Episteme, 16, 3 (2019) 340–356 © Cambridge University Press
doi:10.1017/epi.2018.14

INTRASPECIES IMPERMISSIVISM

SCOTT STAPLEFORD
sstapleford@stu.ca

ABSTRACT

The Uniqueness thesis says that any body of evidence E uniquely determines which doxastic attitude is rationally permissible regarding some proposition P. Permissivists deny Uniqueness. They are charged with arbitrarily favouring one doxastic attitude out of the set of attitudes they regard as rationally permissible. Simpson (*Episteme*, 2017) claims that an appeal to differences in cognitive abilities can remove the arbitrariness. I argue that it can't. Impermissivists face a challenge of their own: The problem of fine distinctions. I suggest that meeting this challenge requires impermissivists to loosen up at higher levels – when comparing belief-forming systems that differ in the fineness of their doxastic outputs. This more relaxed take on Uniqueness is a kind of 'intraspecies impermissivism'.

INTRODUCTION

Epistemic permissivists think that evidence is slack in a certain way. I favour the snug, impermissivist alternative. What I'm offering here is assistance: I'll try to heat things up for permissivists by showing that one of the central arguments against them is still in the running, despite a very elegant attempt to block it. The pressure is still on, in other words, if I can deflect this promising counter. An additional bit of help is a suggestion: Impermissivists can meet a challenge to their own view by allowing slack at a higher level. Labelling these moves will help: The big argument against permissivism is *arbitrariness*. The counter invokes *cognitive abilities*. My parry is called *no excuses*. The challenge is *fine distinctions*.[1]

In Section 1, I'll explain what all of this means. In Section 2, I'll sketch the *arbitrariness* objection from Roger White. In Section 3, I'll sketch the *cognitive abilities* response from Robert Mark Simpson. In Section 4, I'll present my *no excuses* argument for getting *arbitrariness* back on its feet. In Section 5, I'll suggest a possible response to the *fine distinctions* challenge posed by Thomas Kelly.[2] I conclude that permissivism is still at risk but, given the results of Section 5, it might be vindicated when we're comparing humans to other possible belief-forming systems. The slightly attenuated 'snug' view that I'm proposing, therefore, could be dubbed *intraspecies impermissivism*.

1 The name 'permissivism' is due to Roger White (2005: 445). I borrow 'impermissivism' from Dogramaci and Horowitz (2016: 130).

2 Kelly doesn't use the term 'fine distinctions', but it seems to fit.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

1. PERMISSIVISM AND UNIQUENESS

You're a *permissivist* if you *reject* the following thesis; you're an *impermissivist* if you *accept* it:

> **Uniqueness:** If an agent whose total evidence is E is fully rational in taking doxastic attitude D to P, then necessarily, any subject with total evidence E who takes a different attitude to P is less than fully rational. (White 2014: 312)[3]

Notice that the discussion is framed in terms of *ideally rational agents*.[4] The question is not about how people as a matter of fact respond to a body of evidence E. It's about how they *would* respond if they were *fully rational* in their evaluation of E. How much divergence of opinion does ideal rationality allow? How much *latitude* is there in the evidence? That's where permissivists and impermissivists differ.

Looser formulations are possible:

> **Uniqueness:** There is just one rationally permissible doxastic attitude one can take, given a particular body of evidence. (White 2014: 312)[5]

Or maybe:

> **Uniqueness:** The evidence uniquely determines which doxastic attitude is rationally permissible.[6]

Uniqueness can be usefully discussed on a *tripartite* division of doxastic attitudes – belief, disbelief and suspension of judgement – but also in connection with credences, where permissivists allow, and impermissivists disallow, varying *degrees of confidence* in a proposition P. Thus we get two versions of permissivism:

> **Strong Permissivism:** There are cases in which it is rationally permissible to believe P, but it is also rationally permissible to believe not-P instead, given the very same evidence.
> **Moderate Permissivism:** There are cases in which there is more than one rationally permissible *degree of confidence* one can have in P, given the same evidence. (White 2014: 312)

3 A referee for this journal suggests reducing the number of terms from three to two so that we have *uniqueness* on the one hand, and *permissivism* on the other. It's not a bad idea, but my inclination is to follow the crowd here (at least the newer crowd) and define the two positions with respect to the thesis – those who accept it (*impermissivists* or *uniquers*) and those who reject it (*permissivists*). You find this in Dogramaci and Horowitz (2016), Castro and Hart (2018), Schultheis (2018), and others. "*Impermissivists* accept *uniqueness* … *Permissivists* deny *uniqueness*" (Castro and Hart 2018: 1–2). To my mind, that's clearer. The reader who disagrees should replace every occurrence of *impermissivist* with *advocate of uniqueness*.

4 Ideal rationality implies having beliefs that are "rationally impeccable" (White 2014: 312).

5 Instead of saying that a body of evidence E uniquely determines *just one* – or worse, *exactly one* – doxastic attitude it is rational to take towards P, it may be preferable to say that it uniquely determines *at most* one. Where you don't understand, or are incapable of understanding, some proposition P, there is *no* doxastic attitude you can rationally take towards P. I'm getting this from Jonathan Matheson (2011: 360–61), a strong supporter of Uniqueness.

6 This is similar to Richard Feldman's original statement of the principle (Feldman 2007: 205).

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

The objection, the counter and the parry will go through on either version, so I'll stress one or the other as aids comprehension.

Two more preliminaries and we'll be set. First, there's a convenient term from White:

> **Permissive Case:** A case in which an agent's evidence E does not determine a uniquely rational attitude to take to a proposition P. (White 2014: 313; cf. White 2005: 450)

The extreme position, accordingly, is impermissivism: There are *no* permissive cases. Open the door to *one*, and you're *permissive* on matters of evidence.

Finally, there are two ways of specifying the *domain* in which permissiveness does or doesn't occur. You can hold that evidence is permissive or impermissive within the *person*, or you can hold that evidence is permissive or impermissive within the *community of knowers*. Thomas Kelly makes this sharp, distinguishing between permissivists who allow *intrapersonal slack* and permissivists who allow *interpersonal slack* (Kelly 2014: 304).[7] The former hold that an individual could find herself in a situation such that more than one doxastic attitude (or credence) towards some proposition P is rationally permissible given her evidence. The latter hold that "different individuals possessing the same evidence might believe differently, and each be reasonable in believing as they do" (Kelly 2014: 304). A taxonomy of permissiveness emerges, depending on which type of slack is allowed and disallowed: intrapersonal permissivism and impermissivism; interpersonal permissivism and impermissivism.[8] These distinctions will help us later.

Uniqueness seems very intuitive to me – almost *obviously* right. So what am I missing? Why would anyone deny Uniqueness? It loses some of its luster when you start thinking in terms of *fine distinctions*. Kelly's Alpha Centaurian thought experiment, which veers away from the "maximally coarse-grained" conception of belief – the tripartite conception – towards an increasingly fine-grained model of credences, is persuasive:

> Suppose that when we meet the Alpha Centaurians, they differ from us in only one important respect: they routinely take up doxastic attitudes towards propositions that are extremely fine-grained compared to our own. So, for example, the Alpha Centaurians really do have psychological states such as *believing to degree .5436497 that the Democrat will win*, or *believing to degree .5122894 that it will rain tomorrow* … The friend of Uniqueness might insist that, for any possible evidential situation, the evidence in that situation singles out some one, exact degree of belief that it is uniquely reasonable for the Alpha Centaurians to have, any slight deviation from which already counts as a deviation from perfect rationality … But as one cuts up the psychology more and more finely, Uniqueness looks increasingly counterintuitive. (Kelly 2014: 300)

That really does sound strange. So there's definitely something going for permissivism, especially the moderate form.[9] But the *arbitrariness* objection is indifferent to this coarse-grained–fine-grained distinction. That's the real headache for permissivists: Arbitrariness infects their view at any degree of resolution.

---

7 Cf. Douven (2009: 354).

8 I doubt if any actual epistemologists allow intrapersonal slack while disallowing interpersonal slack, but it's a coherent option. (I'm going a bit beyond what Kelly says with this taxonomy.)

9 Additional reasons for doubting Uniqueness can be found in Douven (2009), Conee (2010), Ballantyne and Coffman (2011), Brueckner and Bundy (2012), Rosa (2012), Schoenfield (2014), Castro and Hart (2018), amongst others.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

2. THE ARBITRARINESS OF BEING PERMISSIVE

The *arbitrariness* objection to permissivism starts from the assumption that responsible enquiry is guided by evidence. It then proceeds: *If* there are permissive cases, there is no reason to let evidence guide enquiry. So there are no permissive cases. The key, conditional claim is supported by a thought experiment. It can be framed in terms of doxastic attitudes or credences. To avoid distracting complications (White 2005: 457, note 3), White focuses on doxastic attitudes:

> As a jury member it is my responsibility to do my best to arrive at the truth about the defendant's guilt or innocence. The responsible way to go about this is to examine the available evidence and seek to proportion my belief to it. There are lazier ways of going about it. I have two magical belief-inducing pills, one is marked GUILTY, the other NOT GUILTY. Ingesting a pill will give me the labeled belief. Clearly it is irresponsible to take such a pill. Why? Arbitrarily taking a pill while I have no clue as to whether the defendant is guilty gives me only a 50 percent chance of arriving at the correct verdict … Now suppose that permissivism is true and that the present case is a permissive one and that I know this … What reason do I have to form my belief by an examination of the evidence rather than just popping a pill? (White 2014: 315)

If the case is *strongly* permissive, more than one attitude gets a pass: I can believe that the defendant is guilty *or* disbelieve it. I'm free to take my pick. But then I must view the means by which I arrived at my favoured attitude "to be no better than arbitrarily popping a belief-inducing pill" (White 2005: 451). So why bother with evidence? Seeing as there *is* reason to let enquiry be guided by evidence, permissivism must be rejected.

Here's another angle: An agent who takes herself to be in a permissive case with respect to some proposition P, and at the same time privileges her own attitude or credence towards P, adopts what Simpson helpfully calls a "quasi-akratic stance" (Simpson 2017: 523). A fully akratic stance is one in which an agent believes some proposition P but also holds that P is unsupported by her evidence (Horowitz 2014: 718). Simpson's meaning, I think, is this: Suppose an agent believes P on E with a credence of 0.8. In doing so, she is in effect denying that any other credence is rationally permissible. Suppose she's also permissive. In opening herself to slack, she betrays her preference for 0.8. It's sort of like believing something and admitting that the belief has no support – it's *sort of* akratic – since she's simultaneously saying that the evidence justifies a credence of 0.8 and that it also justifies a credence of 0.3. If 0.3 is what the evidence makes reasonable, then stepping it up to 0.8 is to go 0.5 beyond what's reasonable. It's believing *that much more* without support. And that's downright arbitrary.

You can't have it both ways, according to *arbitrariness*. You can't favour your own credence or attitude towards P *and* recognize the legitimacy of rival credences or attitudes. You might as well settle on a credence by spinning a wheel or picking values from a hat.

3. SLACK IN COGNITIVE DIVERSITY

Simpson sets two tasks for the permissivist in responding to *arbitrariness*:

> *Task 1*: Identify something epistemically non-arbitrary about one's own doxastic attitude DA_1 towards P (or one's own credence of 0.8, for instance) that makes preferring it rational.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

> *Task 2*: Accomplish *Task 1* in a way that doesn't undermine permissivism – a way that doesn't treat DA_1 (or credence 0.8) as in general rationally superior to rival doxastic attitude DA_2 (or to rival credence 0.3, say).[10]

Miriam Schoenfield makes a promising suggestion. She defines 'epistemic standards' like this:

> *Epistemic Standards*: Epistemic standards are methods for forming beliefs based on evidence – they are "functions from bodies of evidence to doxastic states [or credences]" (Schoenfield 2014: 199).

And proposes meeting *arbitrariness* like this:

> Defend DA_1 (or your 0.8 credence) on the grounds that it is enjoined by the epistemic standards – call them ES_1 – that you regard as most truth-conducive.
> Recognize that it may be rationally permissible for a person to have a different attitude DA_2 (or credence 0.3), if that attitude (or credence) is enjoined by a different set of epistemically acceptable standards – call them ES_2 – which the other person regards as most truth-conducive.[11]

This satisfies *Task 1*, because a non-arbitrary reason for adopting DA_1 (or credence 0.8) has been identified. It is imposed by ES_1, which you regard as most truth-conducive. And it satisfies *Task 2* because DA_2 (or credence 0.3) is enjoined by the (rationally acceptable) epistemic standards (ES_2) that the other person regards as most truth-conducive. You needn't regard DA_2 (or credence 0.3) as impermissible since "the rational status of any doxastic attitude is assessed relative to the agent's epistemic standards" (Simpson 2017: 528).

Simpson rejects this proposal: "Where before our problem was arbitrarily favoring one among two rationally permissible doxastic attitudes, now our problem is arbitrarily favoring one among two rationally permissible sets of epistemic standards. The problems … recur for Schoenfield at the level of epistemic standards" (Simpson 2017: 529). Unfortunately, the reference point for any epistemic evaluation is one's own set of epistemic standards. If your standards are ES_1, then you see ES_2 as less accurate, less truth-conducive than ES_1. You should give no weight to ES_2's recommendation of DA_2 (or credence 0.3). How to play favourites with your own epistemic standards without in the process discrediting your peer's?[12] It's not obvious that you can.

But Schoenfield is onto something with her appeals to epistemic standards, according to Simpson, and the core idea can be preserved without falling into arbitrariness. His ingenious move is to tie epistemic standards to cognitive ability:

> The answer lies in an appeal to the idea that there is some kind of interaction between the epistemic standards we adopt … and our cognitive abilities, i.e. the abilities that we put into practice in the application of our epistemic standards. Analogically, the idea is that the truth-conduciveness

10 This is a paraphrase of Simpson (2017: 527).

11 This is a summary of Schoenfield (2014: 199–202). Note that she defends *strongly* permissive cases only (214, note 3). Nothing I say depends on this restriction, so I'm including credences in the conversation.

12 Dogramaci and Horowitz are also critical: "endorsing some standards as rational involves encouraging others to take up those standards" (2016: 144).

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

> of an agent's epistemic software depends on what kind of cognitive hardware that agent is working with. (Simpson 2017: 531)

To put it very simply, my epistemic standards might be most truth-conducive for me given my cognitive abilities and yours might be most truth-conducive for you given your cognitive abilities. We should thus follow the recommendations of our own standards. That's where evidential slack creeps in.

Simpson brings this thought to life with an example he calls *detectives* (Simpson 2017: 531–33). Veronica Mars and Nancy Drew are investigating a crime. They share a body of evidence. Their beliefs differ as to who committed the crime. Mars believes 'X is the culprit' with a credence of 0.8; Drew gives it a 0.3. (I'm simplifying the example and making it about credences.) Can they be permissive? Can they regard the other's credence as rationally permissible while standing firm on their own? Apparently so, since they bring different epistemic standards to bear on the evidence and – here we move beyond Schoenfield – the standards are well-suited to their respective cognitive abilities:

> Mars applies her standards because she's cognitively well-equipped to apply them. She has the kind of perceptual and attentional abilities that enable her to accurately judge the sincerity of testimony. In a similar way, Drew is cognitively well-equipped to apply her standards. She has a particular imaginative facility that enables her to formulate credible explanatory narratives, based on gestalt interpretations of diverse bodies of evidence … Although Mars and Drew are working together for the first time, they have a long track-record of working independently on other cases. And in light of their track-records, both of them have good reason to believe that when they apply their own epistemic standards, in forming beliefs about cases relevantly similar to the case at hand, they are very likely to get the answer right. In short, each of them is in the position of (i) having a combination of cognitive abilities and epistemic standards which reliably generates accurate beliefs, and (ii) having good reason to believe that this is the case. (Simpson 2017: 532)

By relativizing standards to the agent's cognitive abilities and making each aware of the other's situation, Simpson has found a novel way of allowing them to non-arbitrarily prioritize their own preferred credences while acknowledging the rationality of their rival's. This solution is highly attractive. Let's put it through the wringer.[13]

13 I don't want to take a hard line on the nature of epistemic standards, partly because I'm not entirely sure what to make of them. Wedding myself to any particular view would also limit the usefulness of my arguments. Still, if pressed, I think I'd follow Schoenfield and treat them as "functions from bodies of evidence to doxastic states" (2014: 199). Another approach is to treat them as particular ways of assigning weights to various theoretical virtues like simplicity, unification, explanatoriness, fruitfulness, etc. I'm leaning on a referee here, who's just handed me an additional argument against Simpson. It goes something like this: You could make a case for the permissibility of different epistemic standards by pointing out that there is no privileged way of trading off simplicity against fruitfulness, fruitfulness against explanatoriness, and so on. So the assigning weights understanding of epistemic standards is friendly to permissivism. But it's not friendly to Simpson. Differences in cognitive makeup don't obviously imply that different weightings of theoretical virtues are more or less truth-conducive for the agent doing the weighting. This is not the argument I had in mind, but it's a good one. Many thanks to the referee for putting another arrow in my quiver.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

4. NO EXCUSE IN COGNITIVE DIVERSITY

I'm going to argue that the appeal to differing epistemic standards tied to differing cognitive abilities does not block the arbitrariness objection, since it does not accomplish *Task 2* – it does not provide a non-arbitrary reason for privileging one's own credences (or doxastic attitudes) in a way that still recognizes the legitimacy, or rational permissibility, of rival credences (or doxastic attitudes) on the same body of evidence. Rival credences get no validation, in other words, by filtering through epistemic standards that are customized to an agent's cognitive quirks. I'll try to make this clear with a few examples. They are intentionally farfetched – pushing Simpson's brand of permissivism to the breaking point so that its problems become apparent. We can ease our way into the examples with a simple analogy concerning vision.

Suppose that my mother and I are standing on a bridge looking down at a cruise ship going past. And suppose that we have the same relevant background beliefs. How many lounge chairs are there on the deck by the pool, we're wondering? Looks like three to me – plain as day. So I form the belief that there are three lounge chairs by the pool with a credence of 0.8. My mother should form the same belief with the same credence since we have the same body of evidence, right? Wrong. My mother has bad vision. All she sees are three blurry red patches against a bluish background. Maybe there are some blue chairs that she can't make out. Maybe someone redistributed or stacked several red cushions from any number of different chairs. A lower credence, therefore – say, 0.3 – might be more reasonable for my mother, but not on the *same* body of evidence. Her evidence is different than mine precisely because her perceptual abilities are different.

The lounge chair analogy homes in on the obvious: Differing *perceptual* abilities afford different bodies of evidence.[14] Is it the same with other mental abilities, the kind more relevant for Simpson? Let's weigh that in connection with his main example – *detectives*. Mars is good at interpreting signs of insincere testimony in isolation and Drew is not. Say they're hearing testimony regarding the murder of the College Dean. Is it any good? When Mars sees the witness scratching his face, covering his mouth, or smiling without squinting, she takes these as signs of deception. But it's all lost on Nancy Drew. Why is that? If she has difficulty paying attention and simply doesn't *notice* these things – if the actions don't even *register* – then her evidence is clearly not the same as the evidence available to Mars. If *that's* what's going on, then Drew is something like my mother on the bridge: Her perceptual input is comparatively impoverished. So her doxastic output will naturally differ. She'll assume that the witness is credible. Mars will be suspicious. But that won't implicate Uniqueness.

If she *does* notice them, on the other hand, her inability to grasp their meaning is more plausibly construed as a lack of knowledge. What *appears* to be a difference in cognitive competence is in fact a difference in background beliefs. Mars must have some additional information. Either she is expressly aware that these actions are characteristic signs of

14 If we class *perceptual* abilities under *cognitive* abilities – I have no view on whether we should – they're not going to be prime candidates for violating Uniqueness, since they very clearly lead to differences in quality of sensory evidence.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

deception, or else she's built up some implicit knowledge of their significance. So once again the evidence is different. Once again Uniqueness is safe.[15]

Simpson's detectives example, therefore – the backbone of his argument – provides ambiguous support for his thesis. The readings I'm suggesting may be more natural. It's certainly hard for me to see this as a permissive case, since I just don't find it credible that Mars and Drew are working with the same materials. Perhaps some additional details could be tacked on to give the permissive reading a boost. But there's nothing in the description given to make it necessary.[16]

I don't want to push too hard on this point, however, since Simpson might be able to establish evidential parity by expanding the story, and – here's the main thing – I'd still be sceptical that the case is permissive. I'd like to try a different strategy: going to extremes. Sometimes it's easier to sort out the subtler, less obvious cases, by looking at extreme examples and using them as a guide to the harder ones. This will work so long as the extreme and the central cases are relevantly similar. The examples I have in mind are designed to illustrate, and make plausible, the following idea: We shouldn't (and probably don't) regard differences in cognitive abilities as licensing rival doxastic attitudes towards the same proposition on the same body of evidence, since we shouldn't (and probably don't) regard a difference in cognitive abilities as an epistemically valid excuse for not believing what, by our lights, the evidence supports. Alternatively, if differences in cognitive ability can make a body of evidence permissive, then differences in cognitive ability can make (epistemic) irrationality epistemically acceptable. But that's a mistake, I hope you'll agree. Let the extremes pump your intuitions.[17]

My first case makes use of a much discussed mathematical problem devised by Shane Frederick: "A bat and ball cost $1.10. The bat costs one dollar more than the ball. How much does the ball cost?" (Kahneman 2011: 44). The quick and tempting answer – the *System 1* answer, in Kahneman's terms – is 10¢. The correct answer is 5¢. Most college students get this wrong (Kahneman 2011: 44–45). Now suppose that Victoria and Elizabeth are working with the same body of evidence E, and that E includes these two propositions (or the corresponding beliefs),[18] no defeaters, and nothing else that's relevant:

P1: A bat and ball cost $1.10.
P2: The bat costs one dollar more than the ball.

Victoria, we'll suppose, answered correctly: 5¢. Elizabeth answered incorrectly: 10¢. Let's suppose further that Elizabeth is cognitively incapable of focusing on math problems – cognitively incapable of engaging slow, *System 2* thinking – and so is unable to perform even simple calculations.

15 The same sort of analysis will show that Drew's "holistic, gestalt assessments" (Simpson 2017: 532) colour the evidence as well.

16 There's much else of interest in the paper, however, even if I'm right that the example fails.

17 I'm definitely not trying to saddle permissivists with the improbable view that *every* evidential situation permits more than one rational response. I am trying to make permissivists more sensitive to the things that make impermissivists uncomfortable. As I see it, the extreme cases are continuous with the 'tame' ones standardly cited by permissivists.

18 My sense is that psychologism about evidence is true and that propositionalism is false. Psychologism is the view that evidence consists of mental items (including beliefs); propositionalism is the view that evidence consists of propositions. See McCain (2014: Chapter 2) for arguments in favour of psychologism. Nothing rides on the distinction here.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

Since Elizabeth's cognitive abilities differ from Victoria's, she makes use of different epistemic standards – standards more suited to her distinctive capacities. We can specify the operative principle for each of the agents (the programme instruction for solving math problems) like this:

> **ES1-CURSORY MATH (Elizabeth)**: Believe in accordance with your first impression.
> **ES2-CAREFUL MATH (Victoria)**: Believe what appears right after careful calculation and double-checking.[19]

We'll concentrate on Elizabeth's answer (10¢). The output of **ES1-CURSORY MATH** on E is *belief* in **P3**; the output of **ES2-CAREFUL MATH** is *disbelief*:

> **P3:** The ball cost 10¢.

Given these assumptions, should we conclude that belief in **P3** was epistemically permissible for Elizabeth on E but not permissible for Victoria? Should we call the case *permissive*? No way: Elizabeth's inability to understand what her evidence implies doesn't alter its implications. The fact that she has different epistemic standards does not affect the evidential relations themselves, in other words, and does not affect what propositions a body of evidence supports, or makes reasonable.[20] **ES1-CURSORY MATH** might be a good *rule of thumb* for Elizabeth: A pre-reflective intuition could be better in a mathematical context than nothing. Who knows? But it doesn't help Elizabeth here: She doesn't end up believing the proposition that her evidence supports on the basis of the evidence that supports it. That's the bottom line: **P1** and **P2** do *not* provide evidential support for **P3**, and so do *not* make belief in **P3** reasonable. Though people might employ different epistemic standards suited to their differing cognitive abilities, E just doesn't make belief in **P3** rationally permissible. A *fully rational* agent whose body of evidence included **P1** and **P2** (and no defeaters) would *disbelieve* **P3**. No good at math? Can't concentrate? Sorry, it's still irrational for you to believe **P3** on the basis of E. Your evidence is squarely against it.

My second case is radical. Albert has a neurological peculiarity that makes him extremely prone to error in drawing out the logical implications of his beliefs. Accordingly, his epistemic standards include a rule along the lines of:

> **ES3-NO LOGICAL INFERENCES (Albert)**: For any logically complex belief, give no doxastic weight to its logical implications.

Albert justifiably believes **P4** (has no non-defeated defeaters, etc.):

> **P4:** Bertie is in the carriage and Alice is in the carriage.

Since his epistemic standards include no inference licence from 'P ∧ Q' to 'Q', he suspends on:

> **P5**: Alice is in the carriage.

19 The exact principles don't matter. I'm just firing an intuition.

20 I've been following Rosa in assuming that propositional and doxastic justification are conceptually linked. See Rosa (2012: 576–77) for the argument.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

Without knowing anything else about Albert's epistemic circumstances, I think I can say that his suspension on **P5** is at odds with his evidence. So I can say that it's irrational.[21] Whether or not Albert *can* believe **P5** on the basis of **P4**, or can find his way from **P4** to **P5**, he *should*. His cognitive distinctiveness and rogue epistemic standards do not justify *suspending* on a proposition that a fully rational agent with the same evidence would *believe*.[22] Unless Simpson can identify some relevant difference between this *logic* case and *detectives*, his argument is toast.

The final case is inductive. Suppose that Philip's hippocampus was damaged in a riding accident at a fairly young age resulting in severe anterograde amnesia: He can form no new long-term memories. And suppose he never saw a beetle before the accident but sees one now for the first time.[23] So he believes:

P6: This is a blue beetle.

Since Philip can form no new long-term memories, he can't perform an enumerative induction. But since he learned before the accident that regularity is common in nature, one of his epistemic principles – suited to his unusual condition – is **ES4-QUICK INDUCTION**:

ES4-QUICK INDUCTION (Philip): If you observe one x that is F, believe that all xs are F.

So Philip believes **P7**:

P7: All beetles are blue.

While following **ES4-QUICK INDUCTION** might sometimes give Philip a true belief, it's not the case that a body of evidence including one observed x that is F supports the belief that all xs are F. An ideally rational agent would not infer that all xs are F from observation of a single Fx. An ideally rational agent would *suspend* on **P7** given this evidence. So Albert should too. It's the *only* permissible attitude.

If these wacky cases didn't shake your confidence in the *cognitive abilities* strategy for blocking *arbitrariness*, then try this: Take *any* proposition P and any body of evidence E about which you're feeling permissive. The particular values of the rival credences won't matter, so let's just say they're 0.3 for you and 0.8 for someone else. In *your* estimation, a credence of 0.3 in P is justified by your total evidence E. Now to implicate cognitive abilities, suppose the other person, Anoushka, has the same body of evidence E, is vastly more intelligent than you, and better at interpreting evidence. Her epistemic standards recommend a credence of 0.8 in P, so if you're permissive here, you judge that a credence of 0.8 is rationally permissible for her. But doesn't that give you a reason to think that

21 I suspect that Rosa would disagree with me here (see his 2012: 575). The best I can do for now is to point out how utterly strange it would be to say that a person can *rationally* believe that A and B are in a carriage while refusing to believe that B is. That seems to jar with our epistemic practice, since we routinely accuse people of irrationality for far less. See Dogramaci and Horowitz (2016: 131). If anything counts as irrational, then Albert's withholding on **P5** does – or else the rule, **ES3**, that licences withholding. For more argumentation along these lines, see my 2012 and 2014.

22 If that's not convincing, imagine this: Albert's programme tells him to infer '~Q' from 'P ∧ Q'. No cognitive idiosyncrasy could make that rational.

23 Assume he has no relevant background information on beetles either.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

0.8 is the right credence? She's smarter than you and her epistemic standards are better. Your credence of 0.3 is suboptimal precisely because a credence of 0.8 is prescribed by better epistemic standards and a more perceptive interpretation of the evidence. The 0.3 only looks right because you misjudge E. And given your views about Anoushka and her epistemic superiority, you can see that. So it would be arbitrary and irrational for you to stick with 0.3. It's off the mark, according to you.

Probably the best thing for Simpson to say at this point is that it's not about *better* or *worse* standards and outputs. Different standards working on the same evidence yield different – and equally valid – doxastic outputs.[24] It's the equal validity part that puzzles me: However you spin it, one of the credences is going to look better to you, so you'll see the standards and abilities that recommend rival credences as somehow worse and those credences as incorrect. So you can't be permissive. Perhaps this thought will make the impermissivist outlook less puzzling to those who agree with Simpson.

The main point of this section is that differences in cognitive ability produce no evidential slack. The *math*, *logic* and *induction* examples draw out the implications of being permissive in this way. Those implications are bizarre.[25]

5. FINE DISTINCTIONS AND SYSTEM OUTPUTS

A possible worry is that the exaggerated cases I put forward are too clean-cut and artificial to be persuasive. As the evidence gets messier, stubborn insistence on Uniqueness starts to feel somewhat dogmatic. My response, as above, is that easy, conspicuous cases like these can help us sort through harder, messier cases like *detectives*. I see no reason to think that they can't, and I see no contrast between them.[26] But there's a particular kind of messiness that might spell disaster for impermissivists of any description: ineliminable vagueness. When evidential circumstances are highly complex, when defeaters and defeater-defeaters run in both directions, it could be objectively vague whether or not some body of evidence E supports a particular proposition P or its negation (Matheson 2011: 370). When it is objectively vague whether or not some relevant bit of information – a fading memory, say, or a barely discernible sensation – actually belongs to E, the total evidential force of E with respect to P could be indeterminate as well.[27]

24 He says that neither Drew's nor Mars's standards are "rationally superior *per se*" (2017: 14). That may be so, but in any given case where they produce different doxastic outputs from the same evidential input, one of them will be in error. That's the thinking behind my examples, in any case.

25 An editor for this journal wonders whether some of what I say in this section commits me to logical omnipotence for ideally rational agents and the (admittedly awkward) implication that beings who fail to see the logical consequences of their justified beliefs are irrational in withholding on them. I'm afraid so – afraid, because the view is neither popular nor prima facie plausible. I've tried to sugar the pill a bit elsewhere: My strategy is to *deflate* epistemic oughts as applied to particular propositions and to construe (full-fledged) epistemic duties *imperfectly* (see Stapleford 2012, 2013, 2015a). Stapleford (2014) is also relevant, though Smithies (2015) is the place to start.

26 Those sympathetic to the *cognitive abilities* strategy may be able to draw a principled distinction between my cases and *detectives*. But it's work to be done.

27 I'm just assuming that a fading memory or a barely discernible sensation are the sorts of things that Matheson is thinking of. See Matheson (2011: 370, note 30).

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

Matheson, who identifies these two sources of vagueness, leaves us guessing what objective indeterminacy in the evidence really amounts to, so I'll make it intuitive with a lowly example from my childhood. It's an example of the second type – where the force and composition of a given body of evidence is indeterminate owing to vagueness regarding *possession*. As I mentioned above (footnote 18), I think evidence is fully mental. If I have no cognitive access to x, then x is not available evidence. The books on my desk contain plenty of information, but since I haven't read them, it's not information *for me*. Things that I have read and retained, by contrast, are a part of my evidence – as are current perceptions and memories. But sometimes it's unclear whether or not I'm actually in some particular mental state or can access some particular mental item. Now for the personal touch.

When I was about four we went to the circus. I'm quite certain of that. It was an indoor arena and we sat on bleachers. But then there's the question of a panther. It seems to me that one of two things happened: Either (1) they arranged to surprise the audience by letting the panther run up the stairs between the bleachers and do a lap behind us, or (2) I imagined how cool it would be if they did. This is a perfect example, since I can never pin down the evidence. Sometimes I remember it happening. Sometimes I'm convinced that it didn't. It's the sort of shtick that was common in the 70s, and it's just what a kid would make up. So my body of evidence is objectively indeterminate with respect to 'The panther did a lap behind us.' It all depends on that memory. But that's what I can't work out. So what am I supposed to believe? There doesn't seem to be a clear answer, since the boundaries of my evidence are hazy.

Vagueness poking holes in Uniqueness is music to permissivists' ears: Where evidence is indeterminate, the temptation is to argue that "there can be distinct doxastic attitudes that are such that one would be justified in adopting either of them towards a proposition" (Matheson 2011: 370). If the evidence is vague regarding P, the permissivist might reasonably see this as a failure of Uniqueness: More than one doxastic attitude would be permitted. I could reasonably believe that the panther did a lap and I could reasonably believe that it didn't.

Matheson pulls the plug on this reasoning: If it is indeterminate whether believing, disbelieving or suspending on P fits E, then E enjoins *no* doxastic attitude (Matheson 2011: 371).[28] We're simply at a loss for what to do when the evidence is vague, but that hardly implies that we can take our pick of attitudes or credences. Vague evidence permits nothing and enjoins nothing. So Uniqueness can live with indeterminacy.[29]

28 Evidential vagueness does not enjoin suspension, by the way, since being *indeterminate* as to supporting P or supporting not-P is not the same thing as being *evenly balanced* between P and not-P.

29 Castro and Hart (2018) describe a few cases that many would regard as permissive. Impermissivists, they say, may try to "co-opt permissive intuitions by adopting ranged credences" to deal with them (3). In other words, "while it seems like any of a range of attitudes is permissible in cases like this, what is actually required is the single broad attitude that encompasses all of these single attitudes" (2). While Castro and Hart make an impressive case against this impermissivist response, it seems to me that it's the wrong response. The impermissivist should follow Matheson in saying that the evidence is indeterminate in these scenarios and thus provides no instruction. Permissivists who take the range-of-permissible-attitudes route – as opposed to the broad, encompassing attitude route – must face a smart objection from Schultheis (2018). Briefly, Schultheis argues that it's always irrational to assign P the lowest (or highest) credence in a range of supposedly permissible credences.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

This is a crucial point for impermissivists, but our interest is the clue it provides for dealing with *fine distinctions*. Recall Thomas Kelly's Alpha Centaurian thought experiment. Kelly isn't suggesting that we are like the Alpha Centaurians, of course, but let's consider the possibility. Are we like them, epistemically speaking? There are actually two questions here. First, does *evidence* really parse this finely? Does it ever happen that the probability of P given one body of evidence is .83527 while the probability of P given another body of evidence is .83528? Maybe sometimes (in controlled experiments and card games), but I doubt if it happens very often when it comes to everyday events and the bodies of evidence that people typically possess. Did the dog eat my shoe again? It's doubtful that my evidence on this matter would discriminate between its being .83527 likely that the dog ate my shoe and its being .83528 likely.[30] Secondly, do our *credences* really parse this finely? Is there any psychological difference – any difference at all – between believing that the dog ate my shoe with a credence of .83527 and believing it with a credence of .83528? I'm oblivious to any such fine tuning of my beliefs. But it doesn't matter, actually. Dividing the question this way makes it easier to see that Uniqueness gets it right whether our epistemic circumstances are Alpha Centaurian or not. Let's see what Uniqueness has to say – what it should say – about various pairings between *degree of evidential resolution* (that was the first question) and *degree of credential resolution* (that was the second).

To simplify the problem, we'll assume that there are only two scales (this is an idealization): The evidence scale is either *finely parsed*, which means that it's capable of supporting a proposition to such precise degrees as .83527 and .83528, or else it is *coarsely parsed* and can support a proposition to such coarse degrees as …0.7, 0.8, 0.9…, but nothing in between. The credence scales will be the same: *finely parsed* means that our credences can taper to a .83527 degree of confidence and a .83528 degree of confidence (etc.); *coarsely parsed* means they get as sharp as …0.7, 0.8, 09…, but nothing in between.[31]

Now consider the possible pairings: *evidence* and *credences* each parse *finely* or *coarsely* – so four combinations.

Pairing 1: evidence parses *finely*; credences parse *coarsely*.

If the evidence supports 'The dog ate my shoe' to degree .83527, but our credences can only get as precise as …0.7, 0.8, 0.9…, then the rational thing to do is to believe that

30 This could be an example of objective indeterminacy in the evidence that doesn't depend on vagueness regarding possession. It connects with the first type of vagueness alluded to by Matheson – the kind grounded in complexity (Matheson 2011: 370).

31 A 0.8 credence on the *coarse* scale differs from a .83527 credence on the *fine* scale only inasmuch as the 0.8 is one of ten possible doxastic attitudes, .83527 is one of one hundred thousand possible doxastic attitudes, and .83527 is slightly stronger. But as a referee points out, the numbering system is not to be taken very seriously. I agree entirely. We could label the possible attitudes with positive integers, with fractions, or even like this: *belief*, *bbelief*, *bbbelief* …, where each additional 'b' indicates an increase in doxastic strength. It's just a question of *how many* distinct attitudes we're capable of forming (I doubt it's very many). I want to stick with decimals for three reasons: (1) I'm responding to Kelly's objection, and he uses decimals; (2) I want my thoughts to connect with the credence literature, where decimals prevail; (3) it makes it easier to line up my credence scales with my evidence scales. But take it all *cum grano salis*.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

the dog ate my shoe with a credence of 0.8. In other words, if, as a matter of psychological fact, there's no such thing as a credence of .83527, then the evidence doesn't require any such thing. Given the evidence I have, and given the type of system I am, the uniquely rational credence towards 'The dog ate my shoe' is 0.8. An ideally rational agent *whose credences parse coarsely*, would shift her confidence up or down to the next closest evidential value. If the only possible credential outputs for a system are coarse, then a 0.8 degree of confidence in P is the only fitting response to evidence that makes P .83527 likely. To extend Simpson's computational metaphor, this is a *system limitation* that shunts outputs along 'circuits' that are hardwired into it.

But what if the finely parsing evidence scale makes it 0.85 likely that the dog ate my shoe? Wouldn't it be okay for me to take the 'high end' credence of 0.9 *or* the 'low end' credence of 0.8, since E hits the midpoint exactly? One could argue that such a case is permissive.[32] But one shouldn't. If I went with a credence of 0.8, I would in effect be saying (arbitrarily) that the evidence makes a credence of 0.8 more reasonable than a credence of 0.9. But it doesn't. There is no 'closest credence' to degree 0.85 for a creature on the *coarse credence* scale we're imagining. So *neither* credence is prescribed.[33] (It's something like Buridan's ass for belief.)

The second pairing inverts the first:

Pairing 2: evidence parses *coarsely*; credences parse *finely*.

Just as a creature capable of both believing and disbelieving P should do neither when the evidence fails to discriminate between them, so a creature capable of having credences of .83527 and .83528 should do neither when the evidence fails to discriminate between them. Hairsplitting believers like us should form a credence of 0.8 when E makes P 0.8 likely, since a credence of .83527 *goes beyond* E.

The third pairing is an easy sell:

Pairing 3: evidence parses *coarsely*; credences parse *coarsely*.

If the evidence supports 'The dog ate my shoe' with a probability of 0.8, and I believe it with a credence of 0.9, I've departed from ideal rationality. And I shouldn't. I should align my credences with the evidence, and the match here is perfect – no slack.

The final pairing smells like trouble, but it's not:

Pairing 4: evidence parses *finely*; credences parse *finely*.

If *both* evidence *and* credences parse finely, then when the evidence makes it .83527 likely that the dog ate my shoe, I should form *that credence*. Why? Because I can and I ought to be rational. This is the right outcome. It only sounds funny to say that a .00001 misalignment constitutes a departure from ideal rationality because we naturally assume that

32 Ballantyne and Coffman probably would. The 'high end–low end' terminology is theirs (Ballantyne and Coffman 2011: 7).

33 Imagine a crude, *binary* scale with an *on–off* switch for belief and no setting for *suspension*. Now suppose the probability of P on E were precisely 0.5. Not refined enough to suspend, you wouldn't be free to believe *or* disbelieve P willy-nilly. *Neither* attitude would fit. Same idea here.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

neither evidence nor credences are that discriminating. (I'm sure they're not.) But if they both *were*, then clearly we should align our beliefs with the evidence as neatly as we can. Why not get that kind of precision in our beliefs if we're capable of it? Surely we epistemically ought to do what is epistemically best. Forget humans for a second: If two Alpha Centaurians, Binar and Proxima, had the same body of evidence E and E made it .83527 likely that the dog ate Binar's grav boot, then Proxima would think it's *crazy* that Binar's credence is .83528. Uniqueness gives the intuitively correct recommendation here too.

Now the permissivist may be ready with a response. Doesn't this view open the door to permissivism at another level? If the probability that the dog ate my shoe on E is .83527, then the uniquely rational doxastic response from an agent whose credences parse coarsely is 0.8. I doubt that our credences are even that fine-grained, but supposing that they are, E requires a credence of 0.8 from me and from any other agent with the same input. Is that right? *Any* agent? What about Proxima whose credences parse finely? Isn't the uniquely rational response from an Alpha Centaurian a credence of .83527? If so, then we may be facing permissivism at a higher level – permissivism between systems of different types. This mirrors the intrapersonal–interpersonal distinction one level up. The permissivist pushing this line might admit that intraspecies permissivism is ruled out, but that interspecies permissivism isn't.

This makes sense. The evidence is not going to enjoin outputs from us of a kind that systems like us don't produce. Consider Keplerites – aliens living on Kepler-442b (if that's possible). The doxastic outputs of Keplerites, we'll suppose, parse more coarsely than ours. Say they have a three-place control: belief, suspension, disbelief. Now suppose the probability of P on E is 0.8 and the probability of Q on E is 0.6. The uniquely rational response for us is a credence of 0.8 in P and 0.6 in Q. But for the Keplerite, whose outputs are limited to three, the uniquely rational doxastic responses are shunted up for P and down for Q – belief in P and suspension on Q. Sounds right to me. There's no arbitrariness in favouring the doxastic outputs that systems like us deliver and at the same time recognizing that other systems may be capable of issuing outputs that cleave more closely to the evidence. The same body of evidence E could enjoin a credence of 0.8 in P from me, a credence of .83527 from an Alpha Centaurian, and belief from a Keplerite (but never a credence of 0.8 from me, a credence of .51361 from the Alpha Centaurian, and disbelief from the Keplerite). Slack arises from a difference in doxastic resolution, but never from the weight of evidence.

Note that there is no tension here with my rejection of the cognitive abilities proposal for blocking *arbitrariness*. It is *not* a matter of allowing that differences in cognitive abilities change the character of evidential relations. They remain constant. It's simply a recognition that limitations on the *resolving power* of a belief-forming system not only can, but must, change the 'programme instruction' for producing a doxastic output. It's an unsupported empirical assumption I'm making that the resolving power of systems like us is constant, so that the rationally required output is constant.[34] If that assumption turned out to be false – though I can't imagine a good test – then I'd happily admit a degree of interpersonal slack.[35] But that would be perfectly consistent with my claim

34 I suspect that our doxastic resolving power is actually quite low. The difference between believing P with a 0.7 degree of confidence and believing P with a 0.8 degree of confidence is introspectively invisible to me. In other words, I don't think there is a difference (and that sits well with my impermissivism-for-humans idea).

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

that differences in cognitive ability make no impact on the evidential relations themselves.[36]

6. CONCLUSION

The relativizing of evidential relations and the accompanying arbitrariness of favouring a single credence out of the set of permissible credences is what makes permissivism seem so wrong to me. But the permissivist could easily turn this around and accuse me of going too far in the opposite direction. There's something quasi-Platonistic about the conception of evidential relations I've been assuming. This is a fair criticism. Robust Platonism is a non-starter and I'd like to keep my distance. But Platonism may be a remote ancestor of impermissivism insofar as impermissivists are *rigid* on evidential relations. That's sort of what it means to be impermissive. I've tried to make a case for a strong conception of evidential relations elsewhere (Stapleford 2012, 2014), and to work out a theory of epistemic requirements on its basis (Stapleford 2013, 2015a, 2015b, 2016). And though I think the view is defensible, its descent from Platonism is arguably a weakness. It's a debate worth having, in any case.

Recent attacks on impermissivism suggest that it's dead in the water. My arguments to the contrary are very limited. I wanted to show that: (1) the *arbitrariness* objection to permissivism withstands appeals to cognitive abilities; (2) the *fine distinctions* objection to impermissivism is unconvincing. Hardly a grand slam, but it's proof that the game isn't over.

REFERENCES

Ballantyne, N. and Coffman, E. J. 2011. 'Uniqueness, Evidence and Rationality.' *Philosophers' Imprint*, 11: 1–13.

Brueckner, A. and Bundy, A. 2012. 'On 'Epistemic Permissiveness'.' *Synthese*, 188: 165–77.

Castro, C. and Hart, C. 2018 (Forthcoming). 'The Imprecise Impermissivist's Dilemma.' *Synthese*.

Conee, E. 2010. 'Rational Disagreement Defended.' In R. Feldman and T. A. Warfield (eds), *Disagreement*, pp. 69–90. Oxford: Oxford University Press.

Dogramaci, S. and Horowitz, S. 2016. 'An Argument for Uniqueness About Evidential Support.' *Philosophical Issues*, 26: 130–47.

Douven, I. 2009. 'Uniqueness Revisited.' *American Philosophical Quarterly*, 46: 347–61.

Feldman, R. 2007. 'Reasonable Religious Disagreements.' In L. Antony (ed.), *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, 194–214. Oxford: Oxford University Press.

Horowitz, S. 2014. 'Epistemic Akrasia.' *Noûs*, 48: 718–44.

Kahneman, D. 2011. *Thinking, Fast and Slow*. New York, NY: Random House.

35 Interpersonal (intraspecies) permissiveness of that sort would be something different than the interpersonal permissiveness of interest to run-of-the-mill permissivists, however. It wouldn't be based on differences in epistemic standards (Schoenfield), for instance, or differences in cognitive ability (Simpson), but on significant differences in the resolving power of belief systems. (Call it *output dependent permissivism*, since the relativizing would occur entirely at the output end of the system.)

36 Look at it this way: I am fully impermissive within and across species for the evidential relations themselves.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press

Kelly, T. 2014. 'Evidence Can Be Permissive.' In M. Steup, J. Turri and E. Sosa (eds), *Contemporary Debates in Epistemology*, 2nd edition, pp. 298–311. Malden, MA: Wiley-Blackwell.
Matheson, J. 2011. 'The Case for Rational Uniqueness.' *Logos & Episteme*, 2: 359–73.
McCain, K. 2014. *Evidentialism and Epistemic Justification*. New York, NY: Routledge.
Rosa, L. 2012. 'Justification and the Uniqueness Thesis.' *Logos & Episteme*, 3: 571–7.
Schoenfield, M. 2014. 'Permission to Believe: Why Permissivism is True and What it Tells Us About Irrelevant Influences on Belief.' *Noûs*, 48: 193–218.
Schultheis, G. 2018 (Forthcoming). 'Living on the Edge: Against Epistemic Permissivism.' *Mind*.
Simpson, R. M. 2017. 'Permissivism and the Arbitrariness Objection.' *Episteme*, 14: 519–38.
Smithies, D. 2015. 'Ideal Rationality and Logical Omnipotence.' *Synthese*, 192: 2769–93.
Stapleford, S. 2012. 'Epistemic Duties and Failure to Understand One's Evidence.' *Principia*, 16: 147–77.
—— 2013. 'Imperfect Epistemic Duties and the Justificational Fecundity of Evidence.' *Synthese*, 190: 4065–75.
—— 2014. 'Completing Epistemic Oughts.' *Philosophical Forum*, 45: 133–48.
—— 2015a. 'Epistemic Versus All Things Considered Requirements.' *Synthese*, 192: 1861–81.
—— 2015b. 'Why There May Be Epistemic Duties.' *Dialogue*, 54: 63–89.
—— 2016. 'Epistemic Value Monism and the Swamping Problem.' *Ratio*, 29: 283–97.
White, R. 2005. 'Epistemic Permissiveness.' *Philosophical Perspectives*, 19: 445–59.
—— 2014. 'Evidence Cannot Be Permissive.' In M. Steup, J. Turri and E. Sosa (eds), *Contemporary Debates in Epistemology*, 2nd edition, pp. 312–23. Malden, MA: Wiley-Blackwell.

SCOTT STAPLEFORD is Professor of Philosophy at St. Thomas University, Fredericton, New Brunswick. He is the author of *Berkeley's Principles: Expanded and Explained* (2017, with Tyron Goldschmidt), *Kant's Transcendental Arguments: Disciplining Pure Reason* (2008), and various articles in epistemology and early modern philosophy.

https://doi.org/10.1017/epi.2018.14 Published online by Cambridge University Press