



# Cognitive diversity and the contingency of evidence

Jack C. Lyons<sup>1</sup>

Received: 3 March 2020 / Accepted: 16 March 2022  
© The Author(s) 2022

## Abstract

Many epistemologists endorse a view I call “evidence essentialism:” if  $e$  is evidence of  $h$ , for some agent at some time, then necessarily,  $e$  is evidence of  $h$ , for any agent at any time. I argue that such a view is only plausible if we ignore cognitive diversity among epistemic agents, i.e., the fact that different agents have different—sometimes radically different—cognitive skills, abilities, and proclivities. Instead, cognitive diversity shows that evidential relations are contingent and relative to cognizers. This is especially obvious in extreme cases (from pathological to gifted agents) and in connection with epistemic defeat, but it is also very plausibly true of ordinary agents, and regarding prima facie justification.

**Keywords** Evidence · Reliabilism · Evidentialism · Cognitive processes

Externalist epistemologies often claim that justification and knowledge depend on factors that hold only contingently. A simple form of reliabilism, for example, holds that the justification of a belief depends on the reliability of the process that produced it, in the world where that process was used. A process that is reliable, and hence justification-conferring, in this world might be unreliable in another. One famous objection to this view holds that victims of a Cartesian demon might be fully justified in their beliefs, even though those beliefs result from processes that are unreliable in that world (e.g., Lehrer & Cohen, 1983). The idea behind this New Evil Demon Problem is at least in part that so long as my demon world counterpart is forming the same beliefs as I am, on the basis of the same evidence as I am, then her beliefs are equally justified as mine.

Evidence essentialism, as I’ll call it, is the view that evidence relations hold necessarily, in the sense that if  $e$  is evidence of  $h$ , for some agent at some time, then necessarily,  $e$  is evidence of  $h$ , for any agent at any time. Evidence contingency is the

---

T.C.: Neurodivergence (ii) Guest Editor Name: Lead Guest Editor: Dr. Alejandro Vázquez del Mercado

---

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

<sup>1</sup> University of Glasgow, Glasgow, UK

denial of this. The reliabilist view sketched above implies evidence contingency. But even if the New Evil Demon Problem shows that this kind of reliabilism is false,<sup>1</sup> there are other ways in which justification and knowledge might depend on contingent factors, beside external world factors that affect reliability. In particular, internal, psychological factors might help determine what counts as evidence for what, in a way that is incompatible with evidence essentialism. Different cognizers have different skills, capacities, abilities and proclivities, and this cognitive diversity has the effect of, among other things, rendering evidential relations contingent.<sup>2</sup>

The goal of this paper is to argue for evidence contingency. I begin by laying some groundwork in Sect. 1, clarifying what's meant by evidence essentialism and evidence contingency, in particular, getting clear on the operative understanding of evidence. In Sect. 2, I distinguish three general conceptions of evidence: an informational view, a phenomenal view, and an etiological view. In the next section, I discuss some problems for the informational view, one having to do with beliefs about necessary truths and another deriving from considerations of cognitive diversity. In the following two sections, 4 and 5, I consider phenomenal and etiological solutions to these problems, concluding that the etiological view does much better. The etiological view commits us to evidence contingency. In Sect. 6 I argue that the phenomenon of epistemic defeat provides a further argument for the etiological view, for very similar reasons. In Sect. 7, I briefly discuss the sense in which the evidence contingency I've been arguing for is compatible with internalism, and in Sect. 8 I offer some concluding comments.

## 1 Evidence essentialism and contingency

Before getting underway, we should start by asking what we mean by 'evidence'. There has been plenty of recent debate about whether evidence consists of facts, propositions, or mental states; whether only what one knows can count as evidence, etc. I'm not concerned here with these debates, and I'll stay neutral regarding them. If I sometimes proceed as if evidence is limited to (nonfactive) mental states [I happen to think this is the best way to understand evidence (Lyons, 2016)], this is only for expository convenience, and none of what I'm trying to argue here will depend on it.

We can't be completely neutral, however, on what kinds of things can stand in the evidence relation to beliefs. I'm going to be arguing that evidential relations are contingent, in the sense that whether  $e$  is good evidence (for  $S$ ) of  $h$ , depends on factors extrinsic to  $e$  and  $h$ . There's no way to get such an argument off the ground if we're working with an entirely trivial conception of evidence.

Suppose  $A$  and  $B$  each see that the litmus paper turned blue and each conclude that the solution was alkaline. However,  $A$  knew that blue litmus indicates alkalinity, while  $B$  was just guessing. It seems that  $A$  is more justified in believing  $h$  (the solution is alkaline) than  $B$ , even though they were relying on the very same  $e$ : that the litmus turned blue. Is this a counterexample to evidence essentialism? No, to make it seem so

<sup>1</sup> For the record, I think it doesn't (see Lyons, 2013), although I won't address this problem here.

<sup>2</sup> To borrow some potentially helpful terminology from Sosa (2015, forthcoming): even if evidential relations are not contingent on situational factors, they might yet be contingent on shape and especially the agent's innermost seat/skill.

would require framing the case in a way that misrepresents the actual evidence. What's really happening is that *B* is basing their belief entirely on *e*, while *A* is basing their belief on *e*, *plus* the (justified) belief that blue litmus indicates alkalinity. Justification is supposed to be determined only by the agent's *total* evidence, and *A* and *B* have different total evidence relevant to *e*, even though there's partial overlap. So *A* and *B* don't have the same evidence vis-a-vis *h* after all, and it's not a counterexample to evidence essentialism. Let's call the move just made the "expansion defense": the friend of evidence essentialism can often argue that a putative counterexample fails, by *expanding* the evidence base presumed by the objection in question.

This defense is, as the current example illustrates, at least sometimes quite reasonable and appropriate. But we can't let it go without limits, on pain of trivializing the concept of evidence. Imagine the following, very bad, philosophical exchange between a reliabilist and an evidentialist:

*R*: Seeing a blue litmus paper is sufficient all by itself for being (prima facie) justified in believing the solution was basic, just in case the agent is in a world where turning blue is in fact a reliable indicator of alkalinity. Since there are worlds where it's not a reliable indicator, evidence essentialism is false.

*E*: That doesn't show essentialism is false; you've underrepresented the evidence. If the laws of nature affect justification, then they factor into the agent's evidence, so that agents in worlds where blue doesn't indicate alkalinity ipso facto have different evidence than we do.

I wouldn't want to defend anything like what *R* is saying, but *E*'s response obviously relies on a conception of evidence that's both alien and useless. One frequently encounters the claim that evidence is whatever it is that justifies belief (Kelly, 2014; Kim, 1988), but no one should (and I suspect no one does) seriously believe this. That incredibly liberal conception of evidence would make *E*'s response here completely appropriate. But *E*'s response clearly isn't appropriate. We don't and shouldn't really think that evidence is whatever justifies belief; we think that an agent's evidence is restricted to what the agent does or could *base* a belief on. An agent's evidence consists of her grounds, or reasons, or bases for belief. (Or maybe a subset of those bases.) The exact nature of the basing relation is a difficult and controversial issue,<sup>3</sup> but all participants in this debate will agree that ordinary unreflective agents are not basing their beliefs on laws of nature of which they're completely unaware.

The point here isn't just about our ordinary concept of evidence; it's that if there's any work for a concept of evidence to play in epistemology, then it needs to be considerably narrower than "whatever it is that justifies beliefs." Furthermore, there *is* an important concept of *things that justify beliefs by serving as bases for those beliefs*. It is this concept that I (and I think, epistemologists more generally) intend to capture by talking about evidence. In what follows, therefore, I'll assume that one's evidence is limited to that on which one might base a belief. This leaves room for the aforementioned neutrality regarding facts, propositions, knowledge, etc., while avoiding the trivialization that we just saw in the overextension of the expansion defense.

<sup>3</sup> For the latest, see the recent volume by Carter and Bondy (2019).

Understanding evidence in this way keeps evidence contingency from being trivially false, but it's still not obviously true.

Evidence essentialism is very similar to a view that John Pollock called "cognitive essentialism" (my term is an adaptation of his, with hopes of making the meaning a little more transparent).

For example, it might be claimed that reasoning in accordance with *modus ponens* is always correct, whereas arriving at beliefs through wishful thinking is always incorrect. This is implied by the claim that the justification of a belief is a function of one's internal states, because what that means is that we can vary everything about the situation other than the internal states without affecting which beliefs are justifiable. In particular, varying contingent properties of the cognitive processes themselves will not affect whether a belief is justified. This is called *cognitive essentialism* (Pollock & Cruz, 1999; Cf Pollock, 1986).

Pollock and Cruz are talking here about cognitive *processes*, rather than about evidence, but this is probably less of a difference than it looks like. They think that justification is a function of what reasons (i.e., what evidence) a belief is based on, so no difference in cognitive processes could make an epistemic difference, unless that difference amounts to a difference in which beliefs are based on what evidence. In essence, "processes" for Pollock and Cruz are individuated entirely by their input–output mappings, where inputs are reasons/evidence, and outputs are beliefs. Thus, their cognitive essentialism, as they apply it, is my evidence essentialism.

Several epistemologists endorse some version of the "uniqueness thesis": "a body of evidence justifies at most one proposition out of a competing set of propositions (e.g., one theory out of a bunch of exclusive alternatives) and it justifies at most one attitude toward any particular proposition" (Feldman, 2007, p. 205). This is a controversial view, but not for reasons of contingency or cognitive diversity. Most of the resistance to uniqueness is in defense of the idea that a body of evidence might justify a *range* of attitudes (credences) toward a proposition, rather than a precise value. Thus, much of the debate surrounding uniqueness (or this formulation, at least) assumes evidence essentialism.

If evidence essentialism is true, then evidential fit can be a two-place relation between a body of evidence and a hypothesis or set of hypotheses supported to some degree by that evidence. If evidence contingency is true, however, then fit must be a three-or-more-place relation among the evidence, the hypotheses, and those contingent factors that determine what is evidence for what. On this view, any claims about evidential fit must at least tacitly make reference to these contingent factors, perhaps by explicitly or tacitly indexing the fit to some class of agents, worlds, times, etc. Because the reference might be tacit, it's hard to know whether authors who write as if evidential relations are two-place really mean it or not.<sup>4</sup> Nevertheless, it's reasonable to think that evidence essentialism has been pretty widely endorsed. Classical confirmation theory in the philosophy of science, for example, has concerned itself with this two-place

<sup>4</sup> White (2005, p. 445), for example, articulates uniqueness as holding that "[g]iven one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition." He could mean what Feldman seems to mean, or he could mean that once we fix *your* total evidence, we fix what attitudes it's rational for *you* to take. The latter is compatible with evidence contingency.

relation between bodies of evidence and sets of hypotheses. Bayesian treatments of confirmation and of rationality obviously make confirmation relations out to depend on other beliefs of the agent, but once we expand the agent's total evidence to include those beliefs, the justification relations between the total evidence and the beliefs they justify are fully necessary. Any further contingent features of the agent or the world outside seem to drop out as irrelevant.

## 2 Three conceptions of evidence

There are three broad views about the nature of evidence in the literature, which I'll call the informational, the phenomenal, and the etiological views.

According to the informational view, whether  $e$  is evidence of  $h$  is simply a matter of the contents of  $e$  and  $h$  and entailment or (a priori) probabilistic relations among them—for example, if  $e$  entails  $h$  or raises the probability of  $h$ . What makes a modus ponens inference a good one, on this view, is simply the fact that modus ponens is valid.  $e$  is good evidence for  $h$ , on this view, just in case  $h$  fits  $e$ , and this holds whenever  $e$  entails or raises the probability of  $h$ ; if the latter, then  $h$  fits  $e$  to the degree that  $e$  raises the probability of  $h$ .<sup>5</sup> Whether and how much  $e$  raises the probability of  $h$ , on this view, is fixed across possible worlds and is not in any way dependent on any cognitive agent. The informationalist (as we might call them) is clearly committed to evidence essentialism.

The phenomenal view holds that  $e$  is evidence of  $h$  if believing  $e$  (or being in  $e$ , in the case of experience states) gives  $h$  a certain phenomenal character—something like the character of seeming true. On this view, the modus ponens inference is good because—and if—the inferring agent has an intellectual experience as of “seeing” the conclusion to follow from the premises. If we took the evidence in this case to be simply the beliefs that  $p$  and that  $p$  implies  $q$ , the phenomenal view would make evidential relations contingent. But my interest here is in a phenomenal view wielded *in defense* of evidence essentialism. For this, we need the claim that the beliefs of interests are at least partly *based on* seemings; hence the seeming-true is intended to be part (maybe all?) of the agent's evidence. This is how I'll understand the phenomenal view in what follows.

The etiological view includes standard externalist and virtue theoretic epistemologies. It holds that  $e$  is evidence of  $h$  provided that the causal chain from  $e$  to  $h$  is of the right sort—e.g., that it instantiates a reliable type of process, or that it exemplifies intellectual virtue, etc. Here the modus ponens inference is good just in those cases where the relevant etiological factors obtain, e.g., the inference is an expression of intellectual virtue rather than vice, or the process by which the agent moves from the premises to the conclusion is a reliable one, etc. Surely it's a contingent matter whether a given modus ponens inference is undertaken out of intellectual virtue or not, but is it contingent whether the process is reliable? This is more obvious for transitions that are non-deductive, like inductive inference, or memory or perceptual belief if we

<sup>5</sup> There are various ways of understanding degree of fit in terms of ratios of probabilities, differences among probabilities, and so forth. My target here is the generic framework, not the specific proposals.

think of those as involving a move from experiences to beliefs or the like. But it might even be true for *modus ponens*, depending on how we think of processes. Processes need not be individuated simply by the inputs and outputs involved on that particular occasion (as Comesaña, 2010 does); more likely *modus ponens* inferences may be the result of different processes on different occasions and some or all of those processes might be involved in other input/output transitions and thus may have varying degrees of reliability (see Lyons, 2019). Thus, etiological views are compatible with—I think probably committed to—evidence contingency.

The three views are very general; not all those who hold “the” phenomenal view, for instance, agree with each other about much else. This is especially true of the etiological view, which encompasses a number of very different specific proposals. The version of the etiological view that I’ll focus on is a generic process or virtue reliabilist view, as this is the kind of view where cognitive capacities matter most. But I’m not really trying to defend reliabilism here. I won’t be assuming or arguing that reliability is either necessary or sufficient for (even *prima facie*) justification; I’m using it as a convenient illustration of an etiological view, in order to argue for evidence contingency and the general epistemic significance of cognitive diversity.

### 3 Problems with the informational view

An initial problem for the informational view is the well-known issue of necessary truths. Everything entails every necessary truth, and entailment seems to be the best evidence one could have on a strictly informational understanding of evidence. So every agent is fully justified in believing every necessary truth, or at the very least, in believing every tautology.<sup>6</sup> This seems extremely problematic, given that some tautologies are exceedingly obscure.

It is standard to say here that logical omniscience is an idealization, or what perhaps amounts to the same thing, that it’s a property had by ideal agents (or ideally rational agents). If this is supposed to be a response to the problem, it’s not clear how. If you’re trying to understand the behavior of gasses, it’s helpful to imagine particles as perfectly elastic spheres, even though they’re not. But if you’re trying to understand milk production, you won’t get far by imagining cows as perfectly elastic spheres. Perfectly spherical and elastic molecules would act a lot like real ones do, although in a tidier and more predictable way. No one knows how perfectly spherical cows would act, but presumably nothing like real ones.

Part of the difference is that in standard scientific idealization, we are abstracting away from specific, identified factors, which make a known difference to the phenomenon under investigation. If we know how an object would behave in the absence of friction and air resistance, we can figure out how it would act in the presence of small amounts of friction and air resistance (normally, pretty similarly, though slower and hotter), and then in the presence of increasing amounts, etc. There’s nothing like this in the cow case. Similarly, it’s entirely unclear what’s being abstracted from in cases

<sup>6</sup> See Fitelson’s (2010) response to an objection by Pollock (2006) along these lines to Bayesian epistemology.

of “ideally rational” agents. One suspects that “ideal” here simply means “perfect”, which is not at all what it means in the scientific case. An “ideal”  $F$  in the scientific case is simply an otherwise normal  $F$  that is operating in the absence of specific, and often explicitly specified, influences.

And anyway, assuming that it did make sense to talk about logically omniscient agents as epistemic ideals, how is that supposed to inform our epistemological judgments about and advice for ordinary mortals? Are we *approximately* justified in believing in these very obscure logical truths? Is it better to be approximately justified with absolute certainty or actually justified to an only moderate degree? Are the very obscure logical truths epistemically worse off than the mildly obscure ones? What could possibly explain the difference, given that the difference between the more and less obscure logical truths doesn’t seem to be an informational difference?

It’s not just the armchair observation that the Banach-Tarski theorem isn’t anywhere near certain for most of us—though that is important. It’s also that 50 years of research in the heuristics and biases tradition shows us that the psychological processes by which we come to form judgments about a wide range of subjects, including tautologies, are categorically different from the way we’re taught to do it in logic class (e.g., Tversky & Kahneman, 1974). And if we’re concerned with whether an agent knows or is justified in believing some tautologies, it seems clear that which process or capacity is employed is highly relevant. You reason carefully and explicitly to the conclusion that  $p$ , while I employ a fast but unreliable heuristic to reach the same conclusion. Surely this difference between us is reflected in differences in justification.

One possible response to these worries is to invoke the distinction between propositional and doxastic justification, where having propositional justification means that a given proposition is proper for you to believe, whether you do believe it or not (and regardless of why you would believe it if you did); doxastic justification is a matter of a particular belief token being justified, including its having been arrived at properly, if that’s a factor that has any epistemic normative force. Declan Smithies (2015) for example, holds the informational view about propositional justification but denies it for doxastic justification. Thus, even though we’re propositionally justified in believing every necessary truth, given that we’re non-ideally-rational agents, many of our beliefs in necessary truths are less than fully (doxastically) justified. Some aren’t (doxastically) justified at all. When it comes to doxastic justification, Smithies abandons the informational view in favor of an etiological view—in fact, a reliabilist view, at least for necessary truths. I think this is the right thing to say about doxastic justification, but it leaves propositional justification quite obscure. Although Smithies retains the common terminology of ideality, there doesn’t seem to be any sense in which the “ideal” reasoner is just a cleaned up and better version of the “non-ideal” reasoner. Not coincidentally, Smithies’s theory of ideal rationality/propositional justification does nothing to illuminate non-ideal rationality/doxastic justification. The two theories (of ideal/propositional and of non-ideal/doxastic justification) have nothing to do with each other. With this gap between them, it’s hard to see how there’s even anything *normative* about propositional justification on such a view. The propositionally justified propositions in this domain (the domain of necessary truths) are simply the truths in this domain, and it’s hard to see what they have going for them epistemically, over and above their being true.



In addition to these worries about idealization, the disembodied conception of evidence mandated by the informational view doesn't allow cognitive diversity among agents to have any evidential significance. Different possible creatures—indeed different actual humans—have different cognitive capacities and this fact seems to have important epistemic implications, which an informational view of evidence can't account for.

To take some familiar examples, Srinivasa Ramanujan seems to have had the ability to intuit mathematical truths that other, ordinary, mathematicians could only arrive at either by guessing or by means of long and difficult proofs (Littlewood, 1929). Calendar savants can often report, sometimes in less than a second, what day of the week a given date falls on, several decades away from the present (Miller, 1999; Thioux et al., 2006). Ordinary people, even when fully informed about the rules and schedule governing leap years, would be hard pressed to answer such questions at all, certainly not without laborious, perhaps written, calculation. Here the informational view gets the exceptional cases right, since Ramanujan and the savants do have evidence that objectively entails the extraordinary belief. What it gets wrong is the verdict regarding the ordinary person's lucky guess to the same answer. The contrast between the ordinary agent and the exceptional agent (e.g., Ramanujan or the savant) highlights just how wrong it is to credit the ordinary agent with knowledge and justification. At the same time, it shows that there won't be some way of distinguishing certain propositions or certain premise-conclusion pairs as more easily knowable than others, without reference to specific cognizers. Propositions that are not at all obvious to us are intuitively clear to a math genius, and inferences that are difficult for us are easy for others. *Modus tollens* is famously more difficult for us than *modus ponens* (Braine et al., 1984; Evans et al., 1993), but it could have been the other way around. Ease and obviousness correspond to epistemic status as well: the exceptional agents know and are justified in believing things that we ordinary agents can only guess at.

All of this applies not only to doxastic justification, but to propositional justification as well. There is a natural and theoretically important conception of propositional justification—quite unlike the highly idealized one dismissed a few paragraphs back—according to which the exceptional agents have propositional justification for propositions that they haven't yet formed beliefs about, for which we ordinary agents lack propositional justification. This undermines both the idea that the informationalist can simply be interpreted as making plausible claims about propositional, rather than doxastic, justification; and also the radical disconnection between doxastic and propositional justification that resulted from this sort of move. We're not all propositionally justified in the same things on the same evidence, and the differences among us in this respect bear striking relations to the differences among us with regard to doxastic justification.

This doesn't—yet—imply evidence contingency. If an expansion response is plausible, evidence essentialism might still be able to say the right things about these kinds of cases.

But it doesn't seem that an expansion response within the framework of the informational view could work. For instance, maybe the exceptional agents have a justified metabelief to the effect that they're very good at this sort of thing, which would give them additional evidence in support of their original belief, evidence that isn't had by



the ordinary agent. But it is unclear how this inductive evidence could increase the justification already derived from what was a perfectly conclusive inference. That is, such a response already supposes that the first-order inference didn't provide complete justification (perhaps not any), but that supposition is inexplicable on an informational view of evidence. (Furthermore, this response would make the exceptional agent's justification mostly or entirely inductive, which seems wrong.)

An expansion response thereby solves the wrong problem. We didn't need to explain why the exceptional agents are justified; we needed to explain why *we're not*. The contrast with exceptional agents spotlights how badly we fall short of real justification in these cases. What is needed here, and what the informational view cannot provide, is an account of how it is that I'm just *guessing* that January 27, 1927 was a Thursday and the calendar savant is not just guessing in forming the same belief. I'm guessing despite the fact that my knowledge of today's date and of how the calendar works jointly entail that fact and thus constitute the best evidence—on an informational view of evidence—there could be.

#### 4 Phenomenal fix?

A natural reaction to these sorts of problems is to note that, unlike the ordinary agent, the exceptional agent “sees” the truth of the proposition, or “sees” that it follows from the premises. Relatedly, perhaps, some truths or connections are *obvious* for the savant and not obvious for the rest of us. According to the phenomenal view, this “seeing” and this obviousness can be understood as the agent having a certain kind of intellectual experience, or “seeming,” somewhat akin to a perceptual experience (Huemer, 2001, 2007). The most sophisticated version of this view is Chudnoff's (2012, 2013), who emphasizes the epistemic importance of “presentational phenomenology”: “What it is for an experience to have presentational phenomenology with respect to  $p$  is for it to both represent that  $p$  and make it seem as if you are aware of a truth-maker for  $p$ ” (2012, p. 55). Such presentational phenomenology is present when I form simple (for me) intuitions or make simple inferences, like *modus ponens*; it's at least normally absent when I'm simply guessing.<sup>7</sup>

Thus, rather than invoking metabeliefs to distinguish the exceptional agents from the ordinary agents, we could invoke presentational phenomenology instead. This has two advantages. First, it better captures the ordinary difference between guessing and seeing something to be true than the account in terms of metabeliefs: sure, we often have the metabeliefs in the “seeing” case and lack them in the guessing case, but the reverse is often true as well; and besides, it's the presence or absence of presentational phenomenology that *explains* the presence or absence of the metabelief. Second, it works better than the metabelief response because it's being used in the context of a very different overall view about the nature of evidence: a phenomenal rather than an informational view. On the phenomenal view, the seeming—the thing that has presentational phenomenology—doesn't *supplement* some shared body of evidence;

<sup>7</sup> I'm treating intuition and inference here as relevantly similar. Chudnoff does not to my knowledge apply presentational phenomenology to inference, but only to perception and intuition. Huemer thinks seemings underwrite the justification of all three.

the seeming *is* the evidence. So there's no worry as there was for the informational view that the epistemic force of the seeming is going to be swamped by the epistemic force of the evidence that the exceptional and ordinary agents share.

By making phenomenal states part or all of the agent's evidence, the phenomenal view maintains evidence essentialism, even while allowing that different agents can have different degrees of justification for a given tautology or a given premise-conclusion pair. They can have different degrees of justification because they have different seemings and thus different evidence.

To get this line to work, however, I think the proponent of the phenomenal view has to claim that all of the epistemic work is being done by the phenomenal state, with none done directly by the informational properties or relations.<sup>8</sup> Otherwise, my calendar guesses will come out as justified. Informational features can be relevant, of course, but only insofar as they bring about presentational phenomenology.

But then the phenomenal view faces different problems. If it claims merely that presentational phenomenology or seeming-true is *necessary* for justification, then it indeed avoids wrongly ascribing justification to ordinary agents concerning obscure truths. But it doesn't have an account of what's gone right in cases of obvious truths and/or exceptional agents. If it claims that presentational phenomenology or seeming-true is *sufficient* for prima facie justification, then it is far too lenient a view of evidence. I don't get justification from affirming the consequent or hasty generalization, simply because these inferences seem right to me at the time.

Maybe this is an uncharitable reading of the phenomenal view, especially of a version like Chudnoff's. The phenomenal states that Chudnoff is interested in are, after all, quite a bit less common than the ones the Huemer appeals to; not all seemings in Huemer's sense have presentational phenomenology in Chudnoff's sense (see Chudnoff, 2018). I'm not sure exactly what Huemer means his seemings to be, but one sometimes gets the impression that any time one believes anything nonvoluntarily, it's because it seems true. Surely this very liberal conception of seemings is prone to overgeneralize as mentioned in the previous paragraph. For Chudnoff, as mentioned, it's only a particular species of seeming (or a particular aspect of some seemings) that is held to provide immediate justification: those seemings or aspects of seemings that have presentational phenomenology. But now we're talking about something quite rare, rather than something too prevalent. I doubt that I experience presentational phenomenology every time I draw a simple logical inference, although I sometimes do if I'm paying close attention and being unusually reflective. Thus, it's unclear what my evidence would be in such cases: not (simply) the premises, if the view is going to be a phenomenal view; but not the seeming either, because the relevant kind of seeming—one with presentational phenomenology—is absent in this case.

As far as I'm aware, everyone who attributes any epistemic role to seemings (including Huemer and Chudnoff, but also Pollock, 1986; Pryor, 2000; Tucker, 2010, etc.)

<sup>8</sup> I suppose one could endorse a mixed view, where phenomenal factors *and* informational factors both serve as evidence, with the former simply being much stronger than the latter. This would give the right, or close enough, answers for the cases we're interested in; but I don't see that it has anything going for it over the straight phenomenal view. Since the proponents of the phenomenal view that I know of seem to hold a straight view, on which it's only the phenomenology that matters, and that informational relations by themselves do nothing, I'll stick to consideration of this kind of view and ignore the mixed view.

understands them to be nonfactive: it could seem to you that you're aware of a truth-maker for  $p$ , even though you're not. Suppose then, that  $e$  is not at all evidence of  $h$ , but that it seems in the relevant sense that  $e$  is good evidence of  $h$  (that your experience represents  $e$  as evidence of  $h$  and makes it seem to you that you're thereby aware of the fact that  $e$  is evidence of  $h$ ). Should you believe  $h$  on the basis of  $e$  or not? It's not obvious. One line of thought has it that you should: what else should you do in such a case?<sup>9</sup> Another line of thought has it that no, you shouldn't, given that  $e$  is by hypothesis bad evidence of  $h$ . Sometimes, surely, it's reasonable to base a belief on something that you incorrectly but justifiedly *think* is good evidence: if I justifiedly believe that blue litmus indicates acidity, then it's reasonable to infer from the blue litmus that the solution was acidic. Even though in fact, blue is not evidence of acidity. But this is different: here we're considering what's essentially a modus ponens inference with a false but justified major premise—if the litmus turns blue the solution is acidic. That's not the way phenomenal states were supposed to generate justification. Phenomenal states were supposed to justify *directly*, resulting in the immediate justification of epistemically basic beliefs; they weren't supposed to operate by way of doxastic intermediaries. If they did work in this more complicated way, they would generate a Lewis Carroll-type regress.

Finally, there's something a bit surprising about the phenomenal solution to the problem of cognitive diversity, given that the cases that motivate it never hint that the exceptional agents even have the relevant seemings. We do have them when we form intuitions, and maybe even when we perform simple logical deductions in a particular, very reflective way. But is there any reason to think that Ramanujan or the calendar savants have this kind of phenomenology? Again, a lot may depend on how robust we expect this seeming to be. If seemings are as "thin" as Huemer thinks, it's very likely that exceptional agents have them, but very unlikely that they explain anyone's epistemic status. If they're as "thick" as Chudnoff thinks, we have no reason to think that exceptional agents have them.

## 5 Etiological views

The phenomenal view thus inserts seemings into the picture, writing them into a story that neither mentioned them explicitly nor even hinted very strongly of their existence. At the same time, it leaves out something that *is* very explicitly part of the setup regarding cognitive diversity and exceptional agents: that they are *very good at getting the right answers*. This, after all, is part of what we must mean by calling them savants. Seen in this light, it seems very strange to explain our intuitions that these agents are justified by appealing to their seemings and not to their reliability, when seemings haven't been hinted at and reliability is one of the most salient features of the cases. It is worthwhile, therefore, to consider how an etiological view might handle these cases.

Intuitively, the important difference between the savants' justified intuitions and inferences and our unjustified counterparts is that the former are *obvious* while the

<sup>9</sup> Jeremy Fantl has used this general line of argumentation (in conversation) in other connections.

latter aren't. This clearly isn't a matter of informational relations and we've seen that it's not plausibly a matter of accompanying phenomenology. I suggest that an *obvious* answer is one that's *easy* to get *reliably* right.<sup>10</sup> This simple proposal has a good deal of intrinsic plausibility, and it seems to sort the relevant cases in the right way. Most people are reliable at simple deductive inference, but less reliable when it gets complicated and starts pushing the limits of working memory; we're less reliable yet when guessing. The correct answers are obvious to the exceptional agents; that is, they are extremely reliable in their favored domains and are not at all just guessing.

It's necessary that, say, modus ponens and modus tollens are valid, but it's contingent that we're good at them. In fact, as alluded to above, we're better at modus ponens than at modus tollens, perhaps because we have a hard-wired natural deduction-type rule corresponding to MP, while MT has to be accomplished in a less direct way, by piecing together primitive processes (Braine & O'Brien, 1998). This extra computational complexity introduces additional possibility of error. We could have been wired differently, so as to be better at MT than at MP.<sup>11</sup> There could have been an agent who was wired up with just the Sheffer stroke, making a small number of logical inferences very easy for it, while everything else that's easy for us relatively difficult for it. Exceptional agents are presumably wired differently from us, and thus, some of the inferences that they make are far more reliably made and therefore more justified than those very same inferences when we make them. The inferences are of course equally *valid* no matter who makes them, but they're reliably produced only by *some* agents.

As mentioned above, the New Evil Demon Problem argues that contingent factors, like what kind of world you're in, don't have any effect on justification. The kind of contingency that's at issue here is quite different. Certain deductive inferences count as reliably produced not because of contingent facts about the external world,<sup>12</sup> but because of contingent facts about what psychological capacities the agent has

<sup>10</sup> Correctness and reliability are not introspectible, but ease of answering is. This is why obviousness has an introspectible component, in the sense that if an answer is obvious, then it will generally *seem* obvious. Of course, an answer might seem obvious without *being* obvious, if the processes we use to produce the answer are effortless but unreliable, or if the answer is wrong. If a bat and ball together cost \$1.10, and the bat costs \$1 more than the ball, it seems obvious that the ball costs \$.10 (Kahneman, 2011). Of course, since that answer was wrong, it wasn't obvious after all; it only seemed so. (This very distinction precludes a fully phenomenal unpacking of obviousness.) Note that it is the reliability that is doing most of the evidential work, not the ease or the correctness, although the ease might allow some supplemental bolstering from metacognitive processes, in something like the way metabeliefs were suggested to bail out informationalism above. An important difference between the present suggestion and the metabelief proposal above is that the first order process is not, on an etiological view, assumed to be 100% reliable, which means that the bolstering process has space to indeed increase the reliability of the overall process and thus have positive epistemic effect.

<sup>11</sup> I am trying to stay very neutral on how to individuate cognitive processes (i.e., how to solve the generality problem). For the current purposes, however, I need to assume at least the following: that processes are not individuated solely by the actual input–output mappings. If that were the case, we couldn't have a process that performed MT *badly*; it wouldn't be performing MT at all. This assumption seems highly intuitive and highly unobjectionable. Readers who want a more detailed typing of cognitive processes are encouraged to consult Lyons (2019).

<sup>12</sup> This could also happen, although that's a different point from the one I am emphasizing here. How could it make a difference? Depending on how processes are individuated (again, see Lyons, 2019), it might turn out that some kinds of deductive inference are accomplished by way of some of the very same processes

and which of them are employed for a given task. This makes evidence contingency compatible with a certain kind of epistemic internalism, a point I'll return to below.

One of the guiding thoughts behind evidence essentialism in general and the informational view in particular is that certain inferences, like *modus ponens*, should be good inferences no matter what, and others, like affirming the consequent, should be bad no matter what. We saw that the phenomenal view retains the essentialism while rejecting this guiding thought. But its consequent embrace of affirming the consequent wasn't a mark in favor of the view, partly because the view didn't do much to make that embrace seem plausible. There are other considerations, however, that might do a better job of breaking loose the intuition that MP is always good and AC is always bad.

Suppose that, in a range of circumstances (when we're not taking the time and effort to explicitly represent and apply learned rules of logic), we solve simple logic problems in the following way.<sup>13</sup> We (quickly and unconsciously) imagine the premises true and then try to construct a mental model consistent with that, where the conclusion is false. If, after a sufficient length of time (measured in the tens of milliseconds), we're unable to do so, we give up and judge the argument to be valid. Obviously, the reliability of this process will be sensitive to factors like how good the agent is at thinking of counterexamples when there are any, how long they spend searching for counterexamples, etc., but it's *not* sensitive to whether the particular instance was valid. If this process is unreliable, then it shouldn't confer justification even in the cases where it's lucky enough to yield the correct answer. Thus, some *modus-ponens*-conforming inferences would be unjustified. If the process is reliable, then it should confer justification even in cases where it's unlucky enough to have gotten the rare wrong answer. Thus, some affirming-the-consequent-conforming inferences would be justified, so long as the process doesn't produce them too often. If it's genuinely the same process that's being used, it shouldn't matter whether the particular instance is one of MP or AC; both are equally justified. AC-conforming inferences—via a given process—are *just as justified* as MP-conforming inferences—via that same process. In this way, the justification conferring power of *modus ponens* (and affirming the consequent, and the rest) depends on contingent factors.

So says the etiological view, at least in process reliabilist form.<sup>14</sup> I think this line of thinking renders quite plausible the view that the justification conferring power of a *modus ponens* inference is contingent. Again, of course, the *validity* of *modus ponens* isn't contingent. The point is simply that justification isn't wedded in this way to validity.

Earlier, I criticized the phenomenal view for licensing bad reasoning, like affirming the consequent. Now I'm endorsing affirming the consequent under some conditions. Is this a problem? I think not. I complained about the claim (roughly) that AC is good whenever it feels good; that's quite different from the claim that AC is good whenever it's properly (reliably) arrived at. The idea that AC *could* be properly (reliably) arrived

Footnote 12 continued

that are also involved in empirical cognition. In that kind of case, changes in the world that made the whole process more or less reliable would affect the justification of the purely deductive uses of that process.

<sup>13</sup> The substance of this paragraph is treated at greater length and in a different context in Lyons (2019).

<sup>14</sup> I think very similar considerations could be brought to bear in favor of responsibilist or other virtue-theoretic epistemologies. I leave this as an exercise for the interested reader.

at may not initially seem coherent, but the discussion two paragraphs back shows that it is.

Finally, an etiological view can find a way for informational factors to play some role, even if not the role the informational view wants them to play. On the phenomenal view, the reason, say, *modus ponens* confers justification *has nothing to do with its validity*. It confers justification because (if) it produces a relevant seeming. The relationship is at best causal (although it's hard to see how validity per se could bear a causal relation to anything). On an etiological view, the validity of MP can be—partially and indirectly—constitutive of the goodness of the processes that perform MP (and which also performs the occasional AC, although these inferences count against the goodness of the process). Because MP is valid, a process that performs it is, to that extent anyway, reliable. That is, the fact that it performs MP will be part of what makes the process reliable. But only part, because whether the process is reliable overall depends on what else it does. The informational view wanted validity of a particular inference to be fully constitutive of epistemic goodness. The failure of that view need not lead us to deny that there's *any* kind of constitutive relation between validity and justification.

## 6 Defeat and cognitive capacities

I've been arguing that an etiological view handles certain cases better than the informational and the phenomenological views, in virtue of its embrace of evidence contingency and the role it reserves for cognitive capacities. The discussion so far has only concerned *prima facie* justification. Of course to be *ultima facie* justified a belief must not only have *prima facie* justification but must also be undefeated. Here I argue that its treatment of defeat gives the etiological view an even clearer advantage over of the other two views.

Defeat sometimes (perhaps always) results from an agent's having propositional justification to believe something that contradicts or undercuts her current belief.<sup>15</sup> Sometimes this may be because, simultaneously with believing *p*, the subject consciously and occurrently believes something (or undergoes an experience) that immediately and obviously entails that *p* is false or that the agent's reasons for believing *p* are inadequate. More often, however, defeat involves long term memory, inference, or some combination of the two. That is, the "direct" rebutter or undercutter for *p* (the proposition that immediately and obviously entails that *p* is false or that the agent's reasons for believing *p* are inadequate) normally isn't already right there in working memory, but is either stored in long term memory or can be justifiedly inferred from

<sup>15</sup> I'm somewhat inclined to think that *all* defeat derives from propositional justification, although this is controversial. Some epistemologists also allow defeat to come from unjustified beliefs (which therefore don't propositionally justify anything) (e.g., Lackey, 2005) and/or from evidence you don't possess (and thus which doesn't propositionally justify anything) but should have (e.g., Goldberg, 2018). In a very different vein, one might hold that my having been irresponsible in forming a belief might count as a defeater for that belief. For more on these issues, see Graham and Lyons 2021. Setting aside the controversial kinds, I think everyone allows that propositional justification can yield defeat. For the classic statement of rebutting and undercutting defeat, see just about any of Pollock's epistemological writings, e.g., Pollock and Cruz (1999).

the agent's occurrent beliefs. Call these cases "defeat by memory" and "defeat by inference", respectively. Obviously, many cases of defeat will be combinations of the two.

Defeat by memory and by inference are very naturally accommodated by an etiological view, but not by the others. Suppose I believe  $p$ , with prima facie justification, but  $p$  is incompatible with other things I believe with even greater justification. Is  $p$  defeated? It depends, in part, on how obvious or unobvious the incompatibility is. If the defeating belief is simply not- $p$ , then the incompatibility is obvious and the defeat is straightforward. If, however, the only way to render the conflict explicit is by a complex and difficult derivation, then the incompatibility is unobvious and the defeat is much less straightforward.<sup>16</sup> As we saw above, however, how complex and difficult a given derivation is depends on contingent facts about the psychological makeup of the cognizer.

I think the right thing to say about these sorts of cases is that everything else being equal, a belief is *more* defeated—in the sense that its ultima facie justification is reduced by a greater amount—the more obvious the incompatibility. Frege didn't have a defeater for Axiom V before he read Russell's letter, even though the falsehood of Axiom V was entailed by other things he believed. Certainly, continuing to believe Axiom V after reading Russell's letter would have been much worse than believing it before.

I suggested above that we understand obviousness in terms of the ease of producing a reliably true answer. Ease and reliability are quite distinct features, and they play importantly different epistemic roles. Doxastic prima facie justification is almost entirely a function of reliability (though see note 10 above, about metacognitive supplementation). For propositional justification—and thus defeat—ease plays a much more central role, as it is a greater epistemic sin to ignore an easily derived consequence than a difficult and subtle one (similarly for ease and difficulty of memory retrieval).

We saw above that the informational view doesn't have a clear means of making sense of obviousness; this is just what initially motivated the other views. The informational view can only understand the degree of defeat as the *strength* of the reason—the degree of evidential fit—not the *obviousness* of the reason. This view makes defeat too common. Also too strong, in the sense that subtle defeaters will reduce justification a lot, when they should (because they're subtle) only ever reduce it a little.

The phenomenal view, by contrast, will make defeat too rare. In cases of defeat by inference, the things doing the defeating are not occurrent states. But it's only occurrent states that have phenomenology. Even if presentational phenomenology or seeming true offered a decent account of obviousness for beliefs we've already formed (I argued above that it doesn't), what's at issue here are beliefs we haven't formed but could properly arrive at by reasoning. Since those beliefs don't currently seem true, they can't be things that the agent is propositionally justified in believing, according to the phenomenal view, so they can't serve as defeaters.

<sup>16</sup> Once the derivation is done, if it is, the obviousness or difficulty of the derivation makes little further difference. The concern here is with the defeating power of not-yet-performed derivations.



Maybe the phenomenalist could appeal to counterfactuals, for instance: if I were to consider  $p$ , it would seem true. This particular counterfactual only holds for some obvious propositions, and not for those propositions for which a bit of reasoning is required. Relatedly, a counterfactual proposal—I think *any* counterfactual proposal—along these lines will fail to make defeat graded in the desired way. The only source of gradation on offer here is the strength of the seeming, and even if that's epistemically important, it doesn't correlate with degree of obviousness. What might possibly correlate with obviousness, and which is somewhat in the spirit of counterfactuals, is the distance out into the space of possible worlds one must go before the belief in question seems true. Now we have the right kind of gradation, but only because we now have a view that takes contingent facts about the agent's capacities seriously. I still think it's a mistake to interpret obviousness as distance in logical space before the proposition seems true, rather than distance before the agent easily and reliably gets it right. But even if we cede this to the phenomenalist, the distance metric has to be in some sense a reflection of the cognitive capacities of the agent, of what the agent would have to do to arrive at the belief in question.

I've been arguing that the phenomenal and informational views are inferior to an etiological view in accounting for the defeat where the "direct" defeater gets its justification from inference. Many of the same considerations hold when defeat depends on memory. Not everything stored in long term memory is equally easy to access. Ease of access to, for instance, stored counterexamples might depend on surprisingly contingent factors like cueing conditions, typicality, and cognitive set. Again, it seems to me that the greater the ease of recall of a potential defeater, the stronger its defeating force is: a given counterexample to  $p$  reduces my justification for believing  $p$  more if that counterexample is very easily recalled (even though it hasn't yet been recalled) than if it's only very difficult to recall. The informational and phenomenological views will stumble here just as they did concerning defeat from inference, and they will do so for the same reasons.

There is one interesting and important difference, however. The evidence essentialist is committed to the claim that justification is determined by fit with the evidence *one possesses*, and it would be reasonable to hold that what evidence one possesses is in turn determined by such things as ease of recall. This makes what evidence one possesses depend on cognitive capacities, but it was bound to be a contingent matter anyway what evidence one possesses, and the essentialist gets to retain a necessary connection between the evidence possessed and the beliefs supported. This, I think, is the move that the evidence essentialist should make. It has the interesting result of making evidence *possession* something that can come in degrees. This seriously complicates the standard evidentialist picture (what's worse for my belief that  $p$ ? having easy access to a memory that weakly supports not- $p$ , or having difficult access to a reason that strongly supports not- $p$ ?), but these were necessary complications anyway, because the phenomenon in question is complex. I'm not sure this friendly suggestion would be very welcome. Some evidentialists have gone to rather heroic lengths to keep possession from coming in degrees (e.g., Feldman, 1988), although it's unclear to me why. Justification was already bound to come in degrees; who cares if it fails to constitute a simple ordering?

Even so, this move might work quite well for defeat by memory, but it's implausible concerning defeat by inference. It's arguably true that our long term memories are things that we in some sense believe (though non-occurently), but the same is not plausible concerning the inferential consequences of things we (occurently) believe. Inference—especially synthetic inference—is a means of *enlarging* one's belief set, not *revealing* it.

All told, the phenomenon of defeat provides an even stronger argument for an etiological view, and thus for evidence contingency, than the phenomenon of *prima facie* justification. Evidence essentialism is the claim that overall evidential status, not just *prima facie* evidential status, is necessitated by the evidence possessed. Yet the cases of defeat we've been considering are cases where two agents might share the same evidence and yet have different status vis-a-vis *ultima facie* justification, because they have different inferential (and perhaps memory) capacities.

## 7 Contingency, evidence, and internalism

It's worth reiterating that all of this requires that we understand evidence in the somewhat narrow, though standard sense, where evidence is limited to that on which a belief is or might be based. If evidence is simply whatever justifies a belief, then there's no conceptual room for evidence contingency, but evidence essentialism becomes a trivial and uninteresting thesis. Cognitive diversity—differences among agents with respect to their cognitive capacities, skills, processes, and the like—makes for differences in which beliefs are supported for which agents, and how strongly. Because these differences among agents are not (not typically, at least) considerations on which the agent is basing her beliefs, these factors lie outside the agents' bodies of evidence. And yet they determine how strongly that evidence supports various propositions. These factors obtain contingently. Hence evidence contingency.

The argument here has been limited to memory, *a priori* intuition, and certain types of inference, all of which are just as reliable in demon worlds as in the actual world. One could make a case for evidence contingency by appealing to methods of belief formation that do differ in their reliability in demon worlds (Lyons, 2013), but that's not what I've attempted here. Consequently, the kind of evidence contingency I've been arguing for is compatible with a certain kind of internalism.

It is standard these days to distinguish *access internalism*, which holds that all factors relevant to justification must be such that the agent could know by mere reflection whether or not they obtain, from *mentalism*, which holds that justification supervenes on the (non-factive) mental states and conditions of the agent. An etiological view of evidence is incompatible with accessibilism, on the very plausible assumption that one doesn't always know how one is cognitively constructed, but it is compatible with mentalism.<sup>17</sup> It doesn't entail mentalism, of course, because extramental factors, like reliability *in situ*, might contribute to evidential status as well. But I haven't tried

<sup>17</sup> Evidence contingency—and, I think, my argument for it here—is also compatible with a fairly common internalist requirement on inferential justification, that the agent have a justified metabelief that the inference is a good one (Tucker, 2012). It is *not* compatible with the idea that this metabelief suffices to *make* that inference a good one for that agent.

to argue this or assume it. Part of the point here is that for these kinds of beliefs (/processes/inferences) one can be a reliabilist and an internalist at the same time.

## 8 Conclusion

I have been arguing that cognitive diversity—interpersonal differences in cognitive skills, capacities, and proclivities—favors evidence contingency over evidence essentialism. Although it hasn't been my task here, a worthwhile followup to the current work would be a detailed investigation of the various sorts of cognitive diversity and the specific ways in which they affect the epistemic status of real agents, for better and for worse.

Epistemology is a discipline where theories are evaluated partly on the basis of how well they handle cases of brains in vats, mysterious clairvoyant powers, and the like. For my purposes, all I needed was a very few, metaphysically possible, cases of cognitive diversity. But such cases and many more are very plausibly actual as well. The case of Ramanujan, though familiar and illustrative, is largely or entirely anecdotal. But calendar savants have been extensively studied (see references above), and there are surely other sources of the relevant variation. Major research programs hold that there are fairly profound differences in cognitive style between Eastern and Western cognizers (e.g., Henrich et al., 2010; Nisbett et al., 2001), and between political liberals and conservatives (e.g., Hibbing et al., 2014; Mendez, 2017), and any cognitive psychology textbook will review a wide range of individual differences based on age, gender, general intelligence, learning style, etc., all without yet even addressing various types of neuroatypical populations, all of whom have their own specific deficits and/or strengths. (The literature here is so vast and wide-ranging that one doesn't know where to begin citing; the interested reader should consult a cognitive psychology or cognitive neuropsychology textbook.) The discussion of cognitive diversity in connection with defeat by memory and defeat by inference didn't rely on any such exotic forms of cognitive diversity. We don't really need to consult studies of individual differences in inferential or mnemonic power or general intelligence to know that some people have a greater facility for certain inferences than other people, or that you and I might share a memory although it is easier for one of us to recall when needed than for the other.

Evidence essentialism holds that if  $e$  is evidence of  $h$  for  $S$  at  $t$ , then necessarily and for any  $S$  and any  $t$ ,  $e$  is evidence of  $h$ . I have argued that this view underestimates the importance of cognitive diversity among agents: different agents have different psychological capacities, and it's implausible to think that justification is fixed independently of those capacities. I distinguished three broad theories about the nature of evidence—an informational view, a phenomenological view, and an etiological view—and I argued that only the etiological view offers a promising epistemology concerning necessary truths, exceptional agents, defeat by memory, or defeat by inference. Although I used a reliabilist form of the etiological view as a convenient representative, I don't intend for any of the argument to hinge on a controversial form

of reliabilism; the general conclusion should be one that a certain kind of internalist could accept as well.<sup>18</sup>

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Braine, M. D., Reiser, B. J., & Romain, B. (1984). Some empirical justification for a theory of natural propositional logic. In *Psychology of learning and motivation* (Vol. 18, pp. 313–371). Academic Press.
- Braine, M., & O'Brien, D. P. (Eds.). (1998). *Mental logic*. Psychology Press.
- Carter and Bondy. (2019). *Well-founded belief: New essays on the epistemic basing relation*. Routledge.
- Chudnoff, E. (2012). Presentational phenomenology. *Consciousness and subjectivity*, 51–72.
- Chudnoff, E. (2013). *Intuition*. Oxford University Press.
- Chudnoff, E. (2018). Epistemic elitism and other minds. *Philosophy and Phenomenological Research*, 96(2), 276–298.
- Comesaña, J. (2010). Evidentialist reliabilism. *Noûs*, 44(4), 571–600.
- Evans, J. S. B., Evans, J. S. B., Newstead, S. E., & Byrne, R. M. (1993). *Human reasoning: The psychology of deduction*. Psychology Press.
- Feldman, R. (1988). Having evidence. In *Philosophical analysis* (pp. 83–104). Springer.
- Feldman, R. (2007). Reasonable religious disagreements. In L. Antony (Ed.), *Philosophers without Gods* (pp. 194–214). Oxford University Press.
- Fitelson, B. (2010). Pollock on probability in epistemology. *Philosophical Studies*, 148(3), 455–465.
- Goldberg, S. (2018). *To the best of our knowledge*. Oxford University Press.
- Graham, P. J., & Lyons, J. C. (forthcoming). The structure of defeat: Pollock's evidentialism, Lackey's framework, and prospects for reliabilism. In: M. Simion & J. Brown (Eds.), *Reasons, justification, and defeat*. Oxford University Press.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3), 61–83.
- Hibbing, J. R., Smith, K. B., & Alford, J. R. (2014). Differences in negativity bias underlie variations in political ideology. *Behavioral and Brain Sciences*, 37, 297–307.
- Huemer, M. (2001). *Skepticism and the veil of perception*. Rowman & Littlefield.
- Huemer, M. (2007). Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research*, 74(1), 30–55.
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kelly, T. (2014). "Evidence", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (Ed.), <https://plato.stanford.edu/archives/win2016/entries/evidence/>.
- Kim, J. (1988). What is Naturalized Epistemology? In J. Tomberlin (Ed.), *Philosophical Perspectives* 2, *Epistemology* (pp. 381–405). Ridgeview Publishing Co.
- Lackey, J. (2005). Memory as a generative epistemic source. *Philosophy and Phenomenological Research*, 70, 636–658.
- Lehrer, K., & Cohen, S. (1983). Justification, truth, and coherence. *Synthese*, 55, 191–207.
- Littlewood, J. E. (1929). Review of collected papers of Srinivasa Ramanujan. *Math. Gazette*, 14, 425–428.

<sup>18</sup> An earlier version of this paper was presented at a COGITO epistemology conference at the University of Glasgow. Thanks to the audience there, especially to Jessica Brown, Adam Carter, Luca Moretti, Bryan Pickel, and Mona Simion. Thanks also to two anonymous reviewers for this journal.

- Lyons, J. C. (2013). Should reliabilists be worried about demon worlds? *Philosophy and Phenomenological Research*, 86(1), 1–40.
- Lyons, J. C. (2016). Experiential evidence? *Philosophical Studies*, 173(4), 1053–1079.
- Lyons, J. C. (2019). Algorithm and parameters: Solving the generality problem for reliabilism. *Philosophical Review*, 128(4), 463–509.
- Mendez, M. F. (2017). A neurology of the conservative-liberal dimension of political ideology. *The Journal of Neuropsychiatry and Clinical Neurosciences*, 29(2), 86–94.
- Miller, L. K. (1999). The savant syndrome: Intellectual impairment and exceptional skill. *Psychological Bulletin*, 125(1), 31.
- Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001). Culture and systems of thought: Holistic versus analytic cognition. *Psychological Review*, 108(2), 291.
- Pollock, J. L. (1986). *Contemporary theories of knowledge*. Rowman & Littlefield.
- Pollock, J. (2006). *Thinking about acting*. Oxford University Press.
- Pollock, J. L., & Cruz, J. (1999). *Contemporary theories of knowledge* (2nd ed.). Rowman & Littlefield.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4), 517–549.
- Smithies, D. (2015). Ideal rationality and logical omniscience. *Synthese*, 192(9), 2769–2793.
- Sosa, E. (forthcoming). *Epistemic explanations*.
- Sosa, E. (2015). *Judgment and agency*. OUP.
- Thioux, M., Stark, D. E., Klaiman, C., & Schultz, R. T. (2006). The day of the week when you were born in 700 ms: Calendar computation in an Autistic savant. *Journal of Experimental Psychology: Human Perception and Performance*, 32(5), 1155.
- Tucker, C. (2010). Why open-minded people should endorse dogmatism. *Philosophical Perspectives*, 24, 529–545.
- Tucker, C. (2012). Movin' on up: Higher-level requirements and inferential justification. *Philosophical Studies*, 157, 323–340.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19(1), 445–459.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.