You Ought to Have Known: Positive Epistemic Norms in a Knowledge-First Framework

Jonathan Jenkins Ichikawa

August 18, 2022 forthcoming in *Synthese*; please cite published version

Abstract

There are two central kinds of epistemological mistakes: believing things you shouldn't, and failing to believe things that you should. The knowledge-first program offers a canonical explanation for the former: if you believe something without knowing it, you violate the norm to believe only that which you know. But the explanation does not extend in any plausible way to a story about what's wrong with suspending judgment when one ought to believe. In this paper I explore prospects for a knowledge-centering account of positive epistemic norms that describe epistemic duties to believe.

Introduction

There are two fundamental kinds of epistemic norms. *Positive* epistemic norms enjoin us to form beliefs under certain circumstances; *negative* epistemic norms enjoin us to avoid forming beliefs under other circumstances. In "The Will to Believe," William James (1896b) identified the central positive norm as a norm to believe that which is true, and the central negative norm as a norm to avoid believing that which is false. Navigating negative and positive epistemic norms — deciding when to believe, and when not to believe — is central to the practical challenge of lived epistemology.

One of the characteristic ideas of the knowledge first project is to let knowledge play some of the central theoretical roles, as opposed to other notions like truth, belief, and justification. For example, replacing a truth norm with a knowledge norm of belief better respects intuitions about intuitively unjustified or inappropriate beliefs that nevertheless happen to be true. It is unsurprising, then, that a negative knowledge norm on belief has proven an interesting and attractive approach, albeit a controversial one. This paper considers the prospects for knowledge-first positive and negative epistemic norms. It raises a challenge to the standard negative norm, and explores the prospects for developing positive norms. The paper proceeds in five sections.

In §1, I motivate the common-sense idea that there are positive epistemic norms. The aim is to demonstrate that there are positive epistemic norms, and therefore that any approach to epistemology that cannot explain them is thereby deficient.

In §2 I develop a critique of certain formulations of the negative knowledge-first norm, according to which, if one doesn't know something, they shouldn't believe it. The problem is that this norm implies that there can be nothing that one could permissibly have believed, but didn't. This is an

¹This paper, like the knowledge-first project generally, focuses on the epistemic evaluations of committal doxastic states like belief. There are other epistemic norms too, such as those governing credences; they are not my project.

implausible result — and one that flies in the face of the motivations for positive epistemic norms advanced in §1. So knowledge-first theorists interested in explaining negative epistemic norms must be careful to state the norm in other ways.

In §3 I take up the question of whether the knowledge-first approach can explain why it is sometimes, not merely epistemically permissible, but epistemically *obligatory* to believe. The simplest candidate positive norms in terms of knowledge face obvious problems. For example, it will not do to explain the positive norm as the injunction to believe that which one knows. The reason for this is simple: since knowledge entails belief, everyone by necessity believes everything they know. The fundamental positive epistemic norm is no mere triviality; it is a substantive normative constraint on thought.

§4 considers possible solutions to these challenges. However, I'll argue, it remains an open question, to what degree such strategies will ultimately prove consistent with a thoroughgoing commitment to the knowledge first project.

In §5 I consider a different kind of challenge to positive epistemic norms. There are cases where one is not required to consider a question, despite being in a position to know the answer to it. So a norm that requires belief in such cases is mistaken. This is a version of a general attack on positive epistemic norms that already exists in the literature. I will argue that it does require modifying the norms mooted in §4, but that it fails to generalize against all positive epistemic norms. It does, however, result in yet further pressure against certain strong interpretations of the knowledge-first project.

1 Positive Epistemic Norms

I started in the Introduction with the common-sense observation that there are two kinds of epistemic norms — negative ones prohibiting belief in some cases, and positive ones demanding it. However, despite its status as common-sense, and its well-known role in James (1896b), the idea of positive epistemic norms is controversial, and they are often ignored or under-emphasized in the literature. There are certainly some contemporary epistemologists who have explicitly endorsed positive epistemic norms, but they rarely emphasize arguments for them — they tend rather to write as if it were simply obvious that there are positive epistemic norms.²

²For example, Pace (2011, p. 246), Miracchi (2019), Ichikawa (2020), and Gardiner (2021) all commit to positive epistemic norms, but do little to actually establish them beyond pointing to cases that seem like obvious examples. Feldman (2000, p. 678)'s evidentialism implies the existence of positive norms, which Feldman recognizes and seems to consider unremarkable. However, situating Feldman's approach into my question here is somewhat complex, as he appears to employ quite a weak notion of belief, according to which one ought to believe *P* when the evidence makes *P* even slightly more likely than its negation, if one considers the question. In the knowledge-first tradition I'm focused on in this paper, belief is a much more committal attitude; I do not, and should not, believe that the role of a six-sided die will be less than 6. Simion (forthcoming) is a clear example of a project in a similar spirit to mine, explicitly committed to positive epistemic norms, and interested in explaining them in a knowledge-first framework. (In the draft of that paper that is currently available at the time I am writing, Simion does not engage with arguments against positive epistemic norms.) Francesco Praolini has told me he is developing arguments focused on establishing positive epistemic norms in work in progress.

Other authors emphasize positive norms requiring belief, but not specifically *epistemic* ones. Goldberg (2017), for example, defends interpersonal norms demanding knowledge (and so demanding belief), but these are norms that derive from our obligations to one another as people, not from epistemology itself.

Some of this discussion has played out in the context of discussions about the viability of so-called "epistemic deontology," in the light of the apparent involuntariness of belief. Defenders of the former, such as Feldman (2000) and Chrisman (2008), do explicitly defend the view that there are obligations to believe, but their focus is more on whether there are genuine obligations in the doxastic realm. The same is true of related projects like Kornblith (2001), which focuses on the relationship between epistemic obligations and epistemic internalism — they say little or nothing about the difference between obligations to believe, and obligations not to; they do not engage arguments to the effect that there are epistemic obligations, but only negative ones.

Some epistemologists have explicitly argued against positive epistemic norms.³ I do not think those arguments are successful; I will describe and respond to them in §5. This section buttresses the intuitive motivation for positive epistemic norms.

James's own discussion is not especially clear about whether the duty to believe the truth is a moral duty, a practical one, an epistemic one, or something else. These are importantly distinct. Stealing from needy people for selfish and trivial reasons is widely regarded as morally prohibited. Under certain circumstances, however — if detection is sufficiently unlikely, and one is sufficiently safe from sanction, such a theft might be in one's own personal practical interest. If so, stealing might be *practically* permissible, but *morally* prohibited.

Following a standard practice in contemporary epistemology, I assume that there are distinctively *epistemic* norms, too.⁴ Paradigmatically, they apply to beliefs — believing something might or might not be *epistemically* permissible. I make no assumption about whether epistemic normativity is ultimately explicable in terms of moral or practical normativity, but if it is, the connection is a subtle theoretical one: epistemic permissibility is logically independent from both moral and practical permissibility. (Standard cases where they are thought to come apart include irrationally high faith in oneself — which some think gives prudential advantages in some cases⁵ — or in one's partner or family member — which some think can be morally obligatory.⁶ Negative beliefs about a loved one that are well-supported by the evidence may be examples that pull in the other direction.⁷)

My project is epistemic; I defend positive epistemic norms. In light of these distinctions between kinds of norms, it is less obvious than it might at first appear whether there are positive *epistemic* norms. Indeed, it's quite plausible that James was describing positive *pragmatic* norms in his famous discussion, which particularly features the advantages of religious belief. William Clifford, to whom James was most directly replying, wrote in a distinctively moral tone in his (1879, p. 186) emphatic statement of the negative norm that "it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence."

Someone skeptical about positive epistemic norms, then, may try to explain away some of the intuitive motivations for them as responsive to a different kind of normativity. Consider, for instance, cases of testimonial injustice. Suppose someone tells me that my department is hostile towards disabled people. And suppose I know enough to make it clear that I should trust this person; they're familiar with the department, and honest, and have the lived experience that qualifies them to make such a judgment with authority. But suppose I don't just take them at their word; like so many, let's assume, I am reflexively suspicious of accusations of oppressive behaviour, when I haven't noticed it firsthand. (The people here seem so *nice*, after all!) I know that people, especially activists, sometimes get a little carried away in their well-intentioned efforts to improve the world. So I suspend judgment. This is a concerning accusation, and I don't want to rush to judgment! I'll make sure to be on the lookout for anything that looks like ableism going forward, though.

In this thought experiment I have done something wrong. I ought to have believed my colleague;

³E.g. Wrenn (2007), Nelson (2010), and Littlejohn (2012, pp. 46–8).

⁴I take no stand on the question of how or whether epistemic normativity feeds into an "all-things-considered 'ought". See e.g. Feldman (2000, p. 692), Rinard (2019, pp. 1936–7) for discussion.

⁵James (1896a, p. 59), for instance, describes a mountain-climber who will succeed at a difficult jump if and only if they believe that they will. See also e.g. Hazlett (2013, pt. 1), McCormick (2014), Marusic (2015), Rinard (2017, p. 126), Zimmerman (2018, p. 128), Reisner (2018), or Bondy (2019).

⁶See e.g. Keller (2004), Stroud (2006), and Atkins (2021).

⁷Some attribute a similar status for ideas that, they say, are empirically well-supported, but where believing them would be racist. See e.g. Gendler (2011) and Zimmerman (2018, p. 136). The idea that these beliefs are sufficiently empirically well-supported to satisfy the epistemic norms is controversial; some deny it by embracing moral encroachment, holding that the moral features of the situation raise the bar for epistemically permissible belief. See e.g. Basu (2019); Fritz (2020); Moss (2018). Others argue that even by ordinary standards, beliefs of this type are not epistemically justified. See Munton (2019).

⁸Fricker (2007, ch. 1).

that means there is a norm requiring belief in this case. But is it an *epistemic* one? This is a complicated question, because there are plausibly multiple interacting norms here. Perhaps most obviously, there are moral norms at play — I *morally* ought to have believed my colleague. This doesn't mean there isn't also an epistemic norm, but it does mean that someone who wished to deny the latter has an avenue open for doing so: they may accept the obvious intuitive verdict that it's improper to suspend judgment in a case like this, but explain it without invoking positive epistemic norms. Perhaps the *epistemic* norms are exhausted by the negative requirement that I avoid forming beliefs on inadequate epistemic grounds. The sense in which I ought to believe, perhaps, derives from moral norms that require me to respond appropriately to my colleague, or to be proactive in engaging with oppression, or something similar.⁹

I do not accept this strategy for explaining away the appearance of positive epistemic norms. I think epistemic injustices such as these comprise — and can be intuitively recognized as comprising — a *combination* of norm violations, both moral and epistemic. That is, I do not think the intuitions that condemn suspension of belief in this case are indiscriminate between epistemic and moral norms. ¹⁰ But I know from conversations with other philosophers that many do not share my confidence. So I will not rest my judgment here on such intuitions.

Instead, although I am interested in positive epistemic norms that are also morally significant, in making the case for positive epistemic norms, I will focus on examples that are less morally loaded. A case where it feels clear that one would do wrong by failing to believe, where there is no prospect for attributing that feeling to the recognition of a moral violation, can more directly contribute to the case for positive epistemic norms.¹¹

So, consider a boring case. Egerton is sitting in a park on a sunny day, and a person with a dog on a leash walks past him. The dog looks over at Egerton and wags their tail; they continue on their way and walk around the corner and out of sight. Egerton pays attention to all of this; he likes dogs and enjoys watching them. Later, his partner asks him if he saw any dogs today. (He knows that Egerton particularly enjoys watching dogs in the park.) Egerton replies that he doesn't know.

Egerton isn't lying or joking. He is sincere in his expression of ignorance. His memory is fine — he can clearly recollect the experiences just described. Still, he reflects on various possible ways that he might not have seen a dog. Maybe his memory is playing tricks on him. He has no reason to suspect this — he's thinking clearly, and rarely confuses things like this, but he leaves open the possibility that he has suddenly worsened his dispositions, or that he is making an uncharacteristic mistake. Or maybe his memory is working fine, but that wasn't really a dog — could it have been a cleverly disguised pig? Could it have been a robot? He didn't have a chance to check.

This isn't a high-stakes situation of the sort emphasized in the pragmatic encroachment literature. Egerton is just replying to his partner's question about whether he saw any dogs today. He is, I submit, doing something obviously wrong, in suspending judgement on the matter. He should believe that he did see a dog. He is being unduly skeptical.¹²

⁹Dandelet (2021, p. 498) gives one way of developing such a strategy. Moral and practical considerations, on Dandelet's view, come into play within the scope of the epistemically permissible. However, see Simion (forthcoming) for an argument that moral blameworthiness depends on epistemic norms.

¹⁰For the theoretical framework behind explaining away intuitions, see Ichikawa (2009).

¹¹Clifford (1879) seemed to think that epistemic norms — at least the negative ones he emphasized — always carried moral significance. (I say he *seemed* to, because he does not write entirely clearly about the distinction.) He argues, for example, that even insufficiently-justified beliefs about trivial matters are morally prohibited, because they will inevitably lead to careless belief-forming dispositions, which are ultimately likely to lead to moral harm. I do not know whether Clifford is right about this. If he is, the same should go for positive epistemic norms. But since the moral implications of these epistemic failures are controversial and far from obvious, focusing on these cases can still help us to *recognize* the operation of epistemic norms, independently from moral ones.

¹²Cases like this are part of what motivate discussion of normative constraints on suspension of judgment. See e.g. Sosa (2019) and Miracchi (2019). The idea that suspension can be prohibited is closely related to the idea that belief

Can Egerton's failure be explained as moral, rather than epistemic? I don't think so. Moral norms against lying are irrelevant; Egerton is sincerely and accurately reporting his doxastic situation. One might argue that he misleads his partner by conveying the false idea that he doesn't have good grounds on the question, but the case needn't be understood that way. Maybe his partner knows that Egerton is prone to bizarre fits of skepticism from time to time, and so would not be misled. And indeed, his communicative interaction with his partner is unnecessary to recognize the epistemic sense in which he goes wrong — his suspension of judgment *itself* is a mistake, and would be even if he didn't tell anyone about it. It does not seem to be primarily a moral one. Nor is his suspension of belief on this matter pragmatically detrimental to Egerton.

We can find further examples in failures of more abstract forms of reasoning. Consider Lewis Carroll's famous (1895) Tortoise, who was happy to accept that A, and also that if A, then Z, but remained entirely open-minded about whether Z. This was an epistemic failure — not because he believed something he shouldn't, but because he didn't believe something he should. Ichikawa and Jarvis (forthcoming) discusses errors of this kind, and some of their implications for traditional epistemic questions, in some detail. That paper focuses particularly on the implications for virtue epistemology; here, I will discuss the implications for the knowledge first project.

One more kind of example will complete the intuitive case for positive epistemic norms: consider someone in the grip of radical skepticism, who suspends judgment on all matters, or perhaps all but very few matters, in response to skeptical arguments. I do not mean someone who, following Hume's, is convinced that there is no adequate justification for ordinary beliefs, but persists in them anyway. The skeptic I have in mind conducts their doxastic life as the Pyrrhonians advised, actually suspending judgment on all substantive matters. (One might doubt whether there are actual such skeptics in the world, or whether, if they were, they would survive very long. I doubt this myself. But we can *imagine* such an one; I invite you to do so.)

Among the many intuitive things to say about such a skeptic is this: they are not proceeding well, epistemically. Since radical skepticism is mistaken, they err by living by it. Their error isn't the false belief that skepticism is true, for they lack that belief — they are a good skeptic who suspends judgment on that higher-order question, too! Rather, they err by failing to form many beliefs in matters sufficiently well-evident, not only to "justify" belief and make it *permissible*, but to *require* it. Timothy Williamson considers a case like this one, allowing that it might motivate positive epistemic norms, but does not pursue that idea further, instead continuing to focus, as most epistemologists do, on negative norms. Clayton Littlejohn (2012, p. 35) points out that one can easily conform to negative epistemic norms, simply by withholding judgment, embracing the result that epistemic norms should not provide positive guidance.

This is not the place to argue in further detail for the existence of positive epistemic norms.¹⁴ For now, I hope that I have at least made the prima facie case for the need for positive epistemic norms, and I'll ask the reader to accept it, at least for the purpose of argument, through the next few sections. (There are influential arguments in contemporary epistemology against the existence of positive epistemic norms; I'll discuss them in §5.)

How could we explain both positive and negative epistemic norms, in a knowledge-first frame-

can be required. This case also has important features in common with Mona Simion's (forthcoming) "Perceptual Non-Responsiveness" case.

¹³ The Pyrrhonist, if such a person is possible, complies with all three [negative] norms even in the sceptical scenario. ... Non-sceptics may find little to admire in the Pyrrhonist's self-imposed ignorance, especially when that ignorance concerns the needs of others. There may be positive norms for knowledge, such as a norm enjoining knowledge-gathering in various circumstances, and so positive as well as negative norms for beliefs." Williamson (forthcoming) Williamson's invocation of the needs of others may suggest that he thinks positive norms might be primarily moral ones.

¹⁴I take up that project further in Ichikawa (forthcoming).

2 The Negative Knowledge Norm of Belief

There is an obvious role for knowledge in negative epistemic norms. It is a canonical commitment of the knowledge-first program: it is the idea that knowledge is a — or perhaps *the*, or the *central* — norm of belief. Such a norm prohibits beliefs that do not constitute knowledge. This, perhaps, explains what goes wrong when one believes when one (epistemically) shouldn't. Many knowledge-first-friendly epistemologists have defended such norms.

The knowledge norm of belief is often described as analogous to the knowledge norm of assertion, and motivated on similar grounds. ¹⁵ But there are important differences that render the cases disanalogous in interesting ways. ¹⁶ Note that although these two principles behave in quite parallel ways —

 KA_g Knowledge is a governing norm for assertion.

 KB_g Knowledge is a governing norm for belief.

— these two generate very different kinds of verdicts:

 KA_{o-} If one doesn't know that P, one ought not to assert that P.

 KB_{o-} If one doesn't know that P, one ought not to believe that P^{17}

A defender of the knowledge norm of assertion will accept both KA_g and KA_{o-} . But although the knowledge norm of belief implies KB_g , it does not imply KB_{o-} . Indeed, KB_{o-} is quite implausible. My critique of KB_{o-} isn't, like many extant critiques of the negative knowledge norm of assertion, based on intuitions about the appropriateness of justified false beliefs; I am willing to grant, at least for the purpose of argument, that these cases are happily understood as blameless norm violations, just as knowledge-first theorists typically hold them to be. ¹⁸ Instead, the focus of my argument is on KB_{o-} 's implications on beliefs that one does not already hold. Some people are more cautious than they need to be, in forming beliefs — they suspend judgment, even though they had enough evidence to form an appropriate belief. (Given the knowledge norm, this would be a case where they had enough evidence to know that P, but their caution prevented their knowledge.) In such cases, they do not believe, and so do not know, that P. By the negative knowledge norm, therefore, they are not permitted to believe that P, contradicting the plausible assumption of the case.

The problem isn't merely that KB_{o-} isn't or doesn't imply a positive epistemic norm — it's that it is actually inconsistent with the central normative intuitions that motivate positive epistemic norms. Consider Egerton again, who is strangely and selectively skeptical about some of his memories and

¹⁵E.g. by Williamson (2000, 255-6), Adler (2002), Bird (2007, p. 95), and Sosa (2010, p. 48).

¹⁶The point of disanalogy I focus on is quite different from ones other theorists have identified. McGlynn (2014, pp. 32–33), Goldberg (2015, pp. 167–8), and Willard-Kyle (2020), for instance, each emphasize the publicity of assertion, in contrast with the privacy of belief. Like me, they provide reasons to expect these norms to differ from one another, but they will pull in quite different ways. (Indeed, Willard-Kyle's ultimate view ends up close to the converse of my own, accepting a knowledge norm of belief, but rejecting the knowledge norm of assertion in favour of a position-to-know norm of assertion, similar to the one I will discuss for belief in §4 below.)

¹⁷In these abbreviations, K is for knowledge, A and B are for assertion and belief, g is for governing, and o— is for a negative 'ought' norm.

¹⁸For critiques of the knowledge norm of belief along these lines, see McGlynn (2013), Hughes (2017), Schechter (2017, p. 138). For replies, see Littlejohn (forthcoming), Williamson (forthcoming).

perceptual experiences. He doesn't believe that he saw a dog today, and instead suspends judgment on the question. This, I suggested, is an epistemic error. He ought to have believed.

But KB_{o-} implies that Egerton ought *not* to believe that he saw a dog today. Since by stipulation he *doesn't* believe that he saw a dog today, and since knowledge requires belief, ¹⁹ Egerton doesn't know that he saw a dog today. So by KB_{o-} , Egerton *ought not to believe* that he saw a dog today. This is exactly the wrong result.

The problem is general: KB_{o-} implies that there is no such thing as suspending judgment when one ought to believe. If one suspends judgment, one *ipso facto* doesn't know, and so ought not to believe. This is an unacceptable result.²⁰

Some epistemologists have stated a knowledge norm of belief similar to KB_{o-} that inherits this problem. Williamson (2000, p. 256), for instance, gives "one should believe p only if one knows p," which implausibly implies that any time one doesn't believe p, one shouldn't believe p.²¹ (I suspect such theorists really intended a slightly different principle, perhaps KB_g , or a wide-scope norm, or one of the modifications I will consider below.)

The problem is specific to belief; it does not extend to the knowledge norm of assertion. This is because knowledge doesn't entail assertion. Although KA_{o-} does not imply that there are any things it would be an epistemic error not to assert, it is at least consistent with this idea.²²

It is standard, in epistemological discussions of justification outside of the knowledge first context, to distinguish between propositional justification and doxastic justification. Doxastic justification, also sometimes called ex post justification, concerns whether a given belief token meets the epistemic demands having to do with justification; propositional justification, also sometimes called ex ante justification, is less psychologistic; it does not require a belief with the given content.²³ A subject can be propositionally justified in some content P, whether or not they believe P. Epistemologists also often hold that one can be propositionally justified in P, and believe P, without that belief being doxastically justified — if, for instance, one believes on fallacious grounds instead of the good ones that one had available.

This consideration about negative norms motivates the idea that there might be reason for knowledge-first theorists to make use of a notion more similar to propositional justification — one neutral on whether someone actually believes. If one wishes to state the conditions under which it is epistemically prohibited to believe P, given a knowledge norm like KB_g , without presupposing that one believes P, one may need to invoke a notion nonequivalent to knowledge — perhaps the idea that one would know that P if one believed it. Or, perhaps the notion of being in a position to know that P might be an appropriate one to invoke, if we allow that one might not believe that P and nevertheless be in a position to know it.

¹⁹Some epistemologists have denied that knowledge requires belief — see e.g. Radford (1966). But most, including typical defenders of knowledge norms, accept this requirement.

²⁰See Whitcomb (2014, p. 93) for a similar argument against this way of characterizing the knowledge norm of belief.

²¹This quote is ambiguous between this and a different reading, concerning constraints on positive norms — namely, that if one doesn't know p, it is false that one ought to believe p (but allowing that it might be *permissible* to do so). In the broader context it is clear that Williamson's intention is the negative norm. Some philosophers give statements that are ambiguous in English between this problematic norm and a more plausible wide-scope one; e.g. Sutton's (2007, p. 19) "one ought not believe that p unless one knows that p," which, for reasons of charity, we should probably read as $\Box(\neg Kp \to \neg Bp)$, rather than the implausible $\neg Kp \to \Box \neg Bp$. Perhaps Williamson's statement quoted above was also intended that way with a wide-scope obligation, although I'm not sure the English sentence in question can deliver that reading without further punctuation. Elsewhere when he intends wide-scope norms, Williamson punctuates to indicate as much, as in e.g. "One must: assert p only if one knows p". Williamson (2000, p. 243)

²²However, I don't see any particular case for a positive *epistemic* norm of assertion, the way I argued in §1 that there is one for belief. So it would actually be less of a problem if one's assertion norm ruled that out.

²³On the distinction, see Firth (1978), Goldman (1979), Turri (2010), Ichikawa and Jarvis (2013, pp. 32–4, 162), or Melis (2018).

The idea of being in a position to know is sometimes invoked within the knowledge-first program, but it is often thought of as something of a placeholder or hypothetical posit, something less central or less serious. (I'll discuss this distinction in further detail below.) Many knowledge-emphasizing epistemologists don't seem to consider it particularly worthy of extended specific study and explication. Its role in negative epistemic norms on belief may be consistent with this stance — KB_g , which does not require an invocation of being in a position to know, can be an adequate statement of the negative belief norm, so long as one is careful not to conflate it with KB_{ϱ} .

But, I'll argue in the next section, it is difficult to articulate a plausible knowledge-first *positive* norm for knowledge without relying centrally on a state that does not imply belief. We do need something closer to propositional justification. One must either find such a notion that can be characterized in terms of knowledge, or give up on the idea of knowledge explaining all the central epistemic norms.

3 A Positive Knowledge Norm of Belief?

A positive epistemic norm would identify conditions under which belief is epistemically obligatory. This is a non-starter:

 KB_{0+} If one knows that P, one ought to believe that P.

 KB_{o+} is not a norm so much as a triviality; there is no possible way to violate this norm, since knowledge entails belief. It cannot explain what goes wrong in the intuitive cases of epistemically inappropriate suspension of judgment. And adding further necessary conditions to the left-hand-side of the norm will make the problem no better. To be adequate, the norm must be stronger, so its antecedent must be weaker than knowledge. In particular, it must not entail belief.

Something like the counterfactual condition gestured at in the previous section looks like an improvement. There, I observed that a negative knowledge norm of belief, if it is to articulate cases in which belief is epistemically prohibited, without presupposing that one actually has the belief, cannot be expressed directly in terms of knowledge. It must invoke something more like the counterfactual condition that one would know, if one believed. We were looking for something to play something closer to the role of propositional justification.

Only some formulations of the negative norms seemed to require such a notion. So if that were the full case for the theoretical need to invoke a belief-neutral category, knowledge-first theorists might not have much cause for worry. But the problem is much deeper, when it comes to positive epistemic norms. As KB_{o+} illustrates, there is no plausible way to articulate them without invoking something more similar to propositional justification.

Something like CKB_{o+} avoids at least this particular problem:

 CKB_{o+} If, if one were to believe that P, one would know that P, then one ought to believe that P.²⁴

This principle is awkward to read in English; it embeds a counterfactual condition as its antecedent. If we attach a label to the state of being such that, if one were to believe that P, one would know that P, we can state this candidate norm more perspicuously:

S is **knowledge-ready** for P iff: if S were to believe P, then S would know P.

 KRB_{o+} If one is knowledge-ready for P, then one ought to believe that P.

²⁴CK here is for *counterfactual* knowledge.

A virtue of the knowledge-readiness approach is that it coheres well with a version of the negative norm KB_{o-} described above. Using the same new notion, that norm would have it that

 KRB_{o-} If one isn't knowledge-ready for P, one ought not to believe that P.

And KRB_{o+} and KRB_{o-} together make a straightforward biconditional.²⁵

But there are two problems. One problem is that knowledge-readiness may not be the right role for these norms — consideration of some of the complexities of propositional justification and being in a position to know will point to competing candidates. And, I will suggest, it remains an open question, to what degree a satisfactory filling-in of those candidate states will prove consistent with thoroughgoing commitment to the knowledge first project. Another problem is that there are intuitive counterexamples to KRB_{o+} — counterexamples that will extend to all kinds of attempts to modify it by tweaking the knowledge-first stand-in for propositional justification.

Let's consider those challenges in turn. I'll consider a challenge to the suitability of knowledge-readiness for positive norms in §4, and the more general worries about positive norms in §5.

4 Knowledge-readiness, being in a position to know, and propositional justification

Knowledge-readiness, defined counterfactually, was one kind of knowledge-first attempt to capture a positive epistemic standing that does not depend on a subject having a given belief. It is an attempt to fill a role analogous to propositional justification, while still giving a central theoretical role to knowledge. Since the other notions it invokes are not distinctively epistemic — a knowledge-first theorist faces no pressure to explain counterfactual conditionals in terms of knowledge — it fits well with a knowledge-first stance.

But like counterfactual notions generally, knowledge-readiness depends on contingencies of the world in odd ways. And some of those dependences will generate counterintuitive results for a positive knowledge norm. Here is an example.²⁶

Rohan is searching all over his house for the missing necklace. He shouldn't bother — he has excellent reason to believe the necklace never made it into the house. His carefree niece, you see, has a long track record of losing the expensive jewelry she borrows, then concocting far-fetched stories about returning them somewhere in the house. (She can just never the life of her remember where!) His friends tell him they saw her last night, under the influence of a motley variety of party drugs, giving the necklace to a child who just looked "too precious". Rohan loves his niece and would forgive her for losing another expensive necklace if he thought that's what happened, but he gives her the benefit of the doubt — not outright believing, but remaining open-minded about the possibility that she's telling the truth. Maybe, he thinks — wildly optimistically — the witnesses were mistaken; maybe she didn't lose the necklace; maybe she really did come into the house in the middle of the night to return it, but can't remember where she put it, then slipped away unnoticed. So he is looking in every drawer and cupboard to see if she might have deposited it there.

²⁵This approach is similar to the biconditional offered by Sosa (2019, p. 362, n. 9): "I affirm ... on that question if and only if I thereby affirm aptly". Here Sosa focuses on objectives and skills, but he also argues (Sosa, 2010, pp. 48, 51) that the considerations generalize to norms.

 $^{^{26}}$ Whitcomb (2014, p. 96–7) offers a case with a similar structure, making a similar point. I largely agree with Whitcomb's discussion of that case, although I think it is helpful to distinguish, as he does not, the positive from the negative knowledge norms of belief. Whitcomb considers many examples of possible ways to draw the connection between knowledge and whether one ought to believe, but he exclusively considers norms that involve biconditionals that identify a condition as both necessary and sufficient for appropriate belief — either wide-scope norms of the form $\Box(S \text{ believes } P \leftrightarrow ___)$, or narrow-scope norms of the form $\longrightarrow \hookrightarrow \Box(S \text{ believes } P)$. §5 below gives reason to reject the assumption that the positive and negative norms are symmetric in this way.

Perhaps there is something morally virtuous about Rohan's agnosticism,²⁷ but it is epistemically bad. Rohan, who is about to open the third drawer on the right side of the credenza in the dining room, suspends judgment on whether P — the proposition that necklace is inside this drawer. He shouldn't do that — he should believe that $\neg P$. He has plenty of evidence — knowledge-level evidence — that the necklace is nowhere in the house, which is inconsistent with P. This is an example, like those given in §1, motivating the existence of positive epistemic norms.

Is Rohan knowledge-ready for $\neg P$? Is it the case that, if he were to believe $\neg P$, he would know $\neg P$? That depends on yet-unspecified features of the situation. There are various ways that he could have believed $\neg P$. If he hadn't extended such an extreme benefit of the doubt to his niece, for instance, he would have known that she'd lost the necklace, and so known that it wasn't in the house, and so known that $\neg P$. But what if that's not what would have happened, had Rohan believed that $\neg P$?

On standard treatments of counterfactual conditionals, what would have happened, had Rohan believed that $\neg P$, is a matter of what things are like at a particular subset of possible worlds in which, unlike the actual world, Rohan believes that $\neg P$. And that particular subset of possible worlds is characterized by those that are *close* to the actual world, according to a particular similarity metric over modal space. There are various possible things that could have happened, had I dropped the champagne flute — someone *could* have swooped in and caught it; it *could* have been empty, and shattered dry glass over the floor. But those aren't what *would* have happened — the *nearest* worlds in which I drop the flute are ones where it hits the floor and splashes brut on my trousers.

So, whether Rohan would have known that $\neg P$, had he believed it, isn't settled by the fact that he has a possible route to knowledge that was available to him. It needs to be the route he would have taken. In Lewis's framework, possible worlds where he takes that route must be *closer* than those where he takes a different route that would not yield knowledge.

Suppose this is not so. Rohan is actually agnostic about whether the necklace is in the drawer. Stipulate that the possible ways he might have ended up with knowledge are *more distant* than possibilities in which he believes that $\neg P$, but doesn't know it. Perhaps, for instance, Rohan has an unjustified superstition about this particular drawer, which tempts him towards groundless pessimism about its contents. He just has a bad feeling about this drawer. Since he hopes to find the necklace, his pessimism tempts him towards disbelief — towards affirming that $\neg P$.

He doesn't believe $\neg P$; he overcomes his pessimism in this instance, remaining agnostic until he checks. But his resistance to this bias is not particularly robust. Had he been in a slightly worse mood, or slightly more distracted, he could easily have let his pessimistic bias carry the day, resulting in the belief that $\neg P$. This would have been a true belief, but not knowledge. And this, let's assume, could have happened much more easily than the better epistemic scenario, where he believes $\neg P$ due to a more realistic perception of his niece.

So in the nearest worlds where he believes $\neg P$, he doesn't know that $\neg P$ — i.e., he is not knowledge-ready for $\neg P$. Nevertheless, this is still a circumstance where the positive epistemic norms do require him to believe. He should have believed that P. Moreover, he should have believed that P in the good way that he *could* have done, rather than the bad way that he *would* have

²⁷ Perhaps. Proponents of views like those defended in Keller (2004), Stroud (2006), and Atkins (2021) might think so. I do not assume that such views are correct — I am in fact skeptical — but I do not here assume or argue otherwise.

 $^{^{28}}$ The philosophical *locus classicus* is Lewis (1973). I do not assume the details of Lewis's account — the commitments about counterfactual conditionals I rely on are shared by a wide orthodoxy, including Stalnaker (1968), Kratzer (2012), Ichikawa (2011), Moss (2012), Lewis (2015) and many more. (Not all these accounts invoke the language of 'closeness' or 'similarity' of possible worlds, but they do all endorse the idea that the *possibility* of A & C doesn't suffice for its being true that $A \to C$. That's all I assume here. In particular, I make no assumption about the controversial 'conditional excluded middle'.)

done. This is a respect in which the counterfactual notion of knowledge-readiness is an inadequate facsimile of propositional justification.

I have alluded already to a different option. Perhaps, instead of invoking knowledge-readiness, one could invoke *being in a position to know*. Rohan is plausibly in a position to know that $\neg P$, even though the relevant counterfactual conditional doesn't obtain. Positive and negative norms along such lines might look like this:

 PKB_{o+} If one is in a position to know that P, one ought to believe that P.

 PKB_{o-} If one is not in a position to know that P, one ought not to believe that P.

These norms fare better with respect to the challenge I've been focusing on against the counter-factual KRB_o norms. And indeed, I don't have an argument against a role like this for the idea of being in a position to know.²⁹ But I do wish to highlight a challenge for this strategy in this particular dialectical context. I am looking for ways to explain positive epistemic norms particularly for knowledge-first theorists. The exact commitments of the knowledge-first program are rather vague and flexible — see Ichikawa and Jenkins (2017) for detailed discussion — but on at least many of its precisifications, a core idea of the knowledge first program is that knowledge itself is playing the key epistemic explanatory roles. To the degree to which one posits epistemic states that cannot be explained in terms of knowledge, one's commitments run against those of the knowledge first project.

Being in a position to know obviously bears some interesting relationship to knowledge. But I do not think it is a particularly promising candidate for explication primarily in terms of knowledge, without other irreducible epistemic phenomena also playing central roles. There are key questions about which features of a subject do and do not influence what one counts as "in a position to know". For example, there is a sense in which anyone who doesn't believe that *P* isn't in a position to know that *P*, since belief is necessary for knowledge. This is obviously not the intended sense. We're to hold fixed something like the subject's evidential base, but not their beliefs. And for some of a subject's psychological states, it's not particularly intuitive one way or the other whether they interfere with what one is in a position to know. We abstract away from Rohan's excessive charity towards his niece, for instance, when we judge him to be in a position to know. But what about someone's mathematical inability, which leads them to suspend judgment on some truths of arithmetic? Here it is far from obvious just how to apply the notion of being in a position to know.

These reflections give rise to two suspicions, both of which sit uncomfortably with the knowledge-first movement. First, being in a position to know, unless supplemented with richer analysis, seems insufficiently robust to play the key normative roles sketched for it here. Second, insofar as we do have an intuitive grip on its extension, it seems like it might be given in terms of notions like evidence or epistemic support or propositional justification. The position to know may be hiding invocations of other traditional epistemic objects of analysis, rather than replacing or explicating them.

I haven't given an argument against the prospects for an adequate knowledge-first treatment of being in a position to know. Instead, treat these considerations as raising a challenge.³⁰ If you are a knowledge-first theorist, you need to account for positive epistemic norms, using a knowledge-first methodology. If you wish to do so using the position to know, you owe a knowledge-first explanation of that status, which is no small task.

 $^{^{29}}$ I do have an argument against the particular letter of PKB_{o+} — I'll articulate it in §5 below.

³⁰This challenge is similar in spirit to one of the main ideas of Whitcomb (2014), namely that "[i]f a knowledge-first theory of what we should believe (and what we should withhold belief on) ... [cannot] be made to work, then that is some reason to reject the knowledge-first approach to epistemology as a whole." (p. 89.) But Whitcomb is neutral in his paper on whether this condition is met.

A knowledge-first theorist might object that I haven't given a new challenge to the knowledge-first project here, on the grounds that it was already committed to a robust notion of being in a position to know. They might point out that Timothy Williamson himself invokes the term in his (2000) Knowledge and Its Limits, the canonical central text of the project.³¹ But, while the position to know does play a central role in Williamson's famous anti-luminosity argument, that role is wholly negative — Williamson stipulates, in rather vague ways, a notion of being in a position to know, in the service of a reductio argument against the luminosity of nontrivial states. Williamson's commitments about knowledge are clear — it is a mental state, it is prime, it can be used to characterize evidence and the constitutive norm of assertion, etc. He makes no such positive claims about being in a position to know. He is explicit about his disinterest in precisifying it: "Although the notion of being in a position to know is obviously somewhat vague and context-dependent, it is clear enough for present purposes. The vagueness and context-dependence are in any case primarily the result of fudging in attempts to defend the views to be criticized below."32 If a knowledge-first theorist felt they needed to go beyond Williamson's commitments, positing explanatory roles for the position to know, they would face new challenges about further specifying and articulating that notion.³³ Yli-Vakkuri and Hawthorne (forthcoming) defend the idea that 'being in a position to know' has roughly the same meaning as 'could know', including all the imprecision and context-sensitivity of the ordinary-language modal 'could'. Such an approach, combined with a norm like PKB_{o+} , would face the kinds of norm-based challenges that many knowledge-first theorists have raised against contextualists about 'knows'.34

5 Permissible Ignoring, Permissible Ignorance

There is one more complication I'd like to discuss, for attempts to articulate positive knowledge-first norms on belief to stand alongside negative ones. This complication is independent of the worries in §4 about how to characterize an adequate knowledge-first facsimile of propositional justification. Let's assume for the purpose of argument that that challenge has been met — we have some epistemic condition, explicable in terms of knowledge, but not entailing belief, that shows why characters like Rohan ought to believe, in cases like the one described above. Suppose we have an adequate knowledge-first understanding of being in a position to know. Then these norms are back on the table as appealing candidates:

 PKB_{o+} If one is in a position to know that P, one ought to believe that P.

 PKB_{o-} If one is not in a position to know that P, one ought not to believe that P.

But there is another problem with PKB_{o+} : there are counterexamples. There are plenty of situations where, at least intuitively, one is in a position to know something, but there's no epistemic

³¹Yli-Vakkuri and Hawthorne (forthcoming, n. 2) go so far as to attribute the wide contemporary use of the term to Williamson. But this seems to be an exaggeration: Shope (1983) uses the notion at length, attributing it to Ernest Sosa (1974; 1979). Thanks to Steven Diggin for pointing this out to me.

³²Williamson (2000, p. 95). See also Srinivasan (2015, n. 4) for related discussion.

³³Willard-Kyle (2020) defends the view that being in a position to know is the central norm of assertion; his view too fits poorly with the knowledge-first project for this reason. Consistency with that project is not one of his stated ambitions, but he does describe his view as one that preserves "Williamson's (etc.) insight that knowledge is central to assertion." (p. 348) It does so only on the assumption that knowledge is central to being in a position to know, which is not something to take for granted in this context. Indeed, Sosa (1974), which seems to have introduced the phrase 'position to know' into the epistemological canon, attempted to use it in the service of an analysis of knowledge.

³⁴See e.g. Hawthorne (2004, pp. 86–9) and Williamson (2005). See Ichikawa (2017, esp. p. 8) for a catalog of possible contextualist responses.

problem with failing to believe it.³⁵ I have in mind cases where it is permitted to have no opinion whatsoever on a given question.³⁶ Consider Rohan again, who is searching through his house for a lost necklace. A moment ago, he opened the second drawer, and found no necklace. He was focused on his task, and so paid scant attention to the contents of the drawer, beyond observing that it contained no necklace. In fact, it contained a nice ink pen, a cigar box, and his father's leather-bound journal. None of this was at all surprising or interesting to Rohan; he took it in at a glance, then closed the drawer and moved on to the third one.

Focus on the moment just after he closed the second drawer. He believes there was no necklace in it — he was focused on that question, and answered it perceptually. Set aside the necklace; consider the proposition that *J*: *Rohan's father's leather-bound journal is in the second drawer*. Here are three things that I think are true in this case, at the instant just after Rohan closed the second drawer.

- 1. Rohan does not have a belief one way or the other as to whether J.
- This is fine. I.e., it is no epistemic norm violation for Rohan not to have a belief one way or the other as to whether J.
- 3. Rohan is in a position to know J.

I actually don't think any of these three claims are obvious, but I do think they're all true. I'll say a bit in their defence below. First, notice that if they are true, they constitute a counterexample to PKB_{o+} . There's something Rohan doesn't believe (1), that he is in a position to know (3), and this is fine (2) — but this is exactly the sort of belief that PKB_{o+} condemns.

My argument against this positive epistemic norm is actually just an instance of the more general argument that Mark Nelson gives against *all* positive epistemic norms in his aptly-titled (2010) "We Have No Positive Epistemic Duties".³⁷ As I'll explain below, I think the argument is successful against PKB_{o+} , but that Nelson is too quick to assume it will generalize to all positive epistemic norms. Nelson is focused primarily on the norm that one ought to believe that which one's evidence supports; he points out that there are far too many such propositions for any subject to believe, and so dismisses such injunctions as unrealistic. And just the same, I think, goes for propositions that one is in a position to know, or for any other suitable candidate for an epistemic condition that can play this role in the context of a positive norm. (Recall that a central upshot of §3 is that the epistemic condition cannot itself entail belief.)

One might try to resist Nelson's argument, and mine, by denying condition (1) in this and similar cases. Is it so clear that Rohan doesn't believe that J? Maybe Rohan tacitly believes that J — the fact that he isn't consciously entertaining the thought doesn't mean he doesn't have it as a belief. Whether such a defence could be sustained depends on deep and substantive questions about the nature and scope of tacit belief. I do think this line may work for some of Nelson's cases. But I don't think it will work for all cases — in particular, I don't think it will work for this one. The exact timing of our point of evaluation matters. This is a point shortly after he's already closed the

³⁵The same goes for other candidate facsimiles for propositional justification. So my argument does not turn on the invocation of the position to know in particular. It extends, for instance, against the "sufficient evidence" criterion for duties to believe discussed in the drafts of Simion (forthcoming) currently available at the time I'm writing.

³⁶What's important for these counterexamples is that it's *epistemically* permissible not to have an opinion; it doesn't matter, for the evaluation of the principle, whether there might be moral or other reasons to be opinionated. However, as in my §1 discussion of the intuitive case for positive epistemic norms in the first place, the clearest examples of such cases will be ones where no norm *whatsoever* is violated by ignoring the question. So I'll focus on such cases here.

³⁷There are related arguments in Whiting (2012, p. 292). Greenberg (2020) gives a response in a similar spirit to mine. (But he is focused on a truth norm, not a position-to-know norm.)

³⁸For example, a person with normal vision who sees but does not pay attention to the birds flying in front of them might tacitly believe, even tacitly know, that there are birds flying in front of them. Nelson (2010, p. 87)

drawer. The information in question is fleetingly stored in working memory, but it will be gone in a moment. I believe in tacit belief, but I don't think it extends to cases like this, where no lasting cognitive impression is formed at all. I don't think it's particularly plausible that Rohan now believes that *J*, even if he did tacitly believe it a second ago.

A second avenue of resistance would deny the second condition, affirming instead that there *is* a norm violation in Rohan's agnosticism. I myself am much more sympathetic to very demanding epistemic norms than most people are. I'm committed, for example, to the idea that one violates epistemic norms by being actively agnostic about whether a given sentence is a logical truth, even if its proof is far too complex for a normal person to grasp in an instant.³⁹ (If, like most people, you disagree with me about that idea, then I should have an *easier* time convincing you that Rohan's agnosticism is fine; my view on the demandingness of rationality amounts to a concessive assumption in this context.) But focusing on a question and *wondering* about it, despite your having totally conclusive evidence for it, is quite a different matter from refraining from giving a second's thought to a question that one doesn't care about (or even have any *reason* to care about). An ideally rational agent would always judge accurately on any question they consider, if their evidence is conclusive for the answer. But there simply isn't a similar rational requirement to be opinionated on every question, regardless of one's interests.⁴⁰

The third option is to deny (3), saying that Rohan isn't in a position to know that J. Whether this is sustainable, again, will depend on the precise characterization of being in a position to know, but I doubt that in the end this will be a very plausible result. It is relevant once again that we are discussing a time just after he closed the second drawer, when his visual experience of its contents are retained in working memory. Perhaps thirty seconds later, by the time he's rejected his superstitious instincts and opened the third drawer, that information will be inaccessible, and he won't even be in a position to know where the journal is. But now, just after closing the drawer, he has available an epistemic route to that knowledge. If you asked him, at this exact moment, what was in the drawer, he would reflect on his experience and tell you. His actual failure to know that J is a mere result of his indifference to the question. He satisfies Williamson's (2000, p. 95) condition, that "no obstacle must block one's path to knowing p." It would be surprising if the eventual knowledge-first notion of being in a position to know didn't include this case in its extension.

So I do think there is a strong prima facie problem for invoking the position to know, or indeed any similar state, in a positive epistemic norm along these lines. But I do not think, as Nelson does, that this problem is fatal for positive epistemic norms generally. Indeed, I think the explanation of the problem in this case points pretty straightforwardly in the direction of some plausible solutions. The reason it's fine for Rohan not to believe J is that he's not wondering about whether J, nor has he any reason to do so. The positive epistemic norm of belief shouldn't say that anyone must believe anything they're in a position to know; rather, it mandates belief when one is in a position for knowledge, and the question comes up.⁴¹

This is why I didn't say, in my description of the case, that Rohan *suspends judgment* on whether J - I said that he does not have a belief one way or the other on the matter. There is an important distinction between suspending judgment and merely *lacking* judgment. Suspension is a mental activity; lacking judgment is, well, the lack of one.⁴²

³⁹Ichikawa and Jarvis (2013, pp. 284–88), Ichikawa and Jarvis (forthcoming). See also Smithies (2015) for a similar commitment.

⁴⁰In Ichikawa and Jarvis (2013) we did not emphasize the relevance of the question arising, although we did have it in mind, and stipulated that feature in the relevant cases where we said there was a rational requirement to believe, e.g. on pp. 132, 166.

⁴¹Feldman (2000, pp. 678–9) anticipates the Nelson-style objection to positive norms, and makes a move quite similar to the one I here suggest. As indicated in fn. 2, however, Feldman is working with a much weaker conception of belief than I am.

⁴²On the activity of suspension and the distinction, see Friedman (2013, 2017), McGrath (2021).

There are various more specific ways to develop the thought. I've alluded to two, in my disjunctive statements already — the positive norm activates when one considers, *or should consider*, the question. These correspond to two possible candidate norms:

 QP_{o+} If one considers the question whether P, and is in a position to know that P, one ought to believe that P.

 $\Box QP_{o+}$ If one considers or should consider the question whether P, and is in a position to know that P, one ought to believe that P.⁴³

As has been the case throughout this paper, the 'ought' in question is an epistemically normative one; it describes epistemic obligations. This contrasts with e.g. moral obligations, as well as the so-called 'ought of expectation'. (QP_{o+} is not about the idea that we would *predict* that someone in a position to know who considers a question will believe it.)

Deciding between these two norms is beyond the scope of this paper; they each correspond to a plausible way to develop a positive norm of belief.⁴⁴ They each explain why Rohan's lack of belief in J is fine, even though his lack of belief in not-P isn't. The latter, but not the former, is an active suspension of judgment on a question he cares about.

Nelson does not consider positive epistemic norms that have this kind of structure, assuming instead that a positive norm would require cases "where one epistemically *ought* to believe p (i.e. do more than merely withhold regarding $\neg P$, and where this 'ought' is grounded wholly in one's epistemic circumstances and not also in the aims, desires, moral duties, etc. of the agent.)" Nelson also has in mind a quite narrow notion of 'epistemic circumstances', according to which they are "roughly, those aspects of our circumstances that count in favour of the truth or falsity, probability or improbability, of certain propositions." My candidate norms do seem to violate Nelson's constraint.

Nelson is not alone in positing such a constraint. There are similar commitments in the work of Jonathan Sutton, Chase Wrenn, and Clayton Littlejohn.⁴⁷ But what justifies it as a constraint? What is wrong with the idea that there are positive epistemic norms, but that whether they apply depends in part on non-epistemic considerations?

In fact, some of the considerations a positive norm theorist may invoke along these lines arguably *are* epistemic (though not by Nelson's restrictive standards). Questions about attention and what questions one is considering do feel well within the realm of epistemology. There are, to be sure, *some* epistemic traditions according to which epistemology is about the processing of passively-received data, whose possession itself is not subject to epistemic norms. But these approaches are (a) out of fashion, (b) quite implausible, and (c) antithetical to the knowledge-first project with which this paper primarily engages. So I set them aside. I see no dialectically appropriate grounds

⁴³"Q" here is for considering the *question*, and "□" is for the idea that one is *required* do so. I have omitted the "K" from the previous names, letting "P" stand for *position to know*. I also omit the "B" for *belief* here, as I am only considering belief norms at this time.

 $^{^{44}\}Box QP_{o+}$, with its more normative condition, may have a more intuitive extension, insofar as one thinks that one makes an epistemic mistake by failing to consider, and know, something they that was important to their interests. Cf. Greenberg (2020, p. 3287). But it also comes with deeper theoretical challenges, especially if one is committed to a knowledge-first framework — one owes a story about the circumstances under which one should consider a question.

⁴⁵Nelson (2010, p. 89), emphasis in original.

⁴⁶Nelson (2010, p. 86).

⁴⁷Sutton (2007, p. 19): "I will not in general be concerned to argue that there are any positive epistemic obligations, and, indeed, I am inclined to think that the vast majority of beliefs that one ought to hold are such for nonepistemic reasons." Wrenn (2007, p. 117): "Epistemic duties are doxastic duties that are grounded in purely epistemic considerations, such as what evidence one has." Littlejohn (2012, p. 48): "If the ends were not epistemic in nature, it is not clear why there would be distincively epistemic obligations to form the beliefs that suited these non-epistemic ends."

for ruling out the idea that what one is attending to can influence one what may or must believe. So a norm like QP_{o+} , which makes positive norms sensitive to what question one is considering, is a viable candidate.

The triggering conditions for the norm may well get less epistemic, though, if we consider norms like $\Box QP_{o+}$, which invokes normative conditions about what question one *should* consider. Either that 'should' is or is not an epistemic 'should'. If it is, then we owe a further story about substantive epistemic norms that are not captured in norms on belief. (And if we are committed to the knowledge-first project, we face a challenging constraint on that story.⁴⁸) Moreover, it is prima facie just not particularly plausible that the relevant facts about what one ought to consider are independent of nonepistemic practical or moral matters.

If our positive epistemic norm is sensitive to what one should consider, where this 'should' is a *non*-epistemic should — a practical or a moral one, for instance — then $\Box QP_{o+}$ certainly violates Nelson's constraint, even on a looser interpretation of 'epistemic circumstances' that includes attention. Does this mean it can't be a genuinely epistemic norm? I don't see why. Here, once again, is Nelson's (2010, p. 89) statement of what a positive epistemic norm would imply. I have added italic emphasis, to highlight his two separate, non-redundant, invocations of the the 'epistemic' qualifier: "where one *epistemically ought* to believe p (i.e. do more than merely withhold regarding $\neg P$, and where this 'ought' is grounded wholly in one's *epistemic circumstances* and not also in the aims, desires, moral duties, etc. of the agent." The first italicized term characterizes the kind of normativity at issue — this is epistemic normativity, rather than moral normativity, etc. It characterizes the respect in which someone who violates the norm is failing. The second italicization places a restriction on what sort of considerations could ground the relevant norms. It is the first that is clearly a legitimate constraint — our topic is epistemic normativity. But why the second one? Nelson may simply be *assuming* that only epistemic circumstances could ground epistemic obligations. Without that assumption, the second restriction is unwarranted.

Is the assumption true? Is there a problem with the idea that epistemic norms may encode sensitivity to non-epistemic factors? This is far from a trivial matter; it would require argument. Notice, for instance, that it is *not* generally assumed that a similar restriction holds for *negative* epistemic norms. The following is a controversial and substantive thesis, corresponding to a live and active debate in contemporary epistemology:

Purism Only factors relevant to the truth of P can influence whether it is permissible to believe P.

While many epistemologists have defended purism, it is denied by defenders of pragmatic encroachment and moral encroachment.⁴⁹ They hold that whether it is epistemically permissible to believe something can depend on how important the question is (one requires a higher evidential bar for belief in high-stakes situations) or one whether there is a risk of perpetuating a moral harm by believing it (one requires a higher evidential bar for belief in contents that correspond to racist stereotypes). Such views are very much live in the epistemological literature. One does not dismiss

⁴⁸Sosa (2019, p. 361) notes that "[it] is crucial in epistemology to distinguish theory of knowledge from theory of inquiry." If Sosa is right, and if norms on belief depend on norms of inquiry, then the knowledge first project is bound to be seriously incomplete.

⁴⁹Details of theorists' commitments vary; one might defend pragmatic encroachment *about knowledge*, which would commit one to the denial of purism if one adopts a knowledge norm for belief. Or one might accept a different norm on belief, but accept pragmatic encroachment about it. See Ichikawa et al. (2012); Rubin (2015) on the significance of these distinctions. Pragmatic encroachment is defended by Fantl and McGrath (2009); Kim (2016); Stanley (2005); Weatherson (2011, 2017). Moral encroachment is defended by Basu (2019); Bolinger (2020); Fritz (2020); Moss (2018); Pace (2011). The "subject-sensitive invariantism" of Hawthorne (2004) also denies purism, even though it is arguably not a pragmatic or moral encroachment view.

them as incoherent on the grounds that they posit epistemic norms that depend on non-epistemic considerations.⁵⁰

Since epistemologists do not typically assume that, in order to be a genuinely epistemic negative norm, the norm in question must depend only on epistemic considerations, I see no reason to accept the corresponding constraint on epistemic positive norms either. There is no contradiction in the idea that an epistemic norm requires a subject to believe something under circumstances which need to be specified in non-epistemic terms.

Notice also that it is quite ordinary to see positive *moral* norms whose triggering conditions are non-moral. "If one is in a position to save someone's life at small cost to one's own well-being, one is required to do so" is a canonical example of a positive moral norm, despite the left-hand-side invoking many non-moral factors. So too with legal norms. In my community, if snow falls on the sidewalk in front of one's house overnight, one is required to clear the snow by 10am the next morning. This is a genuine legal norm, despite the fact that 'snow falling' isn't a legal notion. ("One's house" is a legal notion, but the principle I'm challenging says that *only* notions corresponding to the kind of normativity in question can play roles.) Such a requirement for epistemic norms seems to me unmotivated.

One might object to this line of argument by insisting that the relevant comparison isn't to whether the these are moral or legal *notions*, but rather that they are legally or morally *relevant considerations*. And it is certainly true that, in part because of these norms, they are relevant. Whether it snowed last night is legally relevant — it can make a difference to one's legal requirements. And whether I am in a position to save someone's life is morally relevant to my moral obligations. But in *this* sense, it is trivial that any condition that plays a role in triggering the norm in question is, in the relevant sense, relevant. According to pragmatic encroachment theorists, the stakes of the question are epistemically relevant, in the sense that they can make a difference as to what one ought to believe. And according to the positive epistemic norms QP_{o+} and $\Box QP_{o+}$, whether one is considering a question, or whether one ought to consider a question, is epistemically relevant in just the same way. To assume otherwise would be blatantly question-begging.

So I do think that those amendments to the PKB_{o+} norm are plausible candidates for positive epistemic norms. The challenge for the knowledge-first theorist is, they take us yet another step away from a central theoretical role for knowledge itself — especially $\Box QP_{o+}$, which builds in considerations about the normativity of inquiry. This isn't necessarily *inconsistent* with the knowledge-first project, whose precise commitments are vague. But it is another respect in which we posit central epistemic roles to states other than knowledge.⁵¹

Bibliography

Adler, J. 2002. Belief's Own Ethics. MIT Press.

Atkins, J. Spencer. 2021. "Epistemic Norms, the False Belief Requirement, and Love." *Logos and Episteme* 12:289–309.

⁵⁰Indeed, the very notion of a non-epistemic consideration becomes somewhat obscure under close scrutiny. See Ichikawa (2017, pp. 32–3).

⁵¹Thanks to two referees for this journal for helpful feedback. For further helpful conversations and advice on topics related to this paper, I'm grateful to Matt Benton, Pat Bondy, Charles Côté-Bouchard, Marc-Kevin Daoust, Will Fleisher, Daniel Greco, Shane Hendrix, Carrie Jenkins, Chris Kelp, Clayton Littlejohn, Graham Seth Moore, Jennifer Nagel, Ángel Pinillos, Kathryn Pogin, Francesco Praolini, Olle Risberg, Mona Simion, Daniele Sgaravatti, Maxim Smyrnyi, Dennis Whitcomb, and Christopher Willard-Kyle. I'm especially grateful to Alisabeth Ayars, Steven Diggin, Devin Morse, Carrie Jenkins, and Jonathan Weinberg for detailed comments on an early draft of this paper. Work on this paper was supported in part by a SSHRC Insight Grant on rape culture and epistemology.

- Basu, Rima. 2019. "Radical Moral Encroachment: The Moral Stakes of Racist Beliefs." *Philosophical Issues* 29:9–23.
- Bird, Alexander. 2007. "Justified Judging." *Philosophy and Phenomenological Research* 74:81–110.
- Bolinger, Renée Jorgensen. 2020. "Varieties of Moral Encroachment." Philosophical Perspectives 34:5–26.
- Bondy, Patrick. 2019. "The Epistemic Norm of Inference and Non-Epistemic Reasons for Belief." *Synthese* 1–21.
- Carroll, Lewis. 1895. "What the Tortoise Said to Achilles." Mind 4:278–280.
- Chrisman, Matthew. 2008. "Ought to Believe." Journal of Philosophy 105:346–370.
- Clifford, William Kingdon. 1879. The Ethics of Belief, volume 2, 177-211. Macmillan and Co.
- Dandelet, Sophia. 2021. "Epistemic Coercion." Ethics 131:489-510.
- Fantl, Jeremy and McGrath, Matthew. 2009. *Knowledge in an Uncertain World*. Oxford University Press.
- Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60:667–695.
- Firth, Roderick. 1978. "Are Epistemic Concepts Reducible to Ethical Concepts?" In Alvin Goldman and Jaegwon Kim (eds.), *Values and Morals: Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt*, 215–229. Dordrecht: Kluwer Academic Publishers.
- Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.
- Friedman, Jane. 2013. "Suspended Judgment." Philosophical Studies 162:165-181.
- —. 2017. "Why Suspend Judging?" *Noûs* 51:302–326.
- Fritz, James. 2020. "Moral encroachment and reasons of the wrong kind." *Philosophical Studies* 177:3051–3070.
- Gardiner, Georgi. 2021. "Banal Skepticism and the Errors of Doubt: On Ephecticism About Rape Accusations." *Midwest Studies in Philosophy* 45:393–421.
- Gendler, Tamar Szabó. 2011. "On the epistemic costs of implicit bias." *Philosophical Studies* 156:33.
- Goldberg, Sanford C. 2015. Assertion: On the Philosophical Significance of Assertoric Speech. Oxford University Press.
- —. 2017. "Should Have Known." Synthese 194:2863–2894.
- Goldman, Alvin I. 1979. "What is Justified Belief?" In G. S. Pappas (ed.), *Justification and Knowledge*, 1–25. Boston: D. Reidel.
- Greenberg, Alexander. 2020. "Should I Believe All the Truths?" Synthese 197:3279–3303.

- Hawthorne, John. 2004. Knowledge and Lotteries. Oxford University Press.
- Hazlett, Allan. 2013. A Luxury of the Understanding: On the Value of True Belief. Oxford University Press.
- Hughes, Nick. 2017. "No Excuses: Against the Knowledge Norm of Belief." *Thought: A Journal of Philosophy* 6:157–166.
- Ichikawa, Jonathan Jenkins. 2009. "Explaining Away Intuitions." *Studia Philosophica Estonica* 2:94–116.
- —. 2011. "Quantifiers, Knowledge, and Counterfactuals." Philosophy and Phenomenological Research 82:287–313.
- —. 2017. Contextualising Knowledge: Epistemology and Semantics. Oxford University Press.
- —. 2020. "Contextual Injustice." *Kennedy Institute of Ethics Journal* 30:1–30.
- —. forthcoming. *Epistemic Courage*. Oxford University Press.
- Ichikawa, Jonathan Jenkins, Jarvis, Benjamin, and Rubin, Katherine. 2012. "Pragmatic Encroachment and Belief-Desire Psychology." *Analytic Philosophy* 53:327–343.
- Ichikawa, Jonathan Jenkins and Jarvis, Benjamin W. 2013. *The Rules of Thought*. Oxford: Oxford University Press.
- —. forthcoming. "Hybrid Virtue Epistemology and the A Priori." In Dylan Dodd and Elia Zardini (eds.), *The A Priori: Its Significance, Grounds, and Extent.* Oxford University Press. In Press.
- Ichikawa, Jonathan Jenkins and Jenkins, C. S. I. 2017. "On Putting Knowledge 'First'." In J. Adam Carter, Emma Gordon, and Benjamin W. Jarvis (eds.), *Knowledge First*, chapter 6, 113–131. Oxford University Press.
- James, William. 1896a. "Is Life Worth Living?" In James (1896c), 32–62.
- —. 1896b. "The Will to Believe." In James (1896c), 1–31.
- —. 1896c. The Will to Believe: And Other Essays in Popular Philosophy. New York: Longmans, Green.
- Keller, Simon. 2004. "Friendship and Belief." Philosophical Papers 33:329-351.
- Kim, Brian. 2016. "In Defense of Subject-Sensitive Invariantism." Episteme 13:233–251.
- Kornblith, Hilary. 2001. "Epistemic Obligation and the Possibility of Internalism." In Abrol Fairweather and Linda Trinkaus Zagzebski (eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, 231–248. Oxford University Press.
- Kratzer, Angelika. 2012. "Conditionals." In *Modals and Conditionals*, 86–108. Oxford University Press.
- Lewis, David K. 1973. Counterfactuals. Malden, MA: Blackwell Publishers.
- Lewis, Karen S. 2015. "Elusive Counterfactuals." Noûs 50:286–313.
- Littlejohn, Clayton. 2012. Justification and the Truth-Connection. Cambridge University Press.

- —. forthcoming. "A Plea for Epistemic Excuses." In Fabian Dorsch Julien Dutant (ed.), *The New Evil Demon Problem*. Oxford University Press.
- Marusic, Berislav. 2015. Evidence and Agency Norms of Belief for Promising and Resolving. Oxford University Press.
- McCormick, Miriam Schleifer. 2014. Believing Against the Evidence: Agency and the Ethics of Belief. Routledge.
- McGlynn, Aidan. 2013. "Believing Things Unknown." Noûs 47:385-407.
- —. 2014. Knowledge First? Palgrave Macmillan.
- McGrath, Matthew. 2021. "Being Neutral: Agnosticism, Inquiry and the Suspension of Judgment." *Noûs* 55:463–484.
- Melis, Giacomo. 2018. "The Intertwinement of Propositional and Doxastic Justification." *Australasian Journal of Philosophy* 96:367–379.
- Miracchi, Lisa. 2019. "When Evidence Isn't Enough: Suspension, Evidentialism, and Knowledge-First Virtue Epistemology." *Episteme* 16:413–437.
- Moss, Sarah. 2012. "On the Pragmatics of Counterfactuals." *Noûs* 46:561–86.
- —. 2018. "Moral Encroachment." *Proceedings of the Aristotelian Society* 118:177–205.
- Munton, Jessie. 2019. "Beyond accuracy: Epistemic flaws with statistical generalizations." *Philosophical Issues* 29:228–240.
- Nelson, Mark. 2010. "We Have No Positive Epistemic Duties." Mind 119:83–102.
- Pace, Michael. 2011. "The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe'." *Noûs* 45:239–268.
- Radford, Colin. 1966. "Knowledge—By Examples." Analysis 27:1-11.
- Reisner, Andrew. 2018. "Pragmatic Reasons for Belief." In Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*. Oxford University Press.
- Rinard, Susanna. 2017. "No Exception for Belief." *Philosophy and Phenomenological Research* 94:121–143.
- —. 2019. "Equal Treatment for Belief." Philosophical Studies 176:1923–1950.
- Rubin, Katherine. 2015. "Total Pragmatic Encroachment and Epistemic Permissiveness." *Pacific Philosophical Quarterly* 96:12–38.
- Schechter, Joshua. 2017. "No Need for Excuses: Against Knowledge-First Epistemology and the Knowledge Norm of Assertion." In J. Adam Carter, Emma Gordon, and Benjamin Jarvis (eds.), *Knowledge-First: Approaches in Epistemology and Mind*, 132–159. Oxford University Press.
- Shope, Robert K. 1983. *The Analysis of Knowing: A Decade of Research*. Princeton: New Jersey: Princeton University Press.
- Simion, Mona. forthcoming. "Resistance to Evidence and the Duty to Believe." *Philosophy and Phenomenological Research* forthcoming.

Smithies, Declan. 2015. "Ideal Rationality and Logical Omniscience." Synthese 192:2769–2793.

Sosa, Ernest. 1974. "How Do You Know?" American Philosophical Quarterly 11:113–122.

- —. 1979. "Epistemic Presupposition." In George Pappas (ed.), Justification and Knowledge, 79–92. Boston: D. Reidel.
- —. 2010. Knowing Full Well. Princeton University Press.
- —. 2019. "Suspension as Spandrel." *Episteme* 16:357–368.

Srinivasan, Amia. 2015. "Normativity Without Cartesian Privilege." *Philosophical Issues* 25:273–299.

Stalnaker, Robert C. 1968. "A theory of conditionals." Studies in logical theory .

Stanley, Jason. 2005. Knowledge and Practical Interests. Oxford University Press.

Stroud, Sarah. 2006. "Epistemic Partiality in Friendship." Ethics 116:498-524.

Sutton, Jonathan. 2007. Without Justification. The MIT Press.

Turri, John. 2010. "On the Relationship Between Propositional and Doxastic Justification." *Philosophy and Phenomenological Research* 80:312–326.

Weatherson, Brian. 2011. "Defending Interest-Relative Invariantism." Logos & Episteme 2:591–609.

- —. 2017. "Interest-Relative Invariantism." In Jonathan Jenkins Ichikawa (ed.), *Routledge Handbook of Epistemic Contextualism*, Routledge Handbooks in Philosophy, chapter 19, 240–253. London: Taylor & Francis.
- Whitcomb, Dennis. 2014. "Can There Be a Knowledge-First Ethics of Belief?" In Jonathan Matheson and Rico Vits (eds.), *The Ethics of Belief: Individual and Social*, chapter 5, 89–111. Oxford University Press.
- Whiting, Daniel. 2012. "Does Belief (Only) Aim at the Truth?" *Pacific Philosophical Quarterly* 93:279–300.
- Willard-Kyle, Christopher. 2020. "Being in a Position to Know is the Norm of Assertion." *Pacific Philosophical Quarterly* 101:328–352.

Williamson, Timothy. 2000. Knowledge and its Limits. Oxford: Oxford University Press.

- —. 2005. "Knowledge, Context, and the Agent's Point of View." In Gerhard Preyer and Georg Peter (eds.), *Contextualism in Philosophy*. Oxford University Press.
- —. forthcoming. "Justifications, Excuses, and Sceptical Scenarios." In Julien Dutant and Fabian Dorsch (eds.), *The New Evil Demon*. Oxford University Press. In Press. Draft online (2 June, 2016) from http://www.philosophy.ox.ac.uk/members/philosophy_panel/tim_williamson.

Wrenn, Chase B. 2007. "Why There Are No Epistemic Duties." *Dialogue: The Canadian Philosophical Review* 46:115–136.

Yli-Vakkuri, Juhani and Hawthorne, John. forthcoming. "Being in a Position to Know." *Philosophical Studies* 1–17.

Zimmerman, Aaron Z. 2018. Belief: A Pragmatic Picture. Oxford University Press.