospects for inferentialism are dim.

Cognitive penetration of perception (henceforth, CP) is the idea that what we see (hear, taste, etc.) is influenced by such "cognitive" states as beliefs, goals, desires, expectations, and so on. More precisely, the idea is that the contents of perceptual states can be influenced in a particular, more or less direct, top-down way, by the cognitive states of the organism.[1] CP is a controversial empirical thesis (Zeimbekis and Raftopoulos 2015), but I won't be concerned here with whether it is true or not. Instead I will focus on the epistemological implications of the view *if* it is true, although of course the question is much more interesting if the kinds of CP at issue are empirically plausible.

It is generally thought that CP might have negative epistemological implications; after all, if perception is determined by our antecedent expectations, then this would threaten to cut us off from the world outside of us, to undermine perception's role as tribunal (Fodor 1984; Raftopoulos 2009).

More recently, it has been suggested that CP might offer counterexamples to certain kinds of internalist epistemologies (Markie 2005, 2006; Lyons 2011; Siegel 2012).

A traditional internalist view holds that perceptual experiences are simply *given*: they are not something we are in any interesting sense responsible for – rather, we simply find ourselves having certain experiences and lacking others. They serve as an epistemological starting point in the sense that epistemic evaluation only enters the picture downstream from experiences. Experiences are not the sort of thing that can be justified or unjustified, but their epistemic consequences are the same as those of justified beliefs: perceptual experiences can confer justification on beliefs, and beliefs that conflict with experience are to that extent epistemically undermined.[2] These epistemic consequences of

1 For more details, see Lyons (2015b).

2 All discussion here of justification will be concerned with doxastic, rather than propositional, justification, and with prima facie, rather than ultima facie, justification.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

experiences obtain simply due to their being perceptual experiences. Experiences may differ in their epistemic potency, on this traditional view, but only in virtue of their intrinsic features (e.g., an experience might be fleeting, or hazy, or ambiguous, etc.), *not* in virtue of their causal histories.[3]

The possibility of CP, however, might offer counterexamples to this traditional view. Consider two influential examples.

> **Angry Jack** (Siegel 2012): Jill believes for no good reason that Jack is angry, and this affects the way Jack looks to her, producing in her a visual experience as of Jack's being angry. She thus believes, even more firmly than before, that Jack is angry.

Intuitively, Jill's belief is not now justified, even though she has the kind of experience that would ordinarily justify the belief that Jack is angry, and the failure to justify is presumably not due to differences in her belief system (she doesn't, for example, believe or have any reason to believe that her experience has resulted from CP).

> **Prospectors** (Markie 2005, 2006): Gus is a seasoned prospector; Virgil is a novice. While out prospecting for gold, they both look at a rock. The rock looks like gold to Gus, as the result of his perceptual expertise; it also looks like gold to Virgil, but only as the result of wishful thinking.

Virgil's belief, like Jill's, seems to lack prima facie justification. Thus, we have putative cases of cognitive penetration of perception, where the experience loses some or all of the justifying power it would ordinarily have had, apparently because the experience has the wrong kind of etiology.

Here then is the problem of cognitive penetration: some beliefs fail to be justified, because they result from CP, even though the experiences those beliefs are responses to are of a type that would ordinarily justify such beliefs (i.e., the respective contents and intrinsic characters of the experiences and beliefs are such that experiences with those contents and characters would normally justify beliefs with those contents). An adequate epistemology has to make room for this fact. The traditional view seems not to be able to do so.

One obvious and natural response to these cases is to reject the traditional view in favor of some kind of externalist epistemology (Lyons 2011). An alternative – and this is what I want to focus on in this paper – is to develop an internalist view I will call "inferentialism," which holds that (a) experiences are indeed epistemically evaluable, on the basis of why the perceiver has that experience, and (b) the familiar canons of proper belief-formation provide the appropriate standards by which experiences are evaluated. Versions of inferentialism have been recently defended by Susanna Siegel (2013a, 2013b, 2013c, Forthcoming, in prep.), Peter Markie (2005, 2006, 2013), and Matthew McGrath (2013a, 2013b). I will argue that all of these inferentialist views are incapable of solving the kinds of problems they set for themselves. If CP really does pose a problem for the traditional view, then the externalist solution looks a lot more promising than the inferentialist solution.

3 Obviously, the agent's background beliefs will partially determine exactly which beliefs an experience might eventually justify. Whether an experience has *any justifying power at all*, however, is determined entirely by its intrinsic features, according to the traditional view.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

1. INTERNALISM AND EXTERNALISM

A reliabilist epistemology offers a natural and straightforward explanation of what's wrong with CP: it makes us less reliable. More exactly, CP is epistemically bad – when it is bad – because and when it cuts us off from the world around us, when it makes us less sensitive to our environments, when it makes us more likely to believe *p* whether or not *p* is actually true. This seems to capture nicely what's going wrong in the central cases: Jill's belief is unreliably formed; she would believe Jack was angry even if he weren't. Virgil would believe of any yellowish rock that it's gold. Jill and Virgil would have had accurate beliefs or been agnostic, were it not for the pernicious influence of cognition on perception; this influence makes them worse at avoiding falsehood, at maintaining a high ratio of true to false beliefs; it makes them less reliable. This reliabilist explanation of what is epistemically wrong with CP leads to a reliabilist view about *when* CP is epistemically pernicious: exactly when it leads to a decrease in reliability (Lyons 2011). This is important, as CP does not always have negative epistemic consequences, as we will see shortly. Any treatment of the problem of CP will have to explain what is epistemically bad about CP and do so in a way that delineates the pernicious cases of CP from the innocuous ones.

This is just a sketch, and a worked-out version of reliabilism would have to make good on this initial promise and show that the bad cases of CP really are cases of reduction in reliability of the relevant sort (different versions of reliabilism understanding the relevant sort of reliability differently). But for the aforementioned reasons, some kind of generic reliabilism seems well suited for handling cases of CP. Notice that the reliabilist view as sketched here finds a middle path between inferentialism and the traditional view. Like the inferentialist, the reliabilist denies that experiences are *ipso facto* beyond epistemic reproach: some experiences have, by dint of their etiology, more justification conferring power than other intrinsically similar experiences.[4] The reliabilist can agree with the traditionalist, however, that experiences are beyond *evidential* reproach: the epistemic status of any beliefs or inferences causally responsible for an experience is irrelevant to the epistemic status of that experience.

I have suggested that the reliabilist thinks of CP (at least in the bad cases) as cutting us off from the world outside of us. Is there an internalist way to understand this cutting off? If not, is there an internalist alternative to understanding CP's epistemic significance in terms of such cutting off?

Internalism is the view that justification is determined entirely by factors "internal" to the agent. The appeal here to the "internal" tends to get unpacked in one of two ways, in terms of reflective access, or in terms of supervenience. *Access internalism* is the view that the only factors relevant to justification are those that the agent can know to be present (or not) by reflection, where reflection is generally understood to include introspection, rational intuition, and a priori reasoning. *Mentalism* is the view that justification supervenes on the (nonfactive) mental states of the agent. We could get different versions of internalism within these two major groups depending on whether we include or omit the above

4 I am granting, for the sake of argument, that experiences really do play an epistemic role. If not, the situation is all the more favorable for reliabilism, since even though many reliabilists endorse the standard view about the role of experience in perceptual justification (e.g., Goldman 2011), reliabilism can be made to work even on the assumption that experiences *per se* are epistemically irrelevant (Lyons 2009).

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

parenthetical expressions, and depending on what exactly we claim can be known by introspection, or what counts as mental, and so on. For the current purposes, these rough characterizations will do.

The relevant cases of CP, like **Angry Jack** or **Prospectors**, can be specified in such a way that the epistemic agent is not in a position to know on the basis of reflection that her perceptual belief is being cognitively penetrated. (In fact, I won't be concerned here with any *other* kinds of cases of CP.) An access internalist will have to say that these penetrated beliefs have the same epistemic status as their non-penetrated counterparts would have, since the only differences are factors not available to reflection. One popular such theory is a view that is sometimes called "dogmatism" (Pryor 2000, Tucker 2010) or "phenomenal conservatism" (Huemer 2007). It claims that a belief that *p* is prima facie justified if it is based on an experience that *p*. Phenomenal conservatism is a claim about what makes a particular belief an epistemically proper response to experience: sameness of content. It is *only* content relations that determine which beliefs an experience justifies and only intrinsic properties of the experience (e.g., forcefulness, clarity, etc.) that determine how strong that prima facie justification is, so penetration can't matter epistemically so long as the penetration isn't reflected in the intrinsic properties of the experience. Where the content comes from – why the agent is having an experience with that content – is deliberately excluded from having any epistemic significance. Coherentism and classical foundationalism offer different, competing accounts of what makes a belief a proper response to experience,[5] but they both imply that the agent with the penetrated experience is in the same epistemic position as the agent with the normal, non-penetrated experience. Consequently, proponents of all of these views have to bite the bullet, and claim that CP doesn't have any epistemic significance (Fumerton 2013; Huemer 2013; Tucker 2014).[6]

I myself find this move of biting the bullet fairly implausible, and this is because I think that wishful *seeing* is very similar epistemically to wishful *thinking*, which I take to be a paradigmatic instance of unjustified belief. After all, someone who wishfully believes that *p* isn't generally aware that she's doing so (or capable of being so on reflection); she generally believes *p* because it seems to her that *p*. Unbeknownst to her, unfortunately, it only seems to her that *p* because she wants *p* to be true (Lyons 2011). Whatever the merits of biting the bullet, however, my concern here is with internalist theories that propose to offer non-bullet-biting responses to the problem of CP. Such a theory would have to be an internalism of the mentalist variety, since the cases of interest are cases where the fact of penetration is not available on reflection to the perceiver.

There are two ways a mentalist might go here. One is to claim that some mental factors, though not accessible by reflection, affect what counts as a proper response to experience. That is, these factors fully or partially determine which beliefs a given experience justifies.

5 The best known versions of coherentism, of course, don't leave any role for experiences (Sellars 1956; BonJour 1985; Davidson 1986), but some coherentist theories do (Conee 1988; Kvanvig and Riggs 1992; Gupta 2006).

6 The claim is not that all coherentists and classical foundationalists are *ipso facto* committed to access internalism; it is that coherentists and classical foundationalists claim that experiences themselves are epistemically beyond reproach and that what makes a particular belief a proper response to such an experience is embodied in other beliefs of the agent. Thus two agents who are experientially and doxastically identical must be epistemically identical, even if one is the victim of CP and the other not.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

Then the mentalist can make the further claim that the causal history of the experience is one such factor.[7] A second route for the mentalist is to claim that the fact of CP (or CP of the pernicious sort) somehow results in a "downgrade" of the experience, where a downgraded experience is one that, perhaps due to nonaccessible factors, has less justification-conferring power than an experience of that type normally would have. This is the inferentialist's move. Furthermore, the inferentialist holds that an experience is downgraded *as the result of having an etiology that is rationally assessable*. On this view, CP, at least in its pernicious forms, provides an etiology for perceptual experiences that undermines those experiences' epistemic potency, because that etiology fares poorly vis-à-vis the standard canons of epistemic rationality.[8] The theory thus seems to be squarely internalist, even if it rejects access internalism in favor of mentalism.

Inferentialism, at least in its extant forms, claims that experiences, even though nondoxastic, are epistemically evaluable. It is worth pausing briefly to explain this unusual claim.

When a belief appraises positively epistemically, it has two logically distinct but connected properties: it is *justified* (i.e., proper, rational, the right thing for the agent to believe), and it is epistemically *potent* (i.e., it has the ability to confer justification on other beliefs to which it's properly related). For beliefs, justifiedness and potency go hand-in-hand; no belief has one of these unless it has both. Most epistemologists hold that things are quite different when it comes to experiences, for these can have potency without having justifiedness. The difference between the traditionalist and the inferentialist is that the traditionalist holds that all experiences have epistemic potency simply in virtue of being experiences, while the inferentialist holds that the potency of an experience is determined by the causal history of that experience. This is all that is normally meant by claiming that experiences are rationally or epistemically assessible.[9]

2. PROCESS INFERENTIALISM

In some recent papers, Siegel (2013a, 2013b) endorses something like the following:

> An experience e with etiology x is downgraded if a belief with the same content as e and etiology x^*, structurally similar to x but involving no intervening experience, would be unjustified in virtue of having x^*.[10]

7 Conee and Feldman (2004 – especially pp. 105–7) seem to have once endorsed the general view that nonaccessible mental factors partially determine what counts as evidence for what, though in Conee and Feldman (2008), they explicitly claim that evidential relations are necessary, and even when they held the earlier view, they never endorsed the further claim. If they had, they could have used it to address the problem of CP.

8 Markie (2005, 2006), as we will see below, pursues both of these two routes, though the first is importantly subservient to the second; this is why he counts as an inferentialist. That is, although other mental states help determine what counts as evidence for what, the primary outcome of this concerns the epistemic status of *experiences*, rather than beliefs.

9 Markie (2013) does hold that experiences can be irrational in much the sense that beliefs can be, but this isn't an essential feature of the inferentialist view.

10 To streamline discussion, I'm combining here Siegel's definitions of "checkered experience," (p. 716) "rationally assessable etiology of experience" (p. 714), and a principle connecting downgrade with checkered experiences (p. 703).

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

The general idea here is clear enough: it's that causal histories matter to the epistemic status of beliefs, and if a given causal history isn't good enough for beliefs, then it isn't good enough for perceptual experiences either. And the general idea is initially plausible: **Angry Jack** seems to be broadly analogous to a case of circular reasoning, and we can easily imagine variants on the case where Jill's experience is influenced by her *fear* that Jack is angry, in which case the overall process is analogous to a case of fearful thinking (Siegel 2013a); **Prospectors** seems analogous to wishful thinking, etc.

What *exactly* the proposal is claiming, however, is quite a bit less clear, because it is unclear how to individuate etiologies and unclear just what is going to count as structural similarity. Obviously, she means for some experiences to be downgraded and some not to be, so if the proposal is going to do any explanatory work, the relevant degree of similarity has to be neither too strict nor too lax. In short, Siegel faces something like the famous generality problem for reliabilism. An etiology is a process by which a belief or experience is causally generated, and because Siegel endorses a process theory of justification, she needs to specify the relevant level of generality for typing inferential/inference-like processes/etiologies.[11] In particular, a proponent of process inferentialism will have to (a) offer some reasonable hint as to what counts as the relevant degree or kind or dimension of similarity or analogy, one that neither trivializes the claim nor attributes downgrade to far too many experiences, and (b) does so in a way that is poised to explain, rather than presupposes, their normative status.

My main emphasis here will be on (a), but regarding (b), note that it is not explanatory in this context to claim that an experience is downgraded because, for example, it results from a process relevantly analogous to wishful thinking. "Wishful thinking" is a thick normative term; to count as an instance of wishful thinking, a belief must be caused in a particular way by a wish or desire, but it must also exhibit a certain kind of normative, epistemic defect. Wishful thinking is thus not a psychological process in the relevant sense (nor, similarly, is fearful thinking, etc.). If there is a class of genuinely psychological processes, all and only the instances of which exhibit this particular epistemic flaw, then we can explain the epistemic status by appeal to membership in this class. Importantly (and as I'll discuss in more detail below), not every belief that is caused by a wish is a case of wishful thinking, and not every perceptual experience that's caused by a wish is a case of wishful seeing. We don't explain why Virgil's experience is downgraded, for example, if all we can say about it is that its etiology is analogous to a case of wishful thinking, unless we can show that in saying this, we are saying more than simply that the experience is causally dependent on desire and that it is epistemically defective.

Once we set aside such appeals to wishful thinking, and fearful thinking, and the like, Siegel is surprisingly silent about just which bad belief-forming processes the bad cases of CP are analogous to, except to note for the one example that **Angry Jack** is structurally analogous to a case of belief perseverance, where one's believing that p at time t causes one to believe that p at $t+n$. On the assumption that belief perseverance is epistemically bad, she claims, this explains why Jill's experience is downgraded (Siegel 2013a: 714). However, belief perseverance and the **Angry Jack** process are structurally similar only at a very coarse level of description. At this high level of abstraction, many perfectly

11 I take it there's no real difference between saying that process tokens a and b are similar/analogous in the relevant respects and saying that they're tokens of the same relevant type, so long as it's not assumed that the "relevant types" must also be psychological (or some other natural) *kinds*.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

good experiences will have bad psychological analogues and will thus wrongly count as downgraded if this very rough degree of similarity is enough to account for the downgrade in **Angry Jack.**

For example, the following case strikes me as a case of innocuous CP, of CP where the resulting belief is perfectly well justified, despite resulting from CP.

> **Snakes** (Lyons 2011): You and I are out hiking in the woods, and I believe (for no good reason) that these woods are full of snakes. This belief primes my visual system for snakes, enabling me to see through their camouflage, and it does so in a way that increases my accuracy without producing any more false positives. I end up visually spotting snakes that you don't see but that are really there in the woods around us.

Intuitively, these perceptual beliefs are perfectly well justified. But my snake beliefs are analogous to belief perseverance beliefs in exactly the way that Siegel claims Jill's belief is; my believing that p at time t causes me to believe that p at $t+n$. So Siegel's proposal seems to downgrade beliefs that shouldn't be downgraded.[12]

Another psychological process that involves believing that p at time t causing belief that p at $t+n$, of course, is memory, although memory clearly doesn't constitute a bad-making etiology for beliefs. In fact, the **Angry Jack** process is psychologically more similar to memory than to belief perseverance. In all three cases, believing that p at t causes one to believe that p at $t+n$. But what is psychologically distinctive and epistemically relevant about belief perseverance is that the agent receives information between t and $t+n$ that undermines her reasons for believing p in the first place, even though she continues to go on believing p anyway. This is absent in Jill's case and in standard memory cases. It seems therefore to be something of a stretch to say that what accounts for the epistemic flaw in Jill's belief is the fact that it has an etiology similar to that of belief perseverance. There is a similarity here, but it's no greater than the similarity to other, epistemically innocuous, etiologies.

Siegel is more impressed by the analogy between **Angry Jack** and perseverance than between **Angry Jack** and memory, even though, as I have just argued, the latter analogy is a closer one than the former. Perhaps it can be pointed out in her defense that her proposal is *not* to derive the epistemic status of an experience etiology from the epistemic status of the *most similar* doxastic etiology. Instead, her view is that as long as there is *any* bad-making doxastic etiology sufficiently similar to a given experience etiology, that experience is downgraded. But this hurts her more than it helps her, for unless the threshold criteria for sufficient similarity are quite strict, the proposal is sure to overattribute downgrade. If they are strict, on the other hand, we have very little idea what they are. She has suggested that having the form 'belief that p at $t \Rightarrow$ belief that p at $t+n$' suffices to make two etiologies sufficiently similar (otherwise the analogy with perseverance

12 I am focusing on a circularity/belief perseverance version of **Angry Jack** and of **Snakes**, rather than wishful thinking or fearful thinking versions of the cases, but the nature of the penetrator doesn't make any difference, if the way it affects experience is the same. It's easy enough to modify **Angry Jack** and **Snakes**, so that the penetrator is a desire or a fear, while everything else is the same. (Siegel's 2013a version of **Angry Jack**, in fact, takes the penetrator to be fear, rather than belief.) These changes don't affect the intuitive epistemic verdicts.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

wouldn't explain **Angry Jack**'s downgrade). But then **Snake** would be a case of unjustified belief, which it isn't.

She also suggests, however (2013a: 716), that merely having the form 'belief that p at $t \Rightarrow$ belief that p at $t+n$' *does not* thereby make two etiologies sufficiently similar, since that structure is not sufficient for downgrade. This, of course, is the right thing to say, but there are two problems with her saying it. First, it rescinds the only real explanation she gives for why Jill's belief is downgraded. We're assured that it's downgraded because her experience etiology is sufficiently similar to some bad-making doxastic etiology, but if it's no longer belief perseverance, then *which* doxastic etiology?

Second, in discussing the relevant cases, she seems to be helping herself to reliabilist considerations that she doesn't have a right to:

> suppose Jill's fear that Jack will get angry has triggered a learning process by which she gets *better* at detecting anger from facial cues, *without many false positives* along the way. … [Or suppose that] Jill's fear that Jack will get angry *improves* her theory of mind (or her theory of Jack's mind), making her *good at* figuring out what will make him angry. Anger-beliefs that are influenced by fear in this way are not, intuitively, ill-founded [i.e., downgraded]. (Siegel 2013a: 716, all emphases mine)

It is natural to invoke reliabilist elements like *getting better* and *avoiding false positives*, but none of these factors can figure into an internalist account, which is what Siegel promises to provide. Instead, the internalist would have to make a convincing case either that our intuitions about downgrade don't really track these reliability considerations (in which case, why mention them?), or that for every reliability difference there is a corresponding difference specifiable in her preferred internalist idiom: namely the canons of proper belief-formation (in which case, we should be told what they are). So far, Siegel hasn't pursued either approach. An alternative would be to find some *other* internalist (mentalist) way to capture the relevant differences, without appeal to rationality, or to canons of good inference. Siegel, however, is explicitly committed here to an internalism that holds the canons that govern proper experience formation are the very same canons that govern proper belief formation.

My argument so far has been this: Siegel claims that a cognitively penetrated experience is downgraded when and because it results from a cognitive process that is sufficiently analogous to a belief-forming process that would violate the canons that govern the latter. But we haven't yet been given any good idea of how to individuate processes or what is required for two etiologies to count as sufficiently analogous. Following up on the brief suggestions Siegel offers, we don't end up with a plausible theory, one that distinguishes the good CP cases from the bad ones. The **Angry Jack** and **Prospectors** cases do bear a certain similarity to cases of circular reasoning or to belief perseverance or to wishful thinking. But we've been given no reason to believe that these similarities *explain* what's epistemically wrong in **Angry Jack** or **Prospectors**, especially, since **Snakes** bears these same similarities but doesn't exhibit epistemic downgrade.

Perhaps the problem is that we've been individuating etiologies at too high a level of abstraction, along the lines of 'belief that p at $t \Rightarrow$ belief that p at $t+n$.' This limits us to a correspondingly lax understanding of structural analogy. Maybe the inferentialist could insist on a finer grained individuation of etiologies and thus a stricter degree of analogy/similarity? It seems independently plausible, after all, that the details will matter to

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

whether a given instance of CP is bad or not. My **Snakes** case is one where the CP is good, or at least innocuous. But we can easily imagine a variant where the prior belief penetrates perception in such a way that I'm mistaking sticks and roots and such for snakes. Now the CP seems epistemically pernicious. Assuming we don't want to account for this difference by appeal to the obvious but internalistically prohibited facts about reliability, we'll have to find some relevant difference in the inner workings of the perceptual processes. And it's clear that there must be some such difference, since two agents could have their perceptual beliefs penetrated in these two – the good and the bad – ways, even though they're confronted with the same stimuli, in the same environment.

Suppose, for example, that visual object recognition involves the construction of high level representations of the current distal stimuli and the comparison of these representations to stored templates in visual long-term memory; if the current stimulus is sufficiently similar to the template, the visual system registers a match and categorizes the stimulus accordingly (e.g., Marr 1982; Riesenhuber and Poggio 1999). A visual object template for *F*s thus serves as an *F* detector. Compare two scenarios:

> **Bad Snake**: Perceiver BS's expectation of snakes cognitively penetrates her perception by selectively shifting down the match criteria or threshold for her snake template. The "tuning curve" of this snake detector is thereby broadened, and a wider range of stimuli will now trigger it, as the result of CP. Suppose further that top-down consequences of snake detection alter low level visual experience of color, shape, etc.

The result will be a large number of false positives and, presumably, epistemically downgraded experience and belief. BS is walking through the woods with her friend GS, whose penetration works differently.

> **Good Snake**: The tuning curve of GS's snake detector remains highly selective, but the detector is, due to CP, queried early and often in the perceptual process, thus leading to increased – but non-biased – processing of snake-like stimuli. That is, the penetrating expectation influences which information gets selected for further processing (namely, the information potentially relevant to snakes) but not the results of this processing.

The penetration here has no distorting effect,[13] but the result is that GS sees many more snakes than a non-penetrated perceiver would (and more than BS does, although the two form an equivalent number of snake *beliefs*). GS's beliefs remain perfectly well justified. The cognitive process involved in **Good Snake** thus offers one possible mechanism for my earlier **Snakes** case.

At a very coarse grain, BS and GS employ the same process: belief that $p \Rightarrow$ experience that $p \Rightarrow$ belief that p. The finer grain just presented offers one way to see the needed difference between the respectively employed relevant processes, but it's not clear that this would be of any help to inferentialism. The main problem is that there doesn't seem to be any obvious candidate purely doxastic process that's strictly analogous to either the **Good Snake** or the **Bad Snake** process. What, for example, is the doxastic analogue of

13 Talk here about *distortion* is inevitably externalist talk about truth ratios in the current environment; talk about *bias*, on the other hand, is perfectly internalistically acceptable talk about whether one event influences, or changes the probability of, another.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

a stored visual template? If the standard canons of rational belief formation are supposed to sort the relevant processes, these won't be processes that involve primal sketches, visual form templates, and the like.

In short, Siegel wants **Angry Jack** and (I presume) **Bad Snake** to end up in the same category, with **Good Snake** in a different one. But there's no reason to think that an inferential analogy can yield this result if the process types are individuated in the internalistic way required by the theory. If and insofar as Siegel's proposal looks at all promising, it's likely only because she has tacitly been sneaking in illicitly reliabilist factors like distortion and rate of false positives, factors that are allowed neither by inferentialism nor by the internalism it has been developed to defend.

Finally, whatever level of generality Siegel settles on for individuating etiologies, there remains a salient and crucial difference between perception and inference: one factor highly relevant to the epistemic status of an inference is the epistemic status of the premises. An inference that takes unjustified beliefs as premises is therefore a bad inference. Perception is quite different. Suppose the penetrating belief in **Good Snake** is completely unjustified and irrational; still, if it makes GS better at spotting snakes, this irrationality doesn't threaten the epistemic status of the resulting perceptual beliefs. And, of course, even if the penetrating beliefs in **Bad Snake** are well justified, this doesn't make BS's perceptual beliefs any better off. This epistemic dependence on the status of the premises is a central and obvious feature of inferential justification, but there's no corresponding dependence in the perceptual case. Perception doesn't seem very similar to inference, epistemically speaking, after all.

In more recent work (Siegel Forthcoming, in prep.), she argues in further detail for a structural similarity between certain bad forms of inference and certain bad forms of CP, but the view is still highly programmatic, and she doesn't reduce the vagueness and open-ended flexibility of the central notion of similarity enough for us to know whether an elaborated version of the proposal can sort the good cases from the bad, let alone whether it can do so in the right way. This means, of course, that nothing I have argued here can decisively refute her position. The complaints I have raised have thus been merely pessimistic explorations of various ways of trying to make the central notion more precise.

The kinds of process difference that distinguish **Good Snake** from **Bad Snake** are, I think, highly relevant to the epistemic status of the resulting beliefs and relevant to the more general problem of PC. I allow that the different processes are internalistically specifiable, but this doesn't mean that an internalist theory can offer a decent answer to the problem of PC. It is not enough for the internalist to enumerate, in internalistic idiom, the good and the bad processes; one must also say in that same internalistic idiom what it is that all the good processes have in common and that the bad processes lack. Reliability is off the table, and the only other suggestion on offer is that the bad experience producing processes are just those that are similar to bad belief producing processes. There is little reason, I have been arguing, to think that such a suggestion can get off the ground.

Before leaving Siegel's view, let me suggest a very different way to look at a *roughly* Siegel-style inferentialism. Because of the way that both she and I have formulated her view (again, see note 10), I have been focusing on the psychological process, rather than the canons, or norms, of belief formation. Perhaps Siegel's view would be better stated, rather than in terms of structural analogies to belief-forming processes, simply in terms of epistemic norms that govern experience in addition to belief. One such norm would prohibit retention of an unjustified belief, whether the mechanism of that retention

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

is memorial, perseverative, or perceptual. Another norm would prohibit not the retention of a belief, whether by memorial, perseverative, or perceptual mechanisms, but increasing one's confidence in a belief in the absence of any independent reason. Both of these norms – call them "Norm 1" and "Norm 2," respectively – seem to be violated in **Angry Jack**.[14] For reasons already discussed, and to be developed further below, I think that even if we didn't have to worry about individuating processes (either because we've figured out how, or we've reformulated the view in such a way that the issue doesn't arise), the norms governing beliefs simply aren't the norms governing experiences. It is clear that the agent in **Snakes** violates Norm 1, since the penetrating belief is unjustified. But the perceptual belief is justified despite this violation. Norm 2 is violated as well, at least on the assumption that CP prevents the experience from contributing a reason that's *independent* in the operative sense. Siegel could perhaps insist that the reason is indeed independent in **Snakes**, despite the CP, but then it is hard to see how she could say anything different about **Angry Jack**. Also, we can imagine versions of **Snakes** and **Angry Jack** where the agent merely persists in the original degree of confidence; the belief isn't strengthened. There is still an epistemic difference (the snake belief becomes justified and the angry Jack belief doesn't), which Norm 2 does nothing to elucidate.

3. EVIDENCE INFERENTIALISM

I call Siegel's view a kind of "process inferentialism" because it claims that the epistemic downgrade of experience results from the overall process by which that experience comes about. Some methods of belief formation are good-making and some are bad-making, and those same (or analogous) methods of experience formation are also good- or bad-making. A different way to spell out the basic inferentialist idea is in terms of the basing relation; one could claim that an experience can be based on other mental states in much the same way as an inferential belief can be, and downgrade happens whenever the content of the state on which an experience is based is not adequate evidence for the content of the experience. This is "evidence inferentialism," and is quite different from the view that Siegel endorses. For Siegel, what downgrades my experience that p is the fact that it is caused – in a particular sort of way – by the belief that p. But this isn't necessarily a matter of bad evidence, and it certainly doesn't have anything to do with any kind of mismatch between the *contents* of my belief and experience; the two match perfectly, after all (this is part of the problem). Instead, what accounts for the downgrade is the particular way in which the belief causes the experience, rather than any evidential connection between the experience and its precursors.

14 Siegel suggests to me, in correspondence, that the epistemic bad-making feature in at least some CP cases is not the retention of a belief *per se*, but the retention of an *unjustified* belief. She also points out that her original complaint in the **Angry Jack** case was not that Jill believed Jack to be angry, but that her penetrated perception made her *more confident than before* that Jack was angry. Thus, the contents of Norm 1 and Norm 2 come from Siegel, although the reframing in terms of norms and canons of belief-formation, instead of in terms of psychological processes, is my own attempt at a friendly modification of her view. The modification isn't friendly in the sense that I think it's any more likely to succeed – in fact, I think only a genuinely process-linked view has any hope of handling the problem of cognitive penetration – but friendly in the sense that it more naturally accommodates certain moves that she wants to make, while avoiding many of the issues raised above.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

Of course, basing a belief or experience on something that isn't good evidence for it is an instance of a flawed cognitive process. In this way, evidence inferentialism can be viewed as a species of process inferentialism. (I'll follow up on this in section 3.2.) My purposes here don't require process and evidence inferentialism to be mutually exclusive. I think the differences between the evidence inferentialist views I discuss and Siegel's version of process inferentialism will be pretty clear.

The first person to work out an evidence inferentialism is Peter Markie (2005, 2006, 2013). Markie's view rests on the idea that perceptual experiences are in an important sense disunified; a typical experience is actually a complex consisting of separable parts. Some version of this view goes back to Reid ([1764] 1997) and has been endorsed more recently by Lyons (2005, 2009), Tucker (2010), Bengson *et al.* (2011), and Brogaard (2013), among others.[15] A fairly standard version of the disunified view (see Lyons 2015a, Forthcoming) holds that sensations are the low level, modality-specific, viewpoint-dependent aspects of experience; they have nonconceptual content and the rich phenomenology that we normally associate with sense perception. Seemings, on the other hand, are high level, generic, viewpoint-independent conceptual representations of things around us.[16] A *percept* is simply a perceptual seeming (I'll use both terms), though there are also mnemonic seemings and perhaps intellectual seemings as well. To the extent that seemings have a phenomenology other than forcefulness and "felt veridicality" (Tolhurst 1998), it is a bland, rather than a rich one; it is more like the phenomenology of belief than like the phenomenology of sensation. Seemings, however, are not beliefs: it might perceptually seem to *S* that there is a cat in front of her, even though *S* thinks she is hallucinating and thus doesn't believe that there is a cat nearby. Percepts are thus construed as nondoxastic experiential states that make up one element of perceptual experience.[17]

A typical perceptual experience thus consists of two separable experiences: a sensation and a percept.[18] This is what enables Markie to claim that some experiences are properly or improperly based; a seeming is properly based on a sensation when it's the right seeming for the cognizer to have in response to that sensation, when the sensation is – for that cognizer – adequate evidence for the content of the seeming. In **Prospectors**, for example, the two agents have the same sensation when they look at the lump of gold: it is a small,

15 Including perhaps, Lamme and Roelfsema (2000), Pylyshyn (2003), and Raftopoulos (2009), although these authors tend to focus on the discontinuity of perceptual processing, rather than the disunity of perceptual experience. And of course, some kind or other of a sensation/perception distinction has a long and respectable history in psychology.

16 It is often pointed out that 'it seems to *S* as if *x* is *F*' could be interpreted in a number of importantly different ways: there is a comparative sense of 'seems', a "hedging" sense, etc. The view currently under discussion holds that one sense of the locution serves to pick out a distinctive type of mental state, which we can then call a "seeming".

17 Markie treats the sensation/percept distinction fairly quickly, and it is unclear just how much of this "standard version" he would be willing to embrace. Whatever differences there are, however, between Markie and other proponents of the disunified view, they shouldn't make any difference to the larger issues of this paper. If we think of visual sensations as the introspectible products of earlier visual processing and visual percepts as the introspectible products of later visual processing, as I think we should, then a division of experience into just two seems crude and simplistic (Lyons 2009). Nonetheless, I think the two-part view is a rather harmless idealization.

18 It is perhaps a bit odd outside of epistemology to treat 'experience' as a count noun. Nothing of substance hangs on it; the only substantive claim embedded in it is that the two aspects of perceptual experience (sensation and seeming) are counterfactually separable.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

shiny, yellowish, object with a certain highly determinate shape and a certain orientation and distance. And they also have the same percept (that it's gold), but that's the problem. It looks (/seems) like gold to Gus because of his expertise, but it only looks like gold to Virgil as the result of wishful thinking. Virgil's seeming experience thus comes about in an epistemically improper way, a way that renders it unsuited for conferring justification on the relevant perceptual belief.

Markie thus claims that the etiology of at least some experiences determines the epistemic potency of that experience. His view is this:

> A perceptual seeming that *p* is epistemically appropriate iff it results from the agent's perceptual know-how.

The difference between Gus and Virgil is that Gus has perceptual know-how linking the sensation to the seeming, rendering the former good evidence for the latter, while Virgil does not. What's wrong with Virgil's seeming, therefore, is not the fact that it's caused by a prior belief; rather, it's that the seeming fails, due to the lack of know-how, to be properly based on the corresponding sensation. The penetrating belief isn't a positively bad-making feature; what it does is causally enables the sensation-seeming transition in the *absence* of the relevant good-making feature: know-how.

An evidence inferentialist is, as far as I can tell, more or less committed to a disunified view of perception; otherwise it is hard to see how perceptual experiences might be *based on* other mental states.[19] Evidence inferentialism per se, however, is neutral with respect to what makes a certain sensation-seeming transition appropriate. One nice feature of Markie's basic know-how proposal is that it can be taken in two importantly different directions, depending on whether he wants to insist that the evidence relation is contingent. In earlier work, Markie (2005, 2006) understood perceptual know-how in terms of procedural knowledge for how to form perceptual beliefs. This renders the evidence relation contingent, because this background knowledge – in part because it is procedural, rather than declarative – does not serve as additional evidence but, rather, serves to determine what counts as evidence (for that agent) of what. This procedural "knowledge" needn't be justified; it needn't be propositional; it needn't be introspectively accessible, etc. The knowledge here plays a role akin to the role that reliability plays for indicator reliabilism, except of course that the evidence-determining factor is internal, at least in the mentalist's sense.

His new view (Markie 2013), on the other hand, understands know-how in terms of factual knowledge about how things look.[20] On this view, the background knowledge

19 Evidence inferentialism in general requires that there be some other mental state that the epistemically evaluable experience is based on. Whatever mental state is first in this chain will not be evaluable, at least not in a way that evidence inferentialism can account for. The traditional view and its default epistemic potency will hold for these states. Perhaps this isn't a problem if, for instance, it is only some unconscious early visual states that are unevaluable. But it is clear that Markie at least (and, as we'll see, McGrath) needs the basis states to be introspectively available, because he wants the basing to be something that the perceiver does, rather than something done for/to her by her subpersonal mechanisms.

20 The relevant factual knowledge is dispositional knowledge, or *ex ante* knowledge. This refinement matters a lot for certain important purposes, as I will explain in section 3.2 below, but it doesn't matter for the present purposes.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

that constitutes the agent's know-how does play an explicitly evidential role. This newer view is therefore compatible with the claim that all evidential relations hold necessarily. On the older view, whether a sensation *s* is evidence for *p* is determined by the contingently possessed procedural knowledge of the cognizer. On the newer view, *s* is only *part* of the agent's evidence for *p*, the rest of the evidence being supplied by the background knowledge.

So "contingent evidence inferentialism" holds that evidential relations are contingent and that what explains downgrade in CP cases is that the high level experience, the one on which the perceptual belief is based, is not itself based on another mental state that is good evidence for that high level content, *for that agent*. "Necessary evidence inferentialism" is similar, except it holds that evidential relations are necessary. To offer non-bullet-biting treatment of CP cases, it needs to hold that the sensation is only *part* of the agent's evidence relevant to the epistemic status of the high level seeming.

The contingent and necessary versions of evidence inferentialism are importantly different. I'll discuss them separately.

3.1. *Contingent evidence inferentialism*

Although Markie was the first to endorse contingent evidence inferentialism, the purposes of this paper are better served by focusing on a version of the view that has been recently defended by Matthew McGrath (2013a, b), who combines elements of Siegel's and Markie's views. McGrath uses Markie's appeal to the disunified theory of perceptual experience to constrain and give a clearer specification of Siegel's inferential analogy. Much of what I say about McGrath's view will apply to Markie's (2005, 2006) view as well.

McGrath sees some experiences as receptive and others as quasi-inferential. An experiential state *e* is quasi-inferential just in case *e* comes about as the result of the agent's treating some other input mental state of hers as adequate support for *e*. An experience is receptive otherwise. McGrath thus thinks that some – though perhaps not all – seemings are quasi-inferential; all sensations, presumably, are receptive.[21] This distinction then allows us to distinguish what is *given* to us in perception and what is of our own making. McGrath claims that whatever is given to us is automatically thereby capable of conferring justification on appropriately related beliefs, just as the traditional view has it; what we have made, however, we must have made properly. That is, receptive experiences are not subject to epistemic appraisal but enjoy a default epistemic potency; quasi-inferential experiences, on the other hand, are subject to epistemic appraisal on the basis of their etiologies.

In particular, McGrath thinks that some transitions from sensations to percepts involve a kind of "free enrichment," very much like jumping to conclusions (2013b: 235–6). Virgil's experience as of gold is such an unlicensed enrichment of his experience as of a yellowish lump. More generally,

> a quasi-inference is downgraded iff an (actual) inference with the same contents would be a bad inference for that subject.

21 McGrath doesn't use the 'sensation/seeming' terminology, calling them all "experiences" or "seemings" apparently interchangeably. I think the 'sensation/seeming' terminology is helpful, so I am translating McGrath into it. Caveats from note 17 apply here as well.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

Relativization to the subject is imperative if the evidence relation is contingent; Virgil and Gus presumably make the same sensation-seeming transitions, although the two contents are appropriately matched for Gus and not for Virgil, because of their respective background knowledge. This promises to handle the **Angry Jack** case as well; if Jill's sensation isn't affected by her belief that Jack is angry, then it is presumably one that should and normally would make it seem to her that he had a neutral facial expression. By having instead the seeming that he is angry, Jill is "jumping to conclusions," although at the level of experience rather than belief.

Like Siegel, McGrath views downgrade in terms of an analogy with belief formation. But his appeal to the sensation/percept distinction solves a central problem for her view by sharply delimiting the relevant analogy. As on Markie's view, the crucial element is the transition from sensation to percept, so it is an analogy of contents, not processes. This obviates a number of the worries raised above for Siegel's view, as these concerned the vagueness and open-endedness of the relevant analogy. However, a whole new set of problems is produced by restricting everything to the sensation-percept transition.

First of all, the theory only addresses the problem of CP for high level experiences, but there may well be CP of sensations as well. We could easily modify the two central cases to stipulate that the CP affects the outputs of early vision.[22] Wishful thinking might cause the rock to look more yellow to Virgil than it really is; Jill's belief that Jack is angry could affect the apparent slant of his eyebrows, and so on. Because these are low level features, the kind that figure in sensations, they are not plausibly quasi-inferential, so McGrath must claim that the beliefs are perfectly well justified after all. One might reply that this isn't so bad from an internalist perspective, since the rock really does look yellow to Virgil, and he's not in a position to know that this is only true because of his wishful thinking. But the same could be said about its looking like gold to Virgil, and McGrath's theory is explicitly designed to allow him to claim that Virgil is unjustified in believing the rock to be gold. The kinds of intuitions that support access internalism and phenomenal conservatism might well mitigate the counterintuitive nature of denying that there is epistemic downgrade in cases of CP of low level perception; but if so, then they render a theory like McGrath's otiose, by mitigating the counterintuitiveness in the high level cases as well. McGrath's theory is supposed to offer an alternative to access internalism and phenomenal conservatism. All this, of course, applies to Markie's version of inferentialism as well.[23]

Second, the proposal is straightforwardly nonvacuous only if there are sensation-to-seemings transitions with the same contents as belief-to-belief transitions. But this is true only if some sensations have the same contents as beliefs. Yet it is very plausible

22 CP of sensations is more controversial than CP of percepts. As stated from the outset, I'm not very worried about whether CP is actual but am interested in the relevant hypotheticals. However, there is some reason to think that CP of low level features, like color, is actual (Hansen *et al.* 2006; Levin and Banaji 2006; Macpherson 2012).

23 Perhaps there is another way of motivating the claim that penetrated sensations don't get downgraded, while penetrated percepts do. The agent might be responsible for the percept in a way that she is not responsible for the sensation, for example (McGrath suggests something like this in personal communication), so deontological or other considerations might justify treating sensations and percepts differently in this regard. As will become clear in my ensuing discussion of the modularity of perception, I don't believe there is any significant sense in which the agent is involved in the construction of percepts but not sensations.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

that although beliefs have fully conceptual content (I take this to be partly definitional of 'conceptual'), sensations do not.[24] It is natural to read McGrath's proposal as claiming that the quasi-inference from E1 to E2 is downgraded iff (i) there are beliefs B1 and B2, with the same contents as E1 and E2, respectively, and (ii) the inference from B1 to B2 is a bad inference for the agent. If so, then no perceptual experiences will be downgraded, because (i) won't be satisfied. Alternatively, he might have a counterfactual reading in mind: E1 to E2 is downgraded iff *were* the agent to have beliefs B1 and B2, with the same contents as E1 and E2, respectively, then the inference from B1 to B2 would be a bad inference for the agent. Evaluating such counterfactuals presumably involves holding the contents of E1 and E2 fixed, and turning to possible worlds where the agent's conceptual repertoire enables beliefs with these same contents. But if my counterpart has vastly different conceptual capacities than I do, it is unclear why the epistemic status of his inferences has anything to say about the epistemic status of my quasi-inferences.[25]

Most importantly, McGrath's view – and the same is true of Markie's view – seems to fit quite poorly with the modularity of perception.

By 'modularity' here I mean that perception is fast, automatic, determined by domain specific principles proprietary to the perceptual system in question, *relatively* immune to cognitive penetration from occurrent beliefs and goals and the like, largely unconscious in its inner workings, and performed by systems that are relatively independent and isolable from other cognitive mechanisms. Although the term 'modularity' and the seminal statement of the view just sketched are due to Fodor (1983, 1985), I want it to be clear that the version I am endorsing is a more moderate, "weakly modular" version of the view (for more on the relevant understanding of weak modularity, see Lyons 2001, 2015b). I don't, for example, want Fodor's (1983) famous nine diagnostic features to be rigidly imposed and *definitional* of modularity; if they were, the issue of CP wouldn't arise except as a purely hypothetical, academic, and uninteresting issue. Some writers (e.g., Pylyshyn 2003, Raftopoulos 2009) hold that early vision (roughly, up to and including Marr's 2 1/2-D sketch) is strictly immune to penetration from cognition, although this view is controversial. I'm thinking of perception as comprising not only early vision, but also later stages of processing, where distal objects are seen as falling under relatively high-level, commonsense concepts, such as *cat*, *cow*, and perhaps fine-grained concepts, like *red oak*. Even the staunchest proponents of modularity allow that perception at this late stage is cognitively penetrable. However, it is nonetheless obvious and uncontroversial that perception even at this late stage is *relatively* immune to CP (hence the above hedge): we can't just see whatever we want (or fear, etc.), and being convinced that the thing I'm looking at isn't really a cat (but is, say, a cleverly disguised Martian robot)

24 The nonconceptual nature of sensations is something most authors agree on (e.g., Crane 1992; Tye 1995; Peacocke 2001; Burge 2010, etc.), and I think it is true (Lyons 2009). Obviously, it is too large an issue to try to argue here.

25 An anonymous referee suggests that McGrath could modify his proposal to hold that a quasi-inference is downgraded iff an (actual) inference with *matching* contents would be a bad inference for that subject, where the matching relation is easier to satisfy than the identity relation. Some such matching relation is required anyway by those who think that perceptual beliefs can be justified by experiences with nonconceptual (and thus nonidentical) contents. One might try to unpack matching in various ways: in terms of conceptualizations (Burge 2010), or similarity (Siegel and Silins 2015), or canonical correspondence (Peacocke 2004), or even proper functionalist or reliabilist terms. I don't have the space here to seriously pursue all or even any of these possibilities.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

won't make it stop *looking* like a cat. If CP exists, and this is a highly contentious, difficult, empirical issue, it is significantly delimited and circumscribed.

The modularity of perception, however, leads to two pervasive, central, and deep facts about perception. The first is that perception quite generally involves moves from low- to high level states that are rationally underdetermined, in the sense that the relevant modules help themselves to "assumptions" that the larger agent might well not share or have a right to, e.g., assumptions about the direction of light, about how light propagates in the current environment (this will include general principles of optics as well as particular assumptions about the current optical medium), about prior probabilities of various arrangements of distal stimuli, etc. Thus, I don't have to know that things nearby are lit from above for them to (properly) look convex to me; I don't have to know how aerial perspective works or know the current meteorological conditions for objects to look a certain approximate distance from me; etc. The second deep fact follows partly from the first and is that high level perceptual contents are therefore "given" to us in the relevant sense. I'm not *doing* anything in having objects look convex to me, even though they look convex on the basis of lower level shading information; it's not *me* that infers distance from a difference in contrast and saturation, it's simply that the one object looks farther away than the other. The Helmholtzian view of perception as a kind of unconscious inference is almost surely true, but it is not inference *on the part of the perceiver*.[26]

In virtually all cases of perception, good and bad, there is a significant gap between the low level sensation and the high level percept, or seeming (Lyons 2009). The gap is generally such that, if an ordinary agent, not deeply trained in the visual arts or perceptual psychology, were to receive an accurate and highly detailed description of the content of sensation, she would be unable to do much more than guess at the content of the corresponding high level, seeming, information. Virtually all such transitions would amount to leaping to conclusions. McGrath, of course, is aware of this and aware of how it might undermine his general theory, so he explicitly invokes a distinction between the personal and the subpersonal. Since it's not *me*, but my visual system, that's making shape-from-shading inferences, the outputs/conclusions count as receptive and are justification conferring even if the visual system is drawing inferences that I wouldn't be justified in drawing. I think this claim about receptivity and givenness is exactly right (Lyons 2009, 2014), but I don't think it helps McGrath much, because *all* genuinely perceptual seemings are going to turn out to be receptive, and he no longer has an account of the epistemic downgrade resulting from CP.[27]

For example, how exactly would McGrath handle **Prospectors**? He claims that Virgil's gold seeming is an improper quasi-inferential seeming. But what about Gus, the expert? Is

26 This, I think, is the explanation for the asymmetry mentioned toward the end of section 2, between perception and (real, ordinary) inference vis-à-vis the epistemic status of the input/premise beliefs. Because perception is modular in the sense described here, the inputs to the perceptual process are screened off from making any evidential contribution; their epistemic status *per se* doesn't matter.

27 Some authors use 'personal' and 'subpersonal' as more or less synonymous with 'conscious' and 'unconscious', respectively. In contrast, the way I am using the terms, and the way I think McGrath uses them, 'personal' means 'attributable to the agent' and 'subpersonal' means, roughly, 'attributable only to parts of the whole agent'. I might have certain unconscious beliefs or knowledge or make certain unconscious assumptions, but so long as those beliefs, assumptions, and knowledge are *mine*, they're person-level phenomena. If it is not me, but my visual system, that assumes that two retinally adjacent points are approximately equidistant, for example, then this is a subpersonal phenomenon.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

his gold seeming quasi-inferential, but proper for him, or is it receptive? If we really insist that quasi-inferences must be *person*-level doings, then it's not plausible that Gus's seeming is quasi-inferential after all, at least not in the standard case of perceptual expertise. In the standard case of perceptual expertise, things really do look to be gold; 'gold' is a modular output, something that is given to – indeed, forcefully thrust upon – the perceiver. Person-level agency enters the picture most clearly when the perceiver *rejects* the modular output; this output is not something she agentially *constructs*. If Gus's case is supposed to be a typical case of perceptual expertise, then his changing his mind about what gold looks like won't change the fact that the rock still looks to be gold, although of course now he won't believe that it really is. Maybe we can imagine a case where Gus is agentially involved in its seeming to him that the rock is gold, but insofar as we do so, I submit, we are imagining a post-perceptual seeming of some sort, rather than a perceptual seeming.

I look across the room, and it seems to me that the way my toddler daughter is climbing on the furniture is dangerous. This seeming of danger is presumably not a genuinely perceptual seeming, in the way that its seeming to me that my daughter is climbing on the furniture presumably is perceptual. Maybe this danger seeming is something that I have, spontaneously although at the personal level, constructed from the perceptual seeming and some inchoate background knowledge, or folk physics, or induction, or something. Perhaps there is a version of **Prospectors** where Gus's gold seeming is like my danger seeming, where Gus has some kind of extra-perceptual, or post-perceptual, expertise that's responsible for his gold seeming. But in such a case, it doesn't literally – visually – *look* like gold to him; this gold seeming isn't a gold *percept*; for that, we need the seeming to result from a specifically *perceptual* expertise. (Of course, we might still say, using some other, nonliteral, sense of 'look', that it looks like gold to him, just as we might say that my daughter's situation looks dangerous or that it looks like my bicycle has been stolen.)

Setting aside the previous objection concerning the nonconceptual contents of sensations, I think McGrath's view could use this kind of move to handle either variant of **Prospectors**, as far as Gus is concerned. The seeming is either perceptual and therefore receptive, in which case it gets default potency, or it is quasi-inferential but, we might suppose, one whose corresponding inference would be justified in virtue of Gus's post-perceptual expertise. The problem is that this view doesn't deal with Virgil in the right way. If Virgil's seeming is a cognitive, post-perceptual seeming, like my danger seeming, then McGrath can claim that it's quasi-inferential and that Virgil wouldn't be justified on the basis of the corresponding inference. But then it's not a case of cognitive penetration of *perception*, and we haven't seen the quasi-inference view offer a solution to the problem of CP. We can construe Virgil's seeming as a genuinely perceptual seeming, but then it has to be a receptive seeming, and McGrath has to count it as having the usual epistemic potency.[28]

McGrath needs to lean hard on the personal/subpersonal distinction in order to avoid the result that virtually all ordinary perceptual seemings are subject to epistemic downgrade, since the inference that corresponds to the typical sensation-seeming transition generally involves a serious leap. But if CP is genuinely CP *of perception*, then this move

28 One needn't accept my subsumption of the perceptual to the modular to embrace this argument; all we need is that *some* perceptual beliefs are modular in the current sense and nevertheless subject to an epistemically pernicious form of CP.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

works against him, for CP then operates at the subpersonal level, affecting one's receptive seemings. McGrath's theory then only accounts for the epistemic downgrade that results from the cognitive penetration of cognition, not the cognitive penetration of perception.

I've raised three objections here to McGrath's version of Contingent Evidence Inferentialism: First, it only countenances downgrade for cognitively penetrated seemings, while there should also be downgrade for cognitively penetrated sensations. Second, McGrath's account of downgrade in terms of an inference with the same content founders on the fact that sensations don't generally have conceptual content and thus won't have corresponding inferences. Third, it fails to do justice to the modularity of perception, either treating perception as too much of an active, person-level phenomenon, or by restricting the scope of downgrade to seemings that aren't literally perceptual.

The first and third of these objections apply to Markie's (2005, 2006) version of the theory as well. Because the evaluable transition for Markie is the one from sensation to seeming, downgrade can only occur at the later stage. And although Markie is less explicit about it than McGrath, his appeal to the agent's knowing how does seem to require that this know-how is a person-level, rather than a subpersonal, phenomenon. This is much more explicit in his more recent (2013) work, to which I now turn.

3.2. *Necessary evidence inferentialism*

Markie's older, contingent evidence, view held that know-how was a matter of the agent's having internalized procedural knowledge for making certain sensation-seeming transitions. His new (2013) view about perceptual know-how is much more intellectualist, for it requires the agent to have *ex ante*, prima facie justified *beliefs* about how things look. More exactly, the agent needs to have evidence/justification for beliefs about how things look; she needn't actually have those beliefs. Thus perceptual beliefs will still be epistemically basic, at least in the sense that their being justified doesn't require the agent to have other (*ex post*, doxastically) justified *beliefs*, although they would not be basic in the sense of being justified independently of evidence for other beliefs.[29]

As before, Markie holds that

> A perceptual seeming that p is epistemically appropriate iff it results from the agent's perceptual know-how.

29 McGrath's new (Forthcoming) view is again like Markie's new view, though again different in important ways. McGrath now claims that experiences immediately justify beliefs about the (objective) looks of things, and these beliefs, when accompanied by justified beliefs about how things look, justify ordinary "perceptual" beliefs about middle sized objects, beliefs like 'x is red' or 'there's a cat in front of me', etc. McGrath's emphasis on how things look is not being pressed into the service of *inferentialism*, because it isn't supposed to explain how *experiences* are downgraded. Rather, the relevant "perceptual" beliefs (e.g., 'there's a pig') are quite strictly inferential, i.e., nonbasic, i.e., mediately justified, for McGrath. Because he is no longer defending what I'm calling inferentialism, McGrath no longer needs the sensation-seeming distinction. CP is no longer viewed as something akin to bad inference; it simply *is* a particular kind of bad inference. The role both theories accord to beliefs about how things look, however, is similar enough that the objections I am about to raise for Markie's new view will mostly apply to McGrath's as well. I won't spell these out in detail, because space is limited and my target here is inferentialism, a view that tries to accommodate the epistemic downgrade of perceptual beliefs as a result of CP, while still allowing that these perceptual beliefs are epistemically basic.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

Now he understands know-how as follows:

> An agent has perceptual know-how regarding *p* iff there is a sensation *s* such that the agent is disposed to have the seeming that *p* in response to *s*, and is so disposed in virtue of having evidence for thinking that *s* indicates *p*.[30]

Know-how has thus gone from being a kind of procedural knowledge, which is a matter of having adopted and internalized certain dispositions, to a kind of declarative, i.e., factual, knowledge, which is a matter of being in good epistemic standing vis-a-vis certain propositions.[31]

To repeat an objection from earlier, neither Markie's old nor new view offer any theory about downgrade for cognitive penetration that occurs at the level of sensations. But this isn't an objection that will worry Markie much. Indeed, he claims (2013: 266) that if Jill's unjustified belief infects her sensations, then she is justified in believing that Jack is angry, after all. Presumably he will say the same about BS in **Bad Snake**. It's a bit surprising that his preferred form of internalism is still mentalism (2013: 248), since his view now seems to be compatible with access internalism, at least on the assumption that the evidence that constitutes know-how, and its status as evidence, are reflectively accessible.

More important and problematic is the role that factual knowledge plays in justifying perceptual belief. Markie claims not only (a) that every time someone has perceptual know-how in the epistemically relevant sense, she also has some person-level, factual knowledge/justification (albeit *ex ante*); he also claims (b) that the disposition toward a certain sensation-seeming transition must hold *in virtue of* the agent's having evidence for beliefs about what things look like; that is, the *ex ante* justification must ground the disposition to form certain seemings, in order for those seemings to be epistemically appropriate. Thus, framing things a bit crudely, we have a $KH \supset K$ principle (a) and a *KH because K* principle (b). Both strike me as far too strong, again because of the modularity of perception.

Let us first consider $KH \supset K$. When I was in college, I drove a 1962 Chevy Impala. I know the model quite well and could always – and still can, I believe – distinguish it from similar-looking cars. The most similar is a '63 Impala, but I can immediately tell whether I'm looking at a '62 or a '63. So I have whatever perceptual know-how is necessary for justified perceptual belief and for epistemically appropriate perceptual seemings. However, my person-level, factual, declarative memory for these matters is a bit fuzzy. I couldn't tell you with any confidence, for example, whether the distinctive, round, bullet-like taillights came in groups of three on each side or in groups of two. In some sense, then, I don't know what a '62 Impala looks like from behind. This is not because what the car looks like outstrips my linguistic and/or conceptual resources; it's simply because I can't remember without the car in front of me whether it has four taillights or six. I don't, in

30 See Markie (2013: 264). Markie does not explicitly endorse the claim that evidential relations hold necessarily, but because he insists that the sensation is only part of the evidence for the perceptual belief, with the rest of the evidence being supplied by evidence for the conditional linking the sensation to the perceptual belief, his view is compatible with that claim. This is why I treat his view under the heading of necessary evidence inferentialism.

31 In what follows, I will sometimes write about factual or declarative *knowledge* rather than justified belief, simply because the former term is less cumbersome. None of my argument will rely on those features (truth, unGettieredness, etc.) that distinguish the two.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

any straightforward sense, possess evidence sufficient to justify any belief I might have about which sensations indicate the presence of a '62 Impala viewed from the rear.

It should be obvious here that I am *not* assuming that declarative/factual knowledge must have a linguistic or discursive format. I'll allow – indeed I'll soon come close to insisting – that much such declarative knowledge has an essentially perceptual format. My point here is not about the representational format of the knowledge; it's that I can't call it up without being cued in appropriate ways by the environment.

Now suppose further (as I'm pretty sure is true, but it doesn't matter, so long as it's possible) that whichever of these two it is, six taillights or four, it is also true that *if* I were to be presented with a photo of the real thing and a photoshopped version with the other number, the real one would look right to me and the photoshopped one would look wrong. This indicates that there is a quite robust sense in which I do know what a '62 Impala looks like, despite the fact that there is that just noted sense in which I don't. This kind of case, I take it, is quite common, and I think it is clear that the resulting perceptual beliefs are perfectly well justified even though I'm lacking the declarative knowledge. If this is right, then it follows from the *generic* know-how view (the one common to Markie's earlier and more recent views, the one that links justification and know-how) that I possess the epistemically relevant perceptual know-how, even though I lack the declarative knowledge. That is, I have KH, even though I lack K.

One might object to this argument on the grounds that it trades on something akin to the distinction between free recall and recognition in studies of memory. Recognition ("Did you see this earlier? Yes or no?") is generally much easier than free recall ("What did you see earlier?"). Recognition implies memory and thus knowledge even in the absence of recall, so if I can recognize one of the taillight configurations as the correct one, then perhaps I do know after all how many taillights are on a '62 Impala, even if this knowledge isn't available for free recall. Thus, it's not a case of KH without K after all.

This objection, however, misses the fact that *right now*, by hypothesis, I have the KH; I'm currently disposed to make the relevant sensation-seeming transitions in a justification-conducive way. But I don't right now have adequate evidence for any beliefs about what sensations indicate the presence of a '62 Impala. Given the way that Markie's view has to unpack the $KH \supset K$ principle, therefore, that principle is false.

This is compatible with the possibility that every *exercise* of KH puts one in a position to declaratively know the relevant sensation-seeming mapping. But this brings us to the next point I want to make, which is that even in cases where we do have factual, declarative knowledge of what things look like, it is far from obvious that this knowledge grounds our perceptual dispositions. I suspect, instead, that a great many cases of KH & K are not cases of KH because K, but more likely the other way around.

I know what my mother looks like, and let's suppose (I think fairly implausibly) that I have enough declarative knowledge of what she looks like that I would be able – if only I were a better artist – to paint a convincing likeness of her from memory. This still wouldn't give us any reason to think that my ability to map sensations to perceptual seemings is dependent on my declarative knowledge. More plausibly, in this sort of case, the declarative knowledge depends on the ability to map sensations to seemings.

My argument proceeds in two stages: first I want to argue that when we do have declarative knowledge, it depends, at least in some cases, on a capacity for forming perceptual imagery; second, this imagery capacity depends on being disposed to form certain

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

perceptual seemings in response to certain sensations. That is, it's K because KH, not KH because K.

As for Stage 1, it is clear that we can sometimes have a certain amount of factual/declarative knowledge of how things look *without* relying on an ability to visualize (i.e., to form perceptual imagery). If you tell me that your car is red, then I have factual knowledge about how your car looks, which doesn't appear to require any capacity for visual imagery (assuming that, e.g., having the concepts *car* and *red* doesn't require that capacity). This knowledge might, in certain highly constrained conditions, allow me to visually identify your car (although I would count the belief as perceptually based, rather than a genuine perceptual belief). But for other cases, our declarative knowledge seems to be bound up with the capacity for imagery. A favorite textbook example involves answering the question how many windows there are in one's current home, to which subjects respond by imaginatively walking, in their mind's eye, along the periphery of their home (Shepard 1966). The idea seems especially plausible if we're interested in the knowledge that would suffice to make a particular percept epistemically appropriate. For Markie, this is a matter of having *ex ante* justification for beliefs of the form 'if $\sum$ then *p*,' where $\sum$ is a collection of low level property instantiations represented by sensations, and *p* is a proposition concerning the external environment, like 'x is gold' or 'that's my mother.' To the extent that I can even approximate this kind of knowledge of my mother's appearance, for example, I have to call up a picture of her in my mind's eye. Of course, there are a few general facts about her appearance that I could recall from memory without using imagery (e.g., hair color, eye color, whether she wears glasses), but these facts, facts that I could easily exhaustively convey verbally to someone else, wouldn't be nearly sufficient to pick her out of a police lineup.

And it's not surprising that we would be dependent on imagery for certain declarative knowledge about how things look. The visual system needs to rely on stored knowledge to identify things, and it is more efficient to keep this knowledge coded in a visual format, especially since the capacity for visual imagery makes this knowledge more generally available whenever needed. If I frequently had reason to answer questions about my mother's facial features, or draw out some of their deductive consequences, it might be worth the time and effort and storage to separately and redundantly code a number of relevant facts in some amodal format, but as things are, leaving most of the information in the visual format seems to work just fine.

Because much of this knowledge hasn't been separately coded, a good deal of our declarative knowledge seems to depend counterfactually on our capacities for imagery. Were I to lose the ability to visually imagine my mother's face, I would not have sufficient factual knowledge about what she looks like to reliably or justifiedly visually identify her.[32] One might prefer the slightly more modest claim that I would then no longer *have access to* sufficient factual knowledge to reliably or justifiedly visually identify her, though the knowledge might still exist. We sometimes think this way about items encoded in long-term memory that we can't retrieve: the knowledge is still there, even though we

32 If this kind of declarative knowledge is stored in a partly perceptual format, then the knowledge would constitutively depend on a capacity for forming perceptual imagery. If (as some might hold, but I am not assuming here) any declarative knowledge must be coded in some more discursive format, then the dependence will be merely causal. All my argument requires is that the declarative knowledge, whatever its nature, depends counterfactually on the ability to form perceptual imagery.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

can't access it. However, I think that in both these cases, we would rightly deny that *the agent* knows whatever it is that she can't access. Knowledge that exists in me but that I can't access is knowledge I don't have.

This brings us to Stage 2. The counterfactual dependence just described strongly suggests that the relevant declarative knowledge depends on the ability to visualize and not the other way around, even in cases of ordinary, justified perceptual belief. If so, then one could have the imagery ability without having the declarative knowledge (if, for instance, something went wrong in the process of forming beliefs by reflecting on the results of the imagery process). Now, if the ability to visualize is materially sufficient for having the disposition to make the appropriate sensation-percept transitions (i.e., if everyone who can visualize also happens to be able to map sensations to percepts), then one could have that disposition without having the declarative knowledge, and this shows that the disposition doesn't depend on the declarative knowledge, contrary to Markie's view. (For that matter, the above argument for KH without K shows this.) But I want to endorse the even stronger claim, that the declarative knowledge depends on the disposition.

I think the stronger claim follows from the fact that the formation of perceptual imagery involves the offline simulative use of the ordinary perceptual systems (Currie 1995; Wheeler *et al.* 2000; Slotnick *et al.* 2012). To produce visual imagery of an *F* is to use the information structurally encoded in the visual system to token, albeit in an offline manner, the sensations one would have while looking at an *F*. What it is for us to be disposed to map certain sensations to certain percepts is to have a perceptual system wired up in such a way that it does, in fact, map these sensations to those percepts. The system may have gotten that way as the result of perceptual learning, or having internalized some norms (Markie 2006), or it might have been innately wired that way. Thus, my ability to form perceptual imagery depends on my having certain now-wired-in dispositions to map sensations to percepts. And since my having declarative knowledge depends, as already noted, on having the capacity for imagery, this declarative knowledge depends – by transitivity of dependence – on the mapping disposition.

I will assume that if having the declarative knowledge depends on having the mapping disposition, then having the mapping disposition does not also depend on having the declarative knowledge. Thus, it is not the case that KH because K.

More explicitly:

(1). For some instances of justified perceptual belief, one's declarative knowledge of how things look depends on one's ability to form perceptual imagery.
(2). The ability to form perceptual imagery depends in turn on being disposed to form certain perceptual seemings in response to certain sensations.
(3). So, some declarative knowledge of how things look depends on the agent's having certain sensation-seeming mapping dispositions.
(4). If the declarative knowledge depends on the disposition, then the disposition doesn't depend on the declarative knowledge.
(5). Therefore, in some instances of justified perceptual belief, our being disposed to produce certain perceptual seemings in response to certain sensations does not depend on our having declarative knowledge about how things look.

And this conclusion directly contradicts Markie's theory.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

None of this will be very surprising if we take at all seriously the modularity and consequent autonomy of perception. To the extent that perceptual systems are modular, they operate independently of the beliefs, goals, and expectations of the larger organism. CP, on this view, is a relatively rare phenomenon. But the causally mediating role that Markie accords to perceptual know-how, construed as agent-level declarative knowledge, requires that perception be consistently and quite pervasively penetrated by cognition. It requires that the perceptual systems simply wouldn't be able to do their jobs of constructing percepts in response to sensations, were it not for person-level knowledge of how things look. This is exactly what the modularity thesis denies.

Here's one way to summarize what I've been arguing: for a substantial range of cases, much of *my* knowledge of what Fs look like is parasitic on *my visual system's* knowledge of what Fs look like, although Markie requires it to be the other way around. You could talk me out of my person-level beliefs about what Fs look like, but this wouldn't change my visual system's knowledge, due to the modularity of perception. If there are two bodies of declarative knowledge here, the one that's connected to procedural knowledge for how to map sensations to percepts is the subpersonal knowledge, not the person-level knowledge. Again, Markie's theory needs it to be the other way around.

Markie could, of course, deny that perception is at all modular. Doing so, however, would be taking on a very substantial empirical burden, especially given the highly heterodox nature of the empirical assumptions implicit in Markie's view. It's not just that the received view embraces a fair degree of modularity for perception; the *kind* of failure of modularity that Markie would be proposing would be a novel and surprising kind. Not even such a staunch and thoroughgoing antimodularist as Andy Clark, for example, espouses this particular kind of cognitive penetration. On Clark's (2013) predictive coding view, perception is highly influenced by the agent's expectations about what's present in the environment. These expectations then filter down the processing hierarchy to be tested against incoming sensory stimulation. But the perceptual know-how that causally mediates this filtering is very much procedural, rather than declarative. The visual system, on Clark's view, is highly plastic and susceptible to learning, but the learning it undergoes is still a slow, laborious process that requires repeated real-world interaction with objects and has little to do with the agent's occurrent, potentially transient, beliefs about the looks of things. The only transient states that play a significant role are the expectations about the layout of the distal environment. My changing my beliefs about how things look would not, on Clark's view, change the internal structure of the perceptual systems and thus not affect the mapping from sensations to percepts. Even on Clark's view, perception is presumably modular in this regard. Not even the most vigorous proponents of CP claim that it is pervasive in the particular way Markie would require it to be.

A substantial difference between me and Markie concerns a difference in how we think as yet unsettled empirical issues will turn out. I believe that, as a matter of empirical fact, actual cases like Gus's belief typically involve the training of a modular visual system to produce 'gold' outputs and that cases like Virgil's involve an untrained module producing something like 'yellowish rock', from which Virgil quickly, automatically, and improperly infers 'gold.' Markie's theory requires that Gus's case is like I've just described Virgil's, except of course that the inference (or transition to seeming) is proper, given Virgil's background evidence. Thus, Markie and I disagree about how far into the processing streams perception remains modular in the relevant sense; I think it remains modular quite a bit later in processing than he does. Recall that the relevant sense of modularity here requires

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

merely that perceptual processing occurs autonomously, in a way that is not synchronically causally dependent on person-level beliefs about how things look, or on person-level evidence for such beliefs.

It is important to clarify that the way I am understanding Gus's case, he is forming a perceptual belief to the effect that the thing in front of him is gold; we might suppose that Gus has a Mentalese term that means *gold* and that he tokens in response to the current stimulus. This is *not* the belief that the thing in front of him is *called* 'gold'; this belief involves a quite different Mentalese term, one that refers to the *word* 'gold'. I expect that this latter belief, even if it arises spontaneously and automatically, is extra-modular and thus has person-level beliefs and inferences in its causal history. If so, we could causally interfere with this belief by convincing Gus that he's been getting the names wrong (say, he's somehow had 'gold' and 'silver' backward all this time), without this interfering with the former belief.

A certain kind of recalcitrance, best exemplified by the persistence of known illusion, is usually taken as evidence of modularity: the lines of the Müller–Lyer illusion still look to be of different lengths, because the processes involved are modular (e.g., Fodor 1983, 1984). But this kind of recalcitrance is highly typical of the perceptual seemings that result from perceptual learning. Once I've learned to recognize my wife's voice, or to spot chanterelles in the forest, things can spontaneously and insistently sound like her voice or look like chanterelles even if I'm carefully avoiding drawing the obvious conclusions. This is at the same time (some) evidence for the claim that the responsible processes are modular in the relevant sense and also evidence that the sensation-seeming transitions are not causally dependent on person-level knowledge or evidence about looks. These sorts of examples suggest that perceptual processing remains modular a good distance into the processing streams, that relatively late stages are still modular and thus that relatively high level properties will be perceptible independently of the agent's declarative knowledge.

When we're new to some perceptual domain, identifying properties or categorizing objects is sometimes a matter of running through a mental checklist, of agent-level inference. But that's not perception; it's inference. Markie seems to think that what perceptual learning does is to automatize this inferential part, in such a way that it's still the agent doing it, but now the agent can do it without thinking about it and without even tokening the operative beliefs. But that still wouldn't be perception. I agree that automatization of this general sort sometimes happens (Lyons 2014: 545–6), but I don't believe that this is typically, or maybe even ever, what happens in perceptual learning. If this learning were person-level/extramodular, it should be susceptible to disruption by testimony, in the way described two paragraphs back. But it generally isn't.[33]

The crucial disagreement here has to do with modularity and the subsequent autonomy and independence from person-level knowledge about how things look. Simple introspection won't tell us whether a given belief is the result of a modular process, because as Markie rightly insists, certain kinds of automatization result in a similar spontaneity. Because the point of contention is a causal claim, the only way to resolve it is to look to the counterfactuals, in this case about whether the seemings would still be produced in the absence of the person-level knowledge.

33 McGrath (Forthcoming) holds a view very similar to Markie's in this respect. Although I'm not treating it in detail here, since it's no longer a version of inferentialism, the problems that I'm raising here for Markie's view apply to McGrath's as well.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

4. CONCLUSION: INFERENCE AND THE DEEP FACTS ABOUT PERCEPTION

More generally, the modularity of perception renders any kind of inferentialism a highly dubious approach. I claimed above that the modularity of perception leads to two deep and pervasive facts about perception: first, that the intramodular computations are rationally underdetermined in an important sense, and second, that the outputs of perceptual processing are given, or receptive, in the sense that the perceiver is not agentially involved in their construction. The second fact makes it possible to argue that, even though the psychological processes leading up to the formation of a percept are inferential in some Helmholzian sense, this does not make them inferential in any epistemic sense. The first fact makes this move a highly desirable one, since the alternative would result in skepticism. Together, however, these facts suggest that the epistemology of perception has very little to do with the epistemology of inference.

The problem of cognitive penetration is an instance of a more general problem, which is that of finding a way to attribute different epistemic statuses to different perceptual beliefs. We have independent reasons for thinking that various kinds of inference yield conclusions with different epistemic statuses, and it would be nice if we could offer a unified account of what makes good inferences good and what makes good perceptual beliefs good. Reliabilism does this by finding some more general feature – reliability – that good inferences and well-formed perceptual beliefs have in common. Inferentialism tries to do it by subsuming the standards of perceptual belief formation under already existing canons of rational inference. Even setting aside the particular counterexamples I've offered here, I think that the general idea of inferentialism is ill-conceived. The first deep fact about perception implies that if perceptual beliefs are justified, then the epistemic rules that govern perception must be quite different from the rules that govern inference. If this is right, then the generic inferentialist project is a singularly unpromising one.

This doesn't mean that it is impossible to find general epistemic principles that encompass the rules of inference and the rules of perception; I think reliabilism offers such principles. It doesn't even mean that it's impossible for an internalist to formulate any such epistemic principles, even if plausible candidates don't leap to mind. But I think the internalist needs to look somewhere else than inferentialism; the rules for good perception aren't likely to fall directly out from, or even have very much to do with, the rules for good inference.[34]

REFERENCES

Bengson, J., Grube, E., and Korman, D. Z. 2011. 'A New Framework for Conceptualism.' *Noûs*, 45: 167–89.

BonJour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

34 Earlier versions of this paper were presented at Katholieke Universiteit Leuven and the University of Geneva. Some of the criticisms of Siegel's view were presented at the SOFIA conference in Huatulco, Mexico. Thanks to these audiences, especially Ariel Cecchi, Pascal Engel, Harmen Ghijsen, Peter Graham, Christoph Kelp, Lisa Mirrachi, Ram Neta, Bram Vaassen, and Wayne Wu. Thanks also to Matt McGrath, Susanna Siegel, and an anonymous referee for comments on an earlier draft of this paper.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

Brogaard, B. 2013. 'Phenomenal Seemings and Sensible Dogmatism.' In C. Tucker (ed.), *Seemings and Justification*. New York, NY: Oxford University Press.
Burge, T. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.
Clark, A. 2013. 'Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science.' *Behavioral and Brain Sciences*, 36: 181–204.
Conee, E. 1988. 'The Basic Nature of Epistemic Justification.' *The Monist*, 71: 389–404.
—— **and Feldman, R.** 2004. *Evidentialism*. Oxford: Oxford University Press.
—— 2008. 'Evidence.' In Q. Smith (ed.), *Epistemology: New Essays*, pp. 83–104. Oxford: Oxford University Press.
Crane, T. 1992. 'The Nonconceptual Content of Experience.' In T. Crane (ed.), *The Contents of Experience*, pp. 136–57. Cambridge: Cambridge University Press.
Currie, G. 1995. 'Visual Imagery as the Simulation of Vision.' *Mind and Language*, 10: 25–44.
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. 307–19. New York, NY: Blackwell.
Fodor, J. A. 1983. *The Modularity of Mind*. Cambridge, MA: MIT Press.
—— 1984. 'Observation Reconsidered.' *Philosophy of Science*, 51: 22–43.
—— 1985. 'Precis of *The Modularity of Mind*.' *Behavioral and Brain Sciences*, 8: 1–42.
Fumerton, R. 2013. 'Siegel on the Epistemic Impact of "Checkered" Experience.' *Philosophical Studies*, 162: 733–9.
Goldman, A. I. 2011. 'Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Problems, Reliabilism's Rescue Package.' In T. Dougherty (ed.), *Evidentialism and its Discontents*, pp. 254–80. Oxford: Oxford University Press.
Gupta, A. 2006. *Empiricism and Experience*. Oxford: Oxford University Press.
Hansen, T., Olkkonen, M., Walter, S., and Gegenfurtner, K. R. 2006. 'Memory Modulates Color Appearance.' *Nature Neuroscience*, 9: 1367–8.
Huemer, M. 2007. 'Compassionate Phenomenal Conservatism.' *Philosophy and Phenomenological Research*, 74: 30–55.
—— 2013. 'Epistemological Asymmetries between Belief and Experience.' *Philosophical Studies*, 162: 741–8.
Kvanvig, J. L. and Riggs, W. D. 1992. 'Can a Coherence Theory Appeal to Appearance States?' *Philosophical Studies*, 67: 197–217.
Lamme, V. A. and Roelfsema, P. R. 2000. 'The Distinct Modes of Vision Offered by Feedforward and Recurrent Processing.' *Trends in Neurosciences*, 23: 571–9.
Levin, D. T. and Banaji, M. R. 2006. 'Distortions in the Perceived Lightness of Faces: The Role of Race Categories.' *Journal of Experimental Psychology: General*, 135: 501–12.
Lyons, J. C. 2001. 'Carving the Mind at its (Not Necessarily Modular) Joints.' *British Journal for the Philosophy of Science*, 52: 277–302.
—— 2005. 'Perceptual Belief and Nonexperiential Looks.' *Philosophical Perspectives*, 19: 237–56.
—— 2009. *Perception and Basic Beliefs: Zombies, Modules, and the Problem of the External World*. New York, NY: Oxford University Press.
—— 2011. 'Circularity, Reliability, and the Cognitive Penetrability of Perception.' *Philosophical Issues*, 21: 289–311.
—— 2014. 'The Epistemological Import of Morphological Content.' *Philosophical Studies*, 169: 537–47.
—— 2015a. 'Critical Notice of Chris Tucker, ed., *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*.' *Analysis*, 75: 153–64.
—— 2015b. 'Unencapsulated Modules and Perceptual Judgment.' In J. Zeimbekis and A. Raftopoulos (eds), *The Cognitive Penetrability of Perception*, pp. 103–22. Oxford: Oxford University Press.
—— Forthcoming. 'Perception and Intuition of Evaluative Properties.' In A. Bergquist and R. Cowan (eds), *Evaluative Perception*. Oxford: Oxford University Press.
Macpherson, F. 2012. 'Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism.' *Philosophy and Phenomenological Research*, 84: 24–62.
Markie, P. J. 2005. 'The Mystery of Direct Perceptual Justification.' *Philosophical Studies*, 126: 347–73.
—— 2006. 'Epistemically Appropriate Perceptual Belief.' *Noûs*, 40: 118–42.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press

—— 2013. 'Searching for True Dogmatism.' In C. Tucker (ed.), *Seemings and Justification*, pp. 248–69. Oxford: Oxford University Press.
Marr, D. 1982. *Vision*. London: Allen and Unwin.
McGrath, M. 2013a. 'Siegel and the Impact for Epistemological Internalism.' *Philosophical Studies*, 162: 723–32.
—— 2013b. 'Phenomenal Conservatism and Cognitive Penetration: The Bad Basis Counterexamples.' In C. Tucker (ed.), *Seemings and Justification*, pp. 225–47. Oxford: Oxford University Press.
—— Forthcoming. 'Looks and Perceptual Justification.' *Philosophy and Phenomenological Research*.
Peacocke, C. 2001. 'Does Perception Have a Nonconceptual Content?' *Journal of Philosophy*, 98: 239–64.
—— 2004. *The Realm of Reason*. Oxford: Oxford University Press.
Pryor, J. 2000. 'The Skeptic and the Dogmatist.' *Noûs*, 34: 517–49.
Pylyshyn, Z. W. 2003. *Seeing and Visualizing: It's Not What you Think*. Cambridge, MA: MIT Press.
Raftopoulos, A. 2009. *Cognition and Perception: How do Psychology and Neural Science Inform Philosophy?* Cambridge, MA: MIT Press.
Reid, T. [1764] 1997. *An Inquiry into the Human Mind on the Principles of Common Sense*. D. Brooks (ed.). Edinburgh: Edinburgh University Press.
Riesenhuber, M. and Poggio, T. 1999. 'Hierarchical Models of Object Recognition in Cortex.' *Nature Neuroscience*, 2: 1019–25.
Sellars, W. 1956. 'Empiricism and the Philosophy of Mind.' In H. Feigl and M. Scriven (eds), *Minnesota Studies in the Philosophy of Science, Volume 1: The Foundations of Science and the Concepts of Psychology and Psychoanalysis*, pp. 253–329. Minneapolis, MN: University of Minnesota Press.
Shepard, R. N. 1966. 'Learning and Recall as Organization and Search.' *Journal of Verbal Learning and Verbal Behavior*, 5: 201–4.
Siegel, S. 2012. 'Cognitive Penetrability and Perceptual Justification.' *Noûs*, 46: 201–22.
—— 2013a. 'The Epistemic Impact of the Etiology of Experience.' *Philosophical Studies*, 162: 697–722.
—— 2013b. 'Reply to Fumerton, Huemer, and McGrath.' *Philosophical Studies*, 162: 749–57.
—— 2013c. 'Can Selection Effects On Experience Influence its Rational Role?' In T. Gendler (ed.), *Oxford Studies in Epistemology, Volume 4*. Oxford: Oxford University Press.
—— Forthcoming. 'How is Wishful Seeing Like Wishful Thinking?' *Philosophy and Phenomenological Research*.
—— in preparation. *The Rationality of Perception*.
—— **and Silins, N.** 2015. 'The Epistemology of Perception.' In M. Matthen (ed.), *Oxford Handbook of Philosophy of Perception*. Oxford: Oxford University Press.
Slotnick, S. D., Thompson, W. L., and Kosslyn, S. M. 2012. 'Visual Memory and Visual Mental Imagery Recruit Common Control and Sensory Regions of the Brain.' *Cognitive Neuroscience*, 3: 14–20.
Tolhurst, W. E. 1998. 'Seemings.' *American Philosophical Quarterly*, 35: 293–302.
Tye, M. 1995. *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
Tucker, C. 2010. 'Why Open-minded People Should Endorse Dogmatism.' *Philosophical Perspectives*, 24: 529–45.
—— 2014. 'If Dogmatists Have a Problem with Cognitive Penetration, You do Too.' *Dialectica*, 68: 35–62.
Wheeler, M. E., Petersen, S. E. and Buckner, R. L. 2000. 'Memory's Echo: Vivid Remembering Reactivates Sensory-specific Cortex.' *Proceedings of the National Academy of Sciences, USA*, 97: 11125–9.
Zeimbekis, J. and Raftopoulos, A. 2015. *The Cognitive Penetrability of Perception*. Oxford: Oxford University Press.

JACK LYONS is Professor of Philosophy at the University of Arkansas. He works mainly in epistemology, philosophy of mind, and cognitive science. He is editor of *Philosophical Topics*.

https://doi.org/10.1017/epi.2015.60 Published online by Cambridge University Press