

Suboptimal Knowledge; or, Groundwork for a Theory of Epistemic Hygiene

Abstract

I argue for the possibility of *suboptimal knowledge*: knowledge that p that's normatively defective given that you ought to inquire into whether p . I do this by appeal to two claims from the literature on inquiry. There's the Ignorance Norm (IGN) according to which you ought not: know that p and inquire into whether p . And there's the claim that, even if you know that p , you may still be obligated to inquire into whether p (Inquiry Beyond Knowledge (IBK)). Since the IGN and IBK are consistent (as I'll explain), we do well to respect them both on account of their plausibility. Notably, though, they entail the possibility of suboptimal knowledge. The following question then arises: what must zetetic normativity be like if suboptimal knowledge is possible? To answer, I suggest that they have *epistemic health* as their anchoring value, where this is a measure of how well agents function with respect to a range of epistemic goods. The IGN exists for the sake of epistemic health maintenance (to maintain knowledge); meanwhile, IBK is true because there are reasons to acquire states of epistemic health that are stronger than mere knowledge (like certainty). Suboptimal knowledge is possible because concerns for epistemic health maintenance and acquisition have a complex interaction. Suboptimal knowledge, I submit, is the epistemic face of a phenomenon that we observe in hygienic normativity more generally.

Key terms: Inquiry; Ignorance Norm; Suboptimal Knowledge; Epistemic Health

1 Introduction

Some questions about knowledge relate to its nature (is knowledge fallible?¹). Others relate to its normative role (is knowledge the norm of assertion?²). Less attention is paid to questions about the evaluative status of knowing itself.³ There’s a reason: it’s admittedly an odd thing to say that knowledge itself could be rightly judged as good/bad, correct/incorrect, right/wrong, justified/unjustified, or whatever.

I think there’s reason to reconsider the felt oddity. As I’ll argue in this paper, there’s such a thing as *suboptimal* knowledge.

1.1 Introducing suboptimal knowledge

Generally, the claim that something is “suboptimal” means that it falls short of the standard. Here I’m concerned with suboptimality that’s traceable to the violation of a *norm of inquiry* (zetetic norm)—rendering one’s knowledge normatively defective in a sense that should be of interest to epistemologists who have taken the zetetic turn.

Provisionally, we can work with the following description:

Suboptimal knowledge: Knowledge that p that’s normatively defective given that you ought to inquire into whether p.⁴

¹Reed (2010); Reed (2012); Reed (2013); Brown (2018).

²Williamson (2000) Ch. 11; Benton (forthcoming).

³Lasonen-Aarnio (2010) is the most notable exception, though her concept of unreasonable knowledge is different than the sort of thing I’m after. More recently, Friedman (2020) comes close to arguing for the sort of claim I try to make about knowledge in this paper. Reed (2013) argues that there are different knowledge relations and that practical circumstances will dictate whether any given type of knowledge is fit for practical reliance on it. Lackey (2011) argues something similar, but with a focus on assertion. In each of these discussions, there’s the tacit assumption that one’s knowledge can be normatively defective. See Paulson (2023) for a more recent discussion of different types of knowledge. None of these authors has quite what I have in mind regarding suboptimal knowledge. But what I have in mind is congenial.

⁴Anyone who thinks that you can (i) know that p while (ii) you ought to inquire into whether p will grant that you’re violating a zetetic norm if you aren’t inquiring into whether p (assuming *ought* is zetetic). But, strictly, all the obtaining of conditions (i) and (ii) would imply is that you violate a zetetic norm by *refraining from inquiring*. That doesn’t tell us anything about your knowledge. So the satisfaction of conditions (i) and (ii) *alone* isn’t enough to vindicate suboptimal knowledge as the provocative concept I intend it to be. Something else has to enter the story. As we’ll see, that “something else” is the truth of a norm that’s familiar to discussions about inquiry: the Ignorance Norm according to which you ought not know that p and inquire into whether p at the same time.

Note that this is a description and not a *definition*. A more formal characterization of suboptimal knowledge will have to wait til §3.

A few points about the description now. First, I'll assume that inquiry is driven by interrogative attitudes.⁵ These are the kinds of attitudes one must have in order to genuinely inquire into whether p. The standard examples are *wondering* whether p and *curiosity* whether p. Second, the “ought” at play in my description of suboptimal knowledge relates an obligation of the epistemic variety. As the reader might've guessed, then, I'm going to assume that zetetic norms are epistemic.⁶ Finally, I want to suggest that it's *knowing* that we should evaluate as suboptimal and not some process or exercise of agency leading to the possession of knowledge. I might eaves drop and come to know something I really shouldn't know. In that case, there's some sense in which my knowing violates a norm. But it's more accurate to say that I shouldn't have *listened* or *paid attention* to the conversation that I overheard. My *knowing* isn't normatively problematic—at least not in a strict sense. What I mean by suboptimal knowledge invokes a stricter sense, however.⁷

1.2 Preview

Consider the following two claims:

Ignorance Norm (IGN): You ought not: know that p and inquire into whether p (at the

⁵Here I'm following Friedman (2013) and Friedman (2019b). These attitudes are question-directed with respect to a proposition. See Falbo (2022) for an argument that not all inquiring attitudes are question-directed. I set that view aside, though see footnote 44 for more on this.

⁶See, for example, Podgorski (2017) and Friedman (2020) for arguments that they're epistemic. See Feldman (2000), Goldberg (2020), and Thorstad (2022) for dissent. See Thorstad (2021) for general discussion. One might question my assumption since it's a “big” assumption. But, in fact, it only makes my job in this paper harder. If I were to assume that zetetic norms were merely practical, then much of what I say could unfold without my having to meander and defend some of the collateral commitments I'll take on (though perhaps with less interest). Besides, even if zetetic norms aren't epistemic in a strict sense, they are surely epistemic in *some* sense. We might think of them as occupying a position along a continuum of norms that's much closer to the position of strict epistemic norms than they are to, say, the norms for cooking. This makes sense since inquiry is an intellectual activity that requires both inquiring attitudes and intellectual/cognitive goals.

⁷Additionally, in saying that there's such a thing as suboptimal knowledge, one isn't necessarily committing to the idea that one should discontinue knowing. For more on this, see §3.

same time).⁸

Inquiry Beyond Knowledge (IBK): Sometimes, even though you know that p, you ought to inquire further into whether p.⁹

The IGN and IBK might seem like an odd pairing. But, strictly speaking, they are consistent. Attention to the scope of the deontic modal in each claim can help us see why: suppose you have an obligation to inquire into whether p even though you know that p (IBK); it's coherent to think of the situation as one where, if you complied with your obligation to inquire, you'd be doing something that's normatively problematic from the perspective of the IGN (knowing p and inquiring into whether p). This is a logically possible view of the zetetic normative landscape.¹⁰

Given their endorsement from divergent streams, I'll take it as a virtue of a position if it can adopt both the IGN and IBK. That's what my case for suboptimal knowledge does. I'll further discuss the IGN and IBK in §2, and then connect the dots to establish the possibility of suboptimal knowledge in §3.

Even though the IGN and IBK are consistent, as I say, they might seem like an odd pairing. A bit about why that is can help set up the preview of the second major movement of the paper.

You might think that reasons to inquire into whether p count against knowing that p.¹¹ This will strike many as an especially plausible claim about epistemic reasons to inquire into whether a proposition is true. Taken as a view about epistemic reasons, let's call the view at issue 'K-Reason.'

⁸Compare to the formulation in van Elswyk and Sapir (2021).

⁹One might consider Brown (2008), Falbo (2022), and Beddor (2023) in this connection (though see my notes about them in §2.2).

¹⁰There's a temptation to think (erroneously) that the IGN and IBK jointly imply that one could be in a normative dilemma in which one ought to inquire into whether p and ought not inquire into whether p. But sensitivity to the wide-scope formulation of the IGN should tell otherwise. I'll say more about this in §3 (see also footnote 30).

¹¹Friedman (2017), Roeber (2018), and Miracchi (2019) are committed to this so long as they affirm that knowing p implies believing p.

K-Reason: Any (epistemic) reason to inquire into whether p counts against knowing that p.

If being obligated to inquire into whether p is a matter of having reasons to inquire into whether p, then K-Reason is bad news for the possibility of suboptimal knowledge insofar as it's a problem for the epistemic reading of IBK.

Now, one could push-back against K-Reason (as I will). But matters aren't so simple. To many it will seem like a truism that the goal of inquiry is to acquire knowledge. But some authors have taken the thought a step further with the hope of explaining zetetic norms solely in terms of knowledge as inquiry's aim.¹² Consider 'K-Aim:'

K-Aim: Knowledge is the aim of inquiry such that, once one has acquired knowledge that p, one cannot have any (epistemic) reason to inquire further into whether p.

If IBK is true, it's not just K-Reason that must be resisted; K-Aim has to go too. But, without K-Aim, the IGN might seem like it lacks an account of its status as a *bona fide* zetetic norm.

My case for suboptimal knowledge will have us rejecting both K-Aim and K-Reason. But I'll try to convince you that these rejections are ultimately worth it on account of the novel view of zetetic normativity that emerges in my defense of suboptimal knowledge.

In response to K-Reason-inspired worries, I argue that not all epistemic reasons to inquire into whether p are epistemic reasons against *believing* p. So, they don't necessarily count against knowing that p. To defend my reasoning, I motivate the claim that some practical reasons to inquire are epistemic insofar as they are "right kind"¹³ reasons for interrogative attitudes. I argue all this in §4.

In §5, I respond to the objection that the IGN lacks an explanation if K-Aim is false. As an alternative, I suggest that we might conceive of the IGN as derivative upon a general

¹²We might consider Williamson (2000), Millar (2011), Kelp (2014), Kelp (2021), and van Elswyk and Sapir (2021) as representative of the position that a knowledge aim can explain the normativity of inquiry. See Friedman (2023) for a more recent broad-side dissent from this approach.

¹³In the sense that's at issue in, for example, Hieronymi (2005).

concern for knowledge *maintenance*. As we'll see, this view of the IGN doesn't require K-Aim and is compatible with epistemic obligations to inquire into whether what one already knows is true (as IBK requires).

By the end, my case for suboptimal knowledge raises this question: what must the zetetic normative landscape be like if it includes both knowledge maintenance norms and reasons to inquire for the sake of something beyond knowledge? I close in §6 by laying the groundwork for an unexplored proposal according to which the norms of inquiry are norms of *epistemic hygiene*: norms promoting *epistemic health*¹⁴ in the way that one inquires given a variety of epistemic goods besides knowledge. If I'm right, then zetetic normativity is a structure in which there are norms for epistemic health maintenance and reasons to acquire epistemic health. Their interaction can license the judgment that one's state of epistemic health is suboptimal. Suboptimal knowledge is therefore the epistemic face of a phenomenon that we observe in hygienic normativity more generally.

2 Claims from divergent streams

Here I motivate the IGN and the IBK.

2.1 The Ignorance Norm

There's debate about what the norms of inquiry are.¹⁵ A lot of that dispute surrounds whether it's permissible to *believe* that p and inquire into whether p at the same time.¹⁶ But less dispute centers around the Ignorance Norm:

IGN: You ought not: know that p and inquire into whether p (at the same time).

¹⁴Here I'm drawing on a concept introduced in Piovarchy and Siskind (2023).

¹⁵Some representative contributions include: Friedman (2019b), Friedman (2019a), Friedman (2019c), van Elswyk and Sapir (2021), Millson (2020), Falbo (2021) and Falbo (2022), and Flores and Woodard (forthcoming).

¹⁶See, for example, Friedman (2019b) and van Elswyk and Sapir (2021).

Though I’m going to be assuming the truth of the IGN, we can still give it some motivation. Perhaps the best way to see the plausibility inherent to it is by considering how expressions of the state of mind you’d have to be in would sound if you violated it (bearing in mind the constitutive link between inquiry and interrogative attitudes):

(1a) # The Warriors won the game, but I wonder whether the Warriors won the game.

(1b) # I’m curious whether the Warriors won the game, though they did win.

I’m not the first to point out that these types of utterances are problematic.¹⁷ But why are they problematic? They have a Moore-paradoxical feel about them. Assuming that proper assertion of *p* requires knowledge that *p*, per the Knowledge Norm of Assertion (KNA),¹⁸ the IGN would help us explain the data. Though it’s *possible* for a speaker, *S*, to know the contents of (1a) and (1b), assertions of (1a) and (1b) on the part of *S* would still seem problematic. The truth of the IGN would explain what is residually problematic; that’s confirmation of it.

This style of argument is not without controversy, however. One might object that, as far as linguistic data are concerned, the likes of (1a) and (1b) aren’t the whole story. Consider examples of what are called “biased interrogatives.”¹⁹

(2a) Gold is an element, isn’t it?

(2b) Isn’t gold an element?

There are settings in which uttering things like (2a) or (2b) seems perfectly natural. What’s more, anyone who utters them with sincerity would seem to have some level of confidence toward a proposition they’re questioning. More specifically, the speaker who utters these would seem to be questioning that Gold is an element while having confidence that it is. So

¹⁷See Woodard (2022) and the references therein for recent (albeit critical) discussion.

¹⁸Williamson (2000) Ch. 11.

¹⁹See especially Millson (2020). I’ve reproduced the examples from Falbo (2021).

maybe it's not so strange to know a proposition and to have an interrogative attitude toward whether that proposition is true after all.²⁰

But do the data of (2a) and (2b) actually cut against the IGN? I grant that anyone who utters (2a) or (2b) with sincerity has confidence in a proposition regarding which they have an interrogative attitude. But the felicity attached to (2a) or (2b) is plausibly accounted for by the fact that anyone who utters them either isn't making an assertion, or else is, but is doing so in a hedged manner.²¹ Thus, the question of whether the speaker has knowledge that Gold is an element doesn't even arise. Therefore, the data represented by (2a) and (2b) are irrelevant to the truth of the IGN.

That said, (2a) and (2b) *are* problematic for norms that proscribe having an interrogative attitude with respect to what one *believes* on the assumption that having a certain level of confidence in p is sufficient for believing that p. Indeed, Millson (2020), Falbo (2021), Falbo (2022), and van Elswyk and Sapir (2021) all contend that the kind of data we have in (2a) and (2b) are problematic for Friedman's norm Don't Believe and Inquire. But, as we shall see later (§4.1), I reject a prohibition on merely believing something and having interrogative attitudes regarding it.²²

2.2 Inquiry Beyond Knowledge

Switching gears, let's get K-Aim before us again:

K-Aim: Knowledge is the aim of inquiry such that, once one has acquired knowledge that p, one cannot have any (epistemic) reason to inquire further into whether p.

Some authors have argued that this makes the wrong prediction about the propriety of

²⁰A different objection to my motivation for the IGN—one I won't address given limitations of space—is that we can account for the infelicities in strictly pragmatic terms, in the style of Grice (1975).

²¹See Benton and van Elswyk (2020) and van Elswyk and Sapir (2021) for relevant discussion.

²²The resulting combination of endorsements might seem puzzling. How is it that I can endorse a prohibition on knowing p and inquiring into whether p *without* endorsing a prohibition on believing p and inquiring into whether p? My discussion of K-Reason in §4 and my account of the IGN in §5 can help sort this matter out.

inquiry in cases where one already has knowledge.²³ In some of these cases, one's knowledge is non-occurrent. Here's one from Falbo (2022):

Rush: Sahar is auditioning for a talent show. Over the past few months, she has perfected a drum solo to Rush's song *Fly By Night*. As she's filling out the audition form, she's having trouble recalling the name of the song. She knows that it starts with 'F' and that it appeared in 1975, but she can't recall anything more than that. This is frustrating because she knows the name of the song—she plays it *all* of the time—but she can't bring it to mind in the current moment. She quickly Googles it. She searches 'Rush album 1975'. As soon as she sees the owl on the album cover, it hits her: *of course—it's Fly By Night!*²⁴

Falbo then remarks that her case “illustrates that it can be rationally permissible to inquire into a question to which one already knows the answer.”²⁵

We might also consider cases where the knowledge is occurrent. Here's one adapted from Beddor (2023):

Murine Research: Mia is a scientist who forms the hypothesis *m*: A particular drug (Accuphine) causes hyperactivity in mice. Mia conducts a number of experiments that support *m*. Eventually, she conducts enough experiments to know that *m* is true. But she still is not completely certain of *m*, and rationally so: she rationally assigns at least some credence, however slight, to the possibility that all of her experimental results are attributable to confounding factors. One day Mia receives an email from a researcher at another university. Their email announces that they have just completed the most comprehensive study to date

²³To be clear, it makes the wrong prediction provided that *one is epistemically permitted* to inquire only if *one has an epistemic reason* to inquire. Whether that's right won't ultimately matter for my purposes since I'm going to suggest that it's possible to be epistemically *obligated* to inquire into whether what one already knows is true, and I assume that being epistemically obligated to inquire *does* imply having epistemic reasons to inquire.

²⁴Falbo (2022) Pg. 306.

²⁵Falbo (2022) Pg. 306.

on the effects of Accuphine on mice, with a specific focus on whether Accuphine causes hyperactivity. They provide assurance that, in their experiments, they've controlled for the relevant confounding factors. As a courtesy, they provide all their data as an attachment. Opening the attachment, Mia inquires into whether hyperactivity correlates with Accuphine in the relevant experiments.

Even here, where Mia's knowledge of m is occurrent, it seems permissible for her to inquire into whether m .²⁶

We could be even more bold than Falbo and Beddor here. We could try to dream up similar cases where one is *obligated* to inquire.²⁷

Here's one of my own invention:

Tom's Predicament: Tom knows that the capital of Finland is Helsinki, for he remembers learning this in school. Yet the capital of Finland isn't really something he's thought about since learning about global geography as a kid. If asked what the capital of Finland is in a casual setting, he'd answer with confidence that it's Helsinki. It's the kind of thing he'd get right at bar trivia. But the present situation is anything but casual. Tom gets one shot to answer correctly the question 'what's the capital of Finland?' or else an evil genius will destroy the planet. Tom has the opportunity to look the answer up online, or ask a friend, and, more generally, to double check.

Tom clearly ought to inquire into whether Helsinki is the capital of Finland before he gives his answer. Evidence for this is that we'd have grounds to criticize Tom for answering without double checking. But, if that's right, then, before he answers, his knowing that Helsinki is the capital of Finland notwithstanding, he ought to inquire into whether Helsinki is the capital.

In light of this last case, I submit that we should accept Inquiry Beyond Knowledge:

²⁶Beddor (2023) Pg.3.

²⁷One might consider the Surgeon case in Brown (2008).

IBK: Sometimes, even though you know that p , you ought to inquire further into whether p .

And, remember, I'm endorsing the epistemic read of the "ought" in IBK.

Now, there are ways to push back on what I'm saying about Tom's Predicament. For example, you might doubt that Tom has anything more than a *practical* reason to inquire further. Or you might object that the case is mis-described because Tom lacks knowledge once he has reasons to inquire further—a natural take given what's stipulated to be at stake in the case and a commitment to pragmatic encroachment in epistemology.²⁸

I'll return to both of these issues (§4.2). But here I'll just say the following. If you take Falbo and Beddor's cases as supporting that it can be *epistemically* permissible to inquire into whether p despite knowing p , then it's hard to see why someone in Tom's circumstances isn't epistemically obligated to inquire into whether p despite knowing p . In each case, we can describe the reasons favoring inquiry as focal points in an assessment of practical rationality. If that's no barrier to counting the *permissions* to inquire in Rush or Murine Research as epistemic, then it's not clear why it would be a barrier to counting the *obligation* to inquire in Tom's Predicament as epistemic.

I want to take Falbo's and Beddor's cases as seriously as they do on the assumption that zetetic normativity is epistemic. So, taking their objection to K-Aim a step further, I say that it can be epistemically permissible to inquire into whether p even if, at the same time, you already know that p . And yet another step further: it can be epistemically obligatory to inquire into whether p even if, at the same time, you already know that p .

3 Connecting the dots: suboptimal knowledge

On the one hand, it seems like you ought not: know that p and inquire into whether p . That's the IGN. On the other hand, it seems as though it's possible to know that p and yet

²⁸Fantl and McGrath (2002); Fantl and McGrath (2009); Bolinger (2020).

be obligated to inquire into whether p . That's IBK.

Let me acknowledge once more now that there at least appears to be something odd about accepting both the IGN and IBK. But, as I explained at the outset, these claims from divergent streams in the literature are consistent. So, whatever the source of the felt oddity, it isn't incoherence. Here we'll get a clearer sense of why that is as we connect the dots to establish the possibility of suboptimal knowledge.

3.1 Suboptimal knowledge: a formal profile

Consider the following type of argument:

(i) S ought not: ϕ and ψ

(ii) S ought to ψ

Therefore,

(iii) S ought not ϕ

(i) and (ii) are formally consistent. Therefore the fact that the derivation of (iii) from (i) and (ii) is valid in standard deontic logic isn't a function of a contradictory setup.²⁹ This argument-type will be instructive. So let's give it a name. Call it 'PATTERN.'

Note that, nothing about a situation in which the likes of PATTERN's premises are true requires that it's impossible that S ϕ s. This means that, without countenancing any incoherence, we could affirm, while upholding (i) and (ii), that S ϕ s and yet, at the same time, S *ought not* ϕ .

Would this have to be a situation in which S should stop ϕ -ing? How we answer this question depends on how we're understanding the type of normativity at issue. But short of saying that S should stop ϕ -ing, the likes of (i) and (ii) would at the very least justify concluding that S's ϕ -ing would be a normatively defective kind of ϕ -ing.

²⁹The argument is valid given a principle of Inheritance for obligation.

Why am I going on about PATTERN? Well, notice that our principles IGN and IBK provide us with the ingredients for an instance of it. Consider:

(iv) S ought not: know that p and inquire into whether p

(v) S ought to inquire into whether p

Therefore,

(vi) S ought not know that p

(iv) represents the IGN; and (v) will hold for any subject, S, witnessing what IBK counts as a possibility: that, even while S knows that p, S ought to inquire into whether p. All that in mind, what our reflection on PATTERN shows is that, if we accept both the IGN and IBK, we'll have to accept as a possibility that a subject both knows that p and yet ought not know that p. S's knowing that p in such a case would be a normatively defective kind of knowing.

This result, I submit, explains why the IGN and IBK are an odd pairing. Though the IGN and IBK are formally consistent, their implication will presumably strike many as odd—so odd that one might think, given logically independent considerations having to do with the nature of knowledge and the normativity of inquiry, that we should reconsider accepting both the IGN and IBK.

Over the next few sections (§4 and §5), I'm going to defend the possibility of suboptimal knowledge against these types of considerations. Meanwhile, we can appreciate the justification of my description of suboptimal knowledge as a kind of knowing that's normatively defective given that you ought to inquire. The description is justified by the strictly formal considerations adduced here. As should be clear by now, any situation in which you know that p and yet ought to inquire into whether p (a case witnessing IBK) will be a situation in which you know that p and yet know in a normatively defective way (given the IGN, you ought not know that p).³⁰

³⁰You might wonder whether the joint satisfaction of the IGN and IBK implies, not just that your knowl-

3.2 Suboptimal knowledge: some clarifications

The formal profile of suboptimal knowledge might raise more questions than it answers. For one thing, you might wonder why I've used the term 'suboptimal' to distinguish knowledge that's normatively defective in the sense at issue. More substantively, though, you might wonder whether I'm committing to the possibility of cases in which a subject should discontinue knowing and, if so, whether this means that I think any case in which a subject has suboptimal knowledge that p is a case in which a subject should refrain from believing that p (knowledge that p implies belief that p , after all). I'll focus on these more substantive questions here. Along the way my rationale for 'suboptimal' should become clear.

As the discussion in the next section will show, I deny that we should conclude, from the fact that one has suboptimal knowledge that p , that they should cease believing p . So, if "stop knowing" is ultimately code for "stop believing," then I deny that we should draw the conclusion, from the fact that one has suboptimal knowledge that p , they they should stop knowing p .

Just as I noted above in my general remarks about PATTERN and the subject who ought not ϕ , whether it's apt to say that a subject should stop knowing in a case of suboptimal knowledge depends on how we're understanding the normativity at issue. For my purposes,

edge that p would be normatively defective in the sense that you ought not know, but also that inquiring into whether p would be normatively defective in the sense that you ought not inquire. This conclusion would be hasty, however, as it would require the added premise that *you ought to know that p* . That's not part of the setup, however.

Might there be a standing zetetic obligation to know that p , for any p ? If that were true, then, indeed, it would follow that a case of suboptimal knowledge that p would also be a case in which one ought not inquire into whether p . Whether we should think this way about the zetetic normative landscape depends on how comfortable we are with normative dilemmas, as the present case would be one in which you ought to inquire into whether p and ought not inquire into whether p (and that you ought to know that p and ought not know that p). I'm not committal on this point, but I have reservations about countenancing this type of normative dilemma. My reason being that, on pain of contradiction, we'd then have to reject (in the case of zetetic normativity) one of the following:

- (a) The equivalence of *S ought to ϕ* and *it's not permissible for S to not ϕ*
- (b) That *S ought to ϕ* entails *S is permitted to ϕ*
- (c) Principles of Inheritance for obligation and permission

I leave the derivation of contradiction as an exercise. (Thanks to an anonymous referee for encouraging me to comment on this).

then, it'll matter how I'm conceiving of the normative structure of inquiry given my endorsements of the IGN, IBK, and, derivatively, the possibility of suboptimal knowledge. Allow me to suggest a framework with the help of some analogies.

Consider the norm that you ought not text while driving. The existence of this norm has a point to it. It exists in order to reduce the risk of accident and to maintain safety. Now, consistent with this norm is the possibility that one ought to send a text message even while driving. We can even fill in the details so that the obligation at issue is of the same normative type as the wide-scope prohibition on texting and driving. Just fill in the case so that sending the text is a matter of safety. It follows, I want to suggest, that the driver is engaged in suboptimal driving. The reason being that the driving is in a normative bind of sorts—there being a complex interaction between a concern for safety maintenance and a concern for safety acquisition. Intuitively, a driver in this predicament really should pull over and stop driving. Indeed, that would seem to be the thing to do given a more general concern for safety. Continued driving would be suboptimal given that concern.

This example gives a sense for how I'm thinking about the interplay between the IGN and IBK. But, even though it's instructive, I hesitate to say that everything about normativity in this example has an analogue in cases of suboptimal knowledge.

Let me be clear about what I think *is* analogous, though. In both cases, I think there's a specific kind of interaction between a wide-scope norm that exists for the sake of *maintaining* some value and a narrow scope requirement of the same type that holds true because of the importance of *acquiring* a commensurate yet distinct value. This kind of interaction leads to intuitive suboptimality judgments.

Note that we *won't* have an intuitive suboptimality verdict if the requirements aren't of the same type and promoting of a commensurate value. To see why, suppose, contrary to fact, that the only reason one shouldn't drive and text is because people tend to send very boring texts when they text while driving.³¹ In that case, any urgent need to text wouldn't

³¹Thanks to an anonymous referee for this example.

render one's driving normatively defective—at least not in any sense worth noting.³²

I would hesitate to give any unqualified endorsement to the proposition that, just as the driver in the above example should stop driving, anyone who has suboptimal knowledge that *p* should stop knowing *p*. The reason being that suboptimality, being the sort of thing that's derivative upon the kind of complex interaction between wide-scope and narrow-scope requirements that I'm interested in, can be underwritten by a variety of negative normative judgments. In the driving case, that the driving is suboptimal is underscored by the judgment that, given a general concern for safety, the driver should stop driving. But there are comparable states of suboptimality that are unaccompanied by judgments of this more specific form.

Consider the norm that you ought not run while wearing poor running shoes. This rule exists to maintain a certain kind of health: being injury free. Consistent with the truth of this norm is the possibility that, for the sake of health, you ought to jog. Even if what you're wearing are poor running shoes, it can still be true that you ought to run for the sake of cardiovascular health. Nevertheless, we can acknowledge that there's something *less than ideal* about your footwear given a more general concern for health. In that respect, your footwear it's suboptimal. That doesn't necessarily mean that you should take your shoes off, however. Wearing *some* running shoes might be better than not wearing any, as far as health is concerned.

Health, not safety, was the subject's concern in this last example. I think the phenomenon of suboptimal knowledge—and, correspondingly, the normativity of inquiry—is better understood with reference to the running shoes analogy on account of that. As I previewed earlier, I'm going to turn to a discussion of *epistemic health* in closing. To preview once more: in my view, suboptimal knowledge is born out of a zetetic normative structure in

³²If, in this alternative possible world, the norms are of the same type—being such that compliance with them serves only to promote the value of entertainment—then, yes, the driver would be engaged in suboptimal driving relative to entertainment norms (if that's a thing). This may well be bizarre. But that seems to be more a function of the fact that the justification for the norms that we're now imagining is itself bizarre.

which concerns for epistemic health maintenance and reasons to acquire epistemic health have a complex interaction. We'll have cause to refer to the norms of inquiry as norms of epistemic hygiene for that reason.

The main point at present is that it's not quite right to say that a subject who has suboptimal knowledge that p should stop knowing that p . Indeed, if the present analogy is instructive on this point, what we should conclude is that the subject should try to know *ideally* or *optimally*. A failure to know that p with certainty, for example, might render one's knowledge that p suboptimal given the norms of inquiry and from the point of view of epistemic health. Insofar as there's a failure to know optimally and yet knowledge is retained, we're dealing with suboptimal knowledge.

4 Worries related to K-Reason

Let's turn now to some objections to suboptimal knowledge as I've developed it up to this point.

It might seem as though any epistemic reason in virtue of which you ought to inquire into whether p would be a reason that necessarily counts against knowing that p . Let's get that idea before us once more:

K-Reason: Any (epistemic) reason to inquire into whether p counts against knowing that p .

The possibility of suboptimal knowledge would seem to fly in the face of this given a particular view of the relationship between obligations and normative reasons. Let's assume that S ought to ϕ , where ϕ is an attitude or action, just in case S has sufficient reason to ϕ .³³ Naturally, then, having an *epistemic* obligation to ϕ is to be understood in terms of the possession of sufficient *epistemic* reason to ϕ .³⁴ Given K-Reason, one would think that sub-

³³We can cash the latter out as S 's having her weightiest undefeated reasons support ϕ ing.

³⁴Nothing essential turns on the choice to understand epistemic obligations in terms of reasons in this way.

optimal knowledge is impossible given that an obligation to inquire into whether p involves having sufficient epistemic reason to inquire into whether p .

Here I'm going to argue that K-Reason is ultimately unmotivated (§4.1), lacking an argument that would rule out a view of epistemic reasons that's friendly to the possibility of suboptimal knowledge. Then I'll respond to the concern that the kinds of reasons I'm countenancing as epistemic cannot be distinguished from mere practical reasons (§4.2).

4.1 Are reasons to inquire necessarily reasons against belief?

K-Reason is surely not a brute fact. So we should consider a way of motivating it.

One thought is that an epistemic³⁵ reason to inquire into whether p would have to be a reason to *suspend judgment* on whether p .³⁶ If such a reason were sufficiently strong, then, presumably, having such a reason would be incompatible with having knowledge that p .

Of course, one could challenge the idea that you can't know that p and have sufficiently strong reason to suspend judgment on whether p . But I wouldn't. The claim strikes me as extremely plausible given the independently plausible assumption that knowledge that p requires *epistemic justification* to believe that p . I'm not quite sure how to make sense of having your belief count as epistemically justified while, at the same time, you ought (epistemically) to suspend judgement on whether p . That your belief that p is epistemically justified implies that you're epistemically permitted to believe p , and so it can't be that you're epistemically obligated to suspend judgment on whether p .

Alternatively, one might challenge the inference from *S has sufficient reason to inquire into whether p* to *S has sufficient reason to suspend judgment on whether p*. Is that inference defensible?

It simply offers a convenient way to flesh out K-Reason related objections to suboptimal knowledge. For classical presentations of reasons-based approaches to obligations in the practical realm, see Parfit (1986), Scanlon (2000), and Ross (2003). For more recent work congenial to it in epistemology, see Kiesewetter (2017), Lord (2022), Schroeder (2011), and Schroeder (2021).

³⁵All references to reasons are epistemic henceforth unless I note otherwise.

³⁶Friedman (2017), Roeber (2018), and Miracchi (2019). Though I should note that the authors aren't necessarily wont to distinguish the reasons as epistemic.

Whether or not it is, what the reasoning behind the objection ultimately implies is that reasons to inquire into whether p are also reasons *against believing* p . With that in mind, we may consider the following argument against the possibility of suboptimal knowledge. Since knowledge that p requires belief that p , any reason in virtue of which you ought (epistemically) to inquire into whether p is also a reason in virtue of which your knowledge that p is destroyed. Why? Because any such reason counts against believing that p .

Now, with the objection in view, I think we should doubt the claim that a reason to inquire into whether p is a reason that counts against believing p . As I'll argue, it ultimately rests on a questionable generalization.

Getting started, though, note that, at least intuitively, belief that p is normatively compatible with both interrogative attitudes toward whether p and inquiry into whether p . Otherwise we couldn't make sense of some linguistic data. Consider a case from van Elswyk and Sapir (2021):

Nikolai Mikhailovich hosted a party to celebrate the start of the new year. Anna Igorovna, who didn't attend, is later asked by a mutual acquaintance about who attended the party apart from Nikolai. Anna considers the question, and offers one of these replies:

(3a) Only Borya was at the party, I believe. But I don't know that. So I can't help but wonder myself which people attended the party.

(3b) Only Borya was at the party, I believe. But I don't know that. So let's ask Kolya, the party's host, to figure out who was at the party.³⁷

(3a) and (3b) do not seem to describe confusion or irrationality. Rather, the speaker's explicit self-attribution of belief while having an interrogative attitude or intention to inquire seem fine. But, if that's right, it seems permissible to believe p while having an interroga-

³⁷van Elswyk and Sapir (2021) Pg. 5848. I've adjusted the enumeration to fit the numbering scheme in this paper.

tive attitude toward whether p/inquiring into whether p—not exactly what you’d expect if reasons to inquire are reasons against belief.

One might still wonder if the interplay between *sufficient* reason to inquire into whether p—the kind that makes inquiry mandatory—and belief that p makes suboptimal knowledge impossible. One might concede that not all reasons to inquire into whether p necessarily count against belief that p in such a way that suspension is mandatory. Perhaps that’s all the linguistic data entitle us to conclude. Still, one might suppose, *sufficient* reason to inquire into whether p may well count decisively against belief that p.

But I think that even this more subtle position rests on a mistake. The bottom line is that it’s not clear why reasons to inquire must, of necessity, count against belief.

Perhaps there’s this rationale. I’m taking for granted a constitutive link between inquiry and interrogative attitudes. With that assumption, it’s natural to think that any reason to inquire into whether p is a reason to have an interrogative attitude toward whether p. With that mind, perhaps something like the following argument lies behind a presumed connection between reasons to inquire and reasons against belief:³⁸

- (1) A reason for or against an attitude regarding p is epistemic only if it’s a marker of the truth value of p.
- (2) A marker of the truth value of p is a reason for or against belief that p.

So,

- (3) A reason for or against an attitude regarding p is epistemic only if it’s a reason for or against belief that p.

We can think of a “marker of the truth value of p” along the lines of evidence for or against p. With that understood, I won’t contest (2). Equipped with (3), a proponent of the objection

³⁸There’s no one person to whom I’m attributing this argument. Yet I suspect it has appeal among many epistemologists—especially those of the more traditional sort who would reduce epistemic reasons and epistemic normativity to strictly evidential considerations.

to suboptimal knowledge under consideration could potentially secure an identification of reasons to have interrogative attitudes toward whether p with reasons against belief that p . After all, if, per (3), a condition on a reason for an interrogative attitude's counting as epistemic is that it counts for or against belief, presumably it counts *against*.

While this argument for (3) might seem formidable at first, premise (1) is questionable. Its chief line of support appeals to adages to the effect that epistemic reasons promote true beliefs and discourage false ones, or that they promote accurate credences. These appeals are unconvincing, however. For one thing, it's hard to distinguish their being reflective of (1) rather than, more simply, the following closely related idea:

- (1*) A reason for or against a *doxastic* attitude toward p is epistemic only if it's a marker of the truth value of p .

But, now, anyone who takes a step back and looks at each of (1) and (1*) should judge that the latter is orders of magnitude more plausible than the former.³⁹ So, given the similarity between the claims—one being a special case of the other—there's an error theory for an intuition favoring (1): we are too hastily generalizing from (1*).

That doesn't show that (1) is false, of course. Yet it does call into question its motivation—especially if there's an alternative equally plausible principle for individuating epistemic reasons that doesn't have it as a consequence. We can perhaps gesture at such a principle by noting what epistemic reasons to inquire would be reasons *for*. If there are epistemic reasons to inquire, they would be, in the first instance, reasons for *question-directed* attitudes and not *propositional* attitudes. That alone might clue us into a more nuanced way of sorting epistemic reasons than what (1) allows.⁴⁰ Epistemic reasons have to be markers of truth value *when they are epistemic reasons for or against belief*, but it's not clear why they would have to be markers of truth value when they favor non-propositional attitudes—as in the

³⁹Though it's not without resistance. See the very innovative work on epistemic reasons in Schroeder (2018) and Schroeder (2021).

⁴⁰For a nice set of discussions on what the epistemic ultimately amounts to, see Cohen (2016a), Cohen (2016b), Lyons (2016), McGrath (2016), Dotson (2019).

case of inquiry.

Here's a view of epistemic reasons you might have, then: though there are epistemic reasons to inquire into whether p , since not all such reasons count against belief that p and, necessarily, *only* epistemic reasons that count against belief that p can destroy knowledge that p , it's possible to know p and be such that you ought (in the epistemic sense) inquire into whether p . No objection so far touches this possibility.

To be clear, the idea wouldn't be that knowledge that p *can't* be undermined by epistemic reasons in virtue of which one ought (epistemically) inquire into whether p . Indeed, the explicit repudiation of knowledge in (3a) and (3b) from van Elswyk and Sapir (2021) may well suggest otherwise.

Rather, what I'm suggesting is that the deciding factor for whether or not knowledge is destroyed by reasons to inquire is whether or not the belief-worthiness of the relevant proposition is one way or another given those reasons. In a slogan: only epistemic reasons that bear on the belief-worthiness of p can destroy knowledge that p , but not all epistemic reasons to inquire into whether p bear on the belief-worthiness of p .⁴¹

4.2 The relationship between practical and epistemic reasons

Remember my motivating case for IBK?

Tom's Predicament: Tom knows that the capital of Finland is Helsinki, for he remembers learning this in school. Yet the capital of Finland isn't really something he's thought about since learning about global geography as a kid. If asked what the capital of Finland is in a casual setting, he'd answer with confidence that it's Helsinki. It's the kind of thing he'd get right at bar trivia. But the present situation is anything but casual. Tom gets one shot to answer correctly the question 'what's the capital of

⁴¹I take comfort in the fact that much of the literature these days is moving in a more expansive direction on the concept of an epistemic reason. See, for examples, Schroeder (2011) and Schroeder (2021). Regarding this trend with respect to reasons to inquire, see especially Fleisher (2022) and Fleisher (2023).

Finland?’ or else an evil genius will destroy the planet. Tom has the opportunity to look the answer up online, or ask a friend, and, more generally, to double check.

I said that Tom is epistemically obligated to inquire. So, if we take the details of Tom’s Predicament seriously, Tom’s knowledge that Helsinki is the capital of Finland is suboptimal.

Here’s a K-Reason-inspired worry about this analysis. Tom surely has strong *practical* reasons to inquire. But if Tom has sufficient epistemic reason to inquire too, then there’s no clear distinction between an epistemic and a practical reason to inquire. Meanwhile, if K-Reason is true, then we can conclude that Tom’s reason to inquire is merely practical. For, if he knows that Helsinki is the capital of Finland but has sufficient reason to inquire nonetheless, K-Reason predicts that those reasons aren’t epistemic.

The worry is that my analysis cannot tell the difference between epistemic and practical reasons to inquire. This recalls my comment about a related point earlier. I said that, if we take seriously that Falbo’s and Beddor’s cases (Rush and Murine Research) relate an epistemic permission to inquire, then it stands to reason that my case (Tom’s Predicament) relates an epistemic obligation to inquire. Isn’t that point adequate to meet the concern here?

Not necessarily. Perhaps not even Falbo or Beddor would’ve signed on to an epistemic reading of the permissions to inquire in their cases. And, so, they wouldn’t have countenanced epistemic reasons to inquire as I am. If that’s right, then a mere dialectical maneuver on my part won’t assuage the worry that I’ve collapsed the distinction between practical and epistemic reasons.

To make matters worse, one might object that, even if it seems like Tom’s reason to inquire is epistemic because he ought to acquire some epistemic state “stronger” than knowledge, there’s an error theory for that seeming. We should distinguish between *epistemic norms*, on the one hand, and *norms with epistemic content* on the other.⁴² To illustrate the difference, suppose you promise someone that you’ll inquire into whether p. If you don’t inquire, you’ll

⁴²For the importance of this distinction, see Simion (2016), Simion (2018), and Simion (2021).

have broken your promise. In that case, your failure to inquire would have violated a *practical* norm with epistemic content; not an epistemic norm. One might wonder, then, why Tom's obligation to inquire is epistemic rather than a practical obligation with epistemic content.

In response to these concerns, I'll start with a concession. No doubt, Tom has strong practical reasons to inquire into whether Helsinki is the capital of Finland. What I want to resist, however, is the suggestion that this undermines the case for his having suboptimal knowledge. So, I deny that this undermines the idea that he has sufficient epistemic reason to inquire into whether that Helsinki is the Capital of Finland.

Here's why. If we're thinking carefully about what Tom's practical reasons favor, we'll conclude that, in the first instance, they favor *improving his strength of epistemic position with respect to the proposition that Helsinki is the capital of Finland*. I want to suggest that this very fact ultimately furnishes Tom with an obligation to inquire that issues in an *epistemic* obligation to *have an interrogative attitude*.

To make the argument, note that, on the assumption that zetetic norms are epistemic, it'll be natural to conclude that an interrogative attitude is a kind of attitude for which one could have epistemic reasons. Indeed, this was presupposed in the previous subsection. With that in mind, we can take cues for how to individuate epistemic reasons for interrogative attitudes by reflecting on how we individuate epistemic reasons for *doxastic* attitudes. It's a familiar fact that any reason of the "right kind"⁴³ for a doxastic attitude is an epistemic reason. So, in the present setting, it's a natural commitment that any reason of the "right kind" to have an interrogative attitude is epistemic.

Now, I think it's fairly obvious that the fact that Tom ought to inquire into whether Helsinki is the capital of Finland means that he has a "right kind" reason to have an interrogative attitude toward whether Helsinki is the capital of Finland.⁴⁴ If Tom has sufficient

⁴³Hieronimi (2005).

⁴⁴See Matthews (Forthcoming) for an argument for a claim of this sort. You might wonder if it's more apt to say that Tom would have a reason, not for an interrogative attitude narrowly construed as wondering or being curious, but for something else. On this view of the case, Tom need not wonder or be curious whether Helsinki is the capital of Finland. Perhaps, instead, he ought merely *consider* the question as he (perhaps grudgingly) double-checks.

reason to inquire (even just a practical one), then he necessarily has a “right kind” reason to have an interrogative attitude given the constitutive link between inquiry and interrogative attitudes. But, if that’s right, then Tom has sufficient epistemic reason to have an interrogative attitude toward whether Helsinki is the capital of Finland. So, Tom ought (epistemically) to have an interrogative attitude toward whether Helsinki is the capital of Finland. We still get the result that Tom has suboptimal knowledge.

One might worry that this argument casts far too wide a net for epistemic reasons for interrogative attitudes, implying that, anytime one has a practical obligation to inquire, one has an epistemic obligation to have an interrogative attitude.⁴⁵ But that doesn’t necessarily follow. If given more space to develop the view of reasons on offer here, we might be able to impose sensible constraints. As I indicated at the beginning of the present line of reasoning, it’s relevant that, fundamentally, Tom’s practical reasons support improving his strength of epistemic position. His obligation to inquire is derivative upon that. It’s not as though he merely promised that he’d inquire. The fact that the buck stops with a need for epistemic improvement might be crucial, then, for separating those practical obligations that furnish epistemic reasons from those that don’t. We can operate on the following assumption, then: you have an epistemic reason to have an interrogative attitude toward whether *p* when you have a practical reason to inquire into whether *p* for the sake of epistemic improvement with respect to *p*.

Taking stock, my contention is that, in Tom’s Answer, even though Tom knows that Helsinki is the capital of Finland, he ought to inquire into whether Helsinki is the capital of Finland in order to improve his strength of epistemic position with respect to that propo-

Even so, that wouldn’t impact my argument in this section or any other. In this section, I’m concerned to vindicate reasons to inquire as epistemic by means of an argument regarding “right kind” reasons for interrogative attitudes. How we ultimately construe those attitudes is beside the point. And if I conceded that Tom ought merely *consider* whether Helsinki is the capital of Finland (though I have my doubts that it isn’t just as intuitive to suppose he should wonder whether it is), that concession wouldn’t undermine my case for the IGN. Additional linguistic motivations for it would apply. ‘I’m considering again whether Helsinki is the capital of Finland, but it is the capital’ sounds odd. (Thanks to an anonymous referee for encouraging to clarify all this.)

⁴⁵Then again, some proponents of moral encroachment might find the idea congenial. See Bolinger (2020) for discussion.

sition. The fact that he ought to inquire, given the reasoning laid out above, implies that he ought (epistemically) to have an interrogative attitude toward whether Helsinki is the capital of Finland. Thus, Tom knows that Helsinki is the capital of Finland but also ought (epistemically) to have an interrogative attitude toward whether Helsinki is the capital of Finland. Given the truth of the IGN, then, Tom’s knowledge that Helsinki is the capital of Finland is suboptimal.

Closing this section out, let me now make good on my promise to weigh in on whether Tom’s Predicament is better analyzed through the lens of pragmatic encroachment. On this take, because of the high stakes attached to Tom’s getting the answer wrong, he doesn’t actually know that Helsinki is the capital of Finland and the case is mis-described.

Now, in a way, the appeal of the pragmatic encroacher’s analysis functions as some evidence for my position. For it requires that there’s something lacking knowledge-wise in Tom’s belief that Helsinki is the capital. The difference is that I claim that Tom knows suboptimally while the pragmatic encroacher would say that Tom doesn’t know at all. How to adjudicate?

Though arguments for pragmatic encroachment are conceptually rich and sophisticated, many of them boil down to intuitions about cases just like Tom’s Predicament. But note the following: if I’m right that Tom has suboptimal knowledge, that gives us an error theory for any intuition favoring a pragmatic encroachment-style diagnosis. For it’s possible that we’re confusing an intuition favoring that Tom doesn’t know *optimally* with the judgment that Tom doesn’t know *at all*. Indeed, it’s a familiar fact that we confuse the judgment that ‘F is not *an ideal* or *optimal* G’ with the judgment that ‘F is not G.’⁴⁶

My appeal here is somewhat partisan. If you’re not partial to pragmatic encroachment to start, then I submit that my position represents an attractive way to reconcile a purist invariantism about knowledge with what animates pragmatic encroachment in the first place:

⁴⁶Further development on this point would contribute yet another reason to think that clarity on the epistemology of inquiry undermines the motivation for pragmatic encroachment (see, for example, Goldberg (2020)). Though my thought would be more in line with the position in Reed (2013).

the idea that what's at stake should be able to count as relevant to one's state of knowledge. So long as the stakes issue in practical reasons to inquire, we can respect the epistemic significance of stakes via the possibility of suboptimal knowledge.

5 A worry related to K-Aim

We're not out of the woods yet. I have officially rejected K-Reason. But remember that I also have to reject K-Aim:

K-Aim: Knowledge is the aim of inquiry such that, once one has acquired knowledge that p , one cannot have any (epistemic) reason to inquire further into whether p .

In fact, my rejection of K-Aim is overdetermined by now.⁴⁷

I'm not in bad company denying K-Aim, as we saw.⁴⁸ But it's not enough to simply deny K-Aim in upholding the possibility of suboptimal knowledge. For if K-Aim is false, then the place of the IGN in the theory of inquiry is somewhat mysterious. Remember, the IGN says that you ought not know that p and inquire into whether p . Though this doesn't imply K-Aim, the truth of K-Aim looks like a natural explanation of the truth of the IGN.⁴⁹

Presumably, the IGN isn't a brute normative fact. But, then, if suboptimal knowledge is possible, how do we explain the truth of the IGN without K-Aim? It's incumbent upon me to state an alternative explanation. That doesn't require an additional article-length defense of an explanation. But it does require something that's not implausible on its face. And whatever alternative explanation I offer, it'll have to be compatible with epistemic reasons to have interrogative attitudes toward whether a proposition one knows is true. Otherwise, I'll have provided an account of the IGN's place in the theory of inquiry that undermines my defense of suboptimal knowledge up til now.

⁴⁷K-Aim puts a constraint on when inquiry is obligatory: it's obligatory only if one hasn't met the aim of inquiry into whether p ; namely, knowledge that p . If you know that p , then, given K-Aim, you can't have reason to inquire into whether p that makes inquiry obligatory. My rejection of K-Reason therefore denies K-Aim by implication.

⁴⁸Again: Millson (2020), Falbo (2022), Friedman (2023), Woodard (forthcoming), and Beddor (2023).

⁴⁹See van Elswyk and Sapir (2021) for development of this explanation.

So, I need to do two things to meet the explanatory challenge at hand: state an alternative to K-Aim for explaining the truth of the IGN that isn't implausible on its face, and explain why that explanation doesn't undermine my defense of suboptimal knowledge. Toward that end, I propose that the IGN is derivative upon a concern for knowledge *maintenance* rather than knowledge as what one tries to acquire in the effort to reach an aim of inquiry. On this view, the IGN is born out of a concern for maintaining knowledge already possessed. The aim of inquiry is orthogonal.

Here's a bit more to the thought. When we inquire into questions about what we know, we run the risk of losing our knowledge—making ourselves more susceptible to a change of mind as a result of inquiry. Even if you're inquiring to confirm what you already know, it's risky business: you might just lose your knowledge on account of, say, acquiring misleading evidence. Losing knowledge is undesirable. So we'll want guidance in the form of rules for maintaining knowledge. And it seems like there are such rules. That there are knowledge maintenance norms is evident from the fact that you shouldn't intentionally self-inflict knowledge loss. Connecting the dots, then, my alternative regarding the grounding of the IGN is that, because we are interested in knowledge maintenance and not just knowledge acquisition, our inquiries are constrained by norms compliance with which reduce the risk of knowledge loss.

How does this square any better with the possibility of suboptimal knowledge? Some maintenance norms are wide-scope rules that issue in normative statuses like that of suboptimal knowledge. I earlier gave the example that one shouldn't drive and text at the same time. This norm exists so that a driver maintains focus while driving to reduce the risk of accident. The norm is entrenched in our evaluations of drivers to the point where, even in a case where a driver has to send a text message because the message is urgent and they can't call, we still treat it as applicable and are inclined to judge that the driver's driving is less than ideal.

I bring this example up again because it shows that the truth of the norm expressed by

‘don’t drive and text’ tolerates the possibility of reasons to text *even while driving*. That is, though it’s true that you shouldn’t drive and text, it’s also true that, even if you’re driving, you might have sufficient reason to text. This is a proof of concept of the normative structure required for suboptimal knowledge. Even if a concern for knowledge maintenance dictates that you not know *p* and inquire into whether *p* at the same time, a concern for the acquisition of some other epistemic good beyond knowledge can make it so that you ought to inquire into whether *p*.

The bottom line: so long as there are epistemic goods beyond knowledge for the sake of which one ought inquire, one can have an epistemic reason to inquire into whether *p* despite knowing that *p*; nevertheless, the IGN has its place in the zetetic normative landscape because we care about the maintenance of epistemic goods and not just their acquisition.

6 Motivating a theory of epistemic hygiene

Having just now proposed that the IGN is a maintenance norm, what we have so far in my case for suboptimal knowledge represents an under-explored take on zetetic normativity. On the one hand, there are norms of knowledge maintenance (the IGN). On the other hand, there are reasons to inquire for the sake of epistemic goods other than knowledge (IBK).

You might be wondering ‘what must zetetic normativity be like more fundamentally if it features both of these?’ This is too big a question to cover adequately in the space that remains. But I think I’ve said enough to motivate a promising and unexplored position.

To lay out the view I have in mind, let’s consider a slightly more specific question: what value or values would unite zetetic normativity in light of the threads developed in this article? What we say in answer to this has to accommodate both knowledge maintenance norms and zetetic normative truths to the effect that one must inquire for the sake of something other than knowledge.

It’s in this connection that I want to draw attention to a recent discussion in Piovarchy

and Siskind (2023) on the concept of *epistemic health*. Equipped with it, we can cite some of the issues surrounding suboptimal knowledge in this paper to motivate the view that the norms of inquiry are norms of *epistemic hygiene*. “Epistemic health,” according to the authors, “is a measure of how well entities (e.g. persons, communities, institutions) are functioning with regard to a range of epistemic goods or ideals.” This evaluative concept nicely compliments the zetetic normative structure required for suboptimal knowledge—and for two reasons.

First, notice that Piovarchy’s and Siskind’s definition of epistemic health invokes a variety of epistemic goods. This is fitting if there are epistemic reasons to inquire, not just for the sake of knowledge, but for the sake of states that are stronger than or supplemental to knowledge. As a measure of how well one is functioning with regard to a variety of epistemic goods, then, epistemic health affords us a suitable unifying value.

But epistemic health isn’t merely evaluative. Hence the second reason epistemic health nicely compliments the zetetic normative structure required for suboptimal knowledge: epistemic health affords us a kind of normative assessment. Consider what Piovarchy and Siskind say about the concept of health in general:

Our ideas of health and uses of the associated concept are often normative. Being healthy is often considered desirable or valuable, motivating us to promote particular ends, and health plays a valuable role helping us co-ordinate to produce those ends. We encourage others to be concerned with their health, we develop medical systems to restore health, and we have public health campaigns aimed at preventing ill health... [O]ur applications of the concept are often implicitly prescriptive. For example, when we say smoking is unhealthy, this is not purely descriptive, like when we say that cigarettes contain plant matter. We mean that smoking is *bad*, and one has reasons to abstain from it.⁵⁰

The authors then point to a variety of cases in which health-related concepts have been said

⁵⁰Piovarchy and Siskind (2023) Pg. 2332.

to play this normative role in the social epistemology literature.⁵¹ Of particular interest for us is the relationship between epistemic health and knowledge. Piovarchy and Siskand document literature with claims to the effect that it’s epistemically unhealthy to threaten one’s knowledge. As an example, they cite Williamson’s description of skepticism as a “disease” of our “cognitive immunity system.”⁵²

I think Piovarchy and Siskand are on to something. And, because of that, I think it’s fitting that we consider epistemic hygiene as a complimentary notion.

Here’s the idea. Just as there are norms of physical hygiene designed to promote physical health through the encouragement or discouragement of physical activities, there are norms of epistemic hygiene designed to promote epistemic health through the encouragement or discouragement of epistemic activities—inquiry being one of the more salient examples of such an activity.⁵³ Our discussion of suboptimal knowledge motivates this.

Here’s the rationale. In general, hygienic normativity features reasons to acquire health and norms for maintaining health. The interaction between reasons for health acquisition, on the one hand, and norms for health maintenance, on the other, have normative implications that, while banal in some contexts, are potentially illuminating in light of our discussion of suboptimal knowledge. I gave the running shoes example in §3. I think it illustrates the point. But to avoid repetition, consider that being operated on involves risks to one’s physical health—risks that ground norms proscribing doing the sorts of things characteristic of open-heart surgery (cutting yourself open isn’t exactly physically healthy!). But, for the

⁵¹Joshi (2022), Begby (2021), and Aird (2023), for example.

⁵²Williamson (2005). But one needn’t be partisan in debates about skepticism to appreciate a general concern for health in the way that we relate to knowledge. Here are Piovarchy and Siskind once more (Piovarchy and Siskind (2023) Pg. 2334):

[W]hat kinds of states we think are valuable can influence what we think of as epistemically healthy, independent of whether those states are in fact epistemically healthy or not. For example, given you think that endorsing Cartesian scepticism leads to missing out on all kinds of epistemic goods, you might think that someone who lives according to Cartesian scepticism is epistemically unhealthy, while they might think the epistemic risks you are taking in admitting beliefs you are not absolutely certain of is unhealthy.

⁵³Though I don’t claim that this is the only such activity.

sake of physical health, the reasons for having surgery might be strong enough so that one ought to have those things done. Yet no one will deny that the person undergoing open-heart surgery is in a suboptimal state of physical health.⁵⁴

A fuller theory of epistemic hygiene promises to provide us with an even deeper understanding of suboptimal knowledge. It would be a theory of inquiry according to which the norms of inquiry promote epistemic health in the way that one inquires given a variety of epistemic goods. What's more, zetetic norms would be grounded in concerns for both acquisition and maintenance of epistemic goods constitutive of epistemic health.

7 Conclusion

In this paper, I've argued that there's such a thing as suboptimal knowledge: knowledge that p that's normatively defective given that you ought to inquire into whether p . As we saw, I made the case for this by appeal to two claims that, at least initially, appear to be in some tension. There was the Ignorance Norm (IGN) according to which you ought not: know that p and inquire into whether p . There was also the claim that, at least sometimes, you ought to inquire into whether p despite the fact that you know p (IBK). Having diffused objections based on both K-Reason and K-Aim, we saw that there's a consistent and defensible way to uphold both the IGN and IBK—a way that offers insight to discussions about the relationship between practical and epistemic reasons, pragmatic encroachment, and the explanation of the IGN.

I've also suggested that, in order to makes sense of the normative profile of suboptimal knowledge, we might do well to think of it as comparable to what we see in hygienic normativity more broadly. With epistemic health as a unifying value—one that implies the existence of a variety of epistemic goods—we should expect norms for epistemic health maintenance as well as reasons to acquire all kinds of epistemic goods constitutive of epistemic health (not

⁵⁴It's worth pointing out that there's something of a precedent set for my use of the term “suboptimal” as it pertains to health. See, for example, Wang and Yan (2012).

just knowledge). If our reasons to acquire these goods are strong enough, it may be that we ought to inquire into whether something is true despite already having knowledge that it is. Hence IBK. But, as I say, we should expect some epistemic health-promoting norms to be norms for epistemic health maintenance. Hence the IGN. The interactions between the zetetic normative facts thus construed are perhaps what give rise to suboptimal knowledge.

References

- R. Aird. A puzzle of epistemic paternalism. *Philosophical Psychology*, 36(5):1011–1029, July 2023.
- B. Beddor. Inquiry beyond knowledge. *Philosophy and Phenomenological Research*, page phpr.13042, Dec. 2023.
- E. Begby. Evidential Preemption. *Philosophy and Phenomenological Research*, 102(3):515–530, May 2021. ISSN 0031-8205, 1933-1592. doi: 10.1111/phpr.12654.
- M. Benton. Knowledge is the Norm of Assertion. In *Contemporary Debates in Epistemology*. Wiley-Blackwell, 3 edition.
- M. Benton and van Elswyk. Hedged Assertion. In *The Oxford Handbook of Assertion*. Oxford University Press, Mar. 2020.
- R. J. Bolinger. Varieties of Moral Encroachment. *Philosophical Perspectives*, 34(1):5–26, 2020.
- J. Brown. Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning. *Nous*, 42(2):167–189, June 2008.
- J. A. Brown. *Fallibilism: Evidence and Knowledge*. Oxford University Press, 2018.
- S. Cohen. Replies to my commentators. *Inquiry*, 59(7-8):906–922, Nov. 2016a.
- S. Cohen. Theorizing about the epistemic. *Inquiry*, 59(7-8):839–857, Nov. 2016b.
- K. Dotson. Tales from an apostate. *Philosophical Issues*, 29(1):69–83, Oct. 2019.
- A. Falbo. Inquiry and Confirmation. *Analysis*, 81(4):622–631, 2021.
- A. Falbo. Inquiring Minds Want to Improve. *Australasian Journal of Philosophy*, 2022.

- J. Fantl and M. McGrath. Evidence, Pragmatics, and Justification. *The Philosophical Review*, 111(1):67–94, 2002.
- J. Fantl and M. McGrath. *Knowledge in an Uncertain World*. Oxford University Press, Oxford, 2009.
- R. Feldman. The Ethics of Belief. *Philosophy and Phenomenological Research*, 60(3):667–695, 2000.
- W. Fleisher. Pursuit and Inquisitive Reasons. *Studies in History and Philosophy of Science Part A*, 94(C):17–30, 2022.
- W. Fleisher. Intellectual Courage and Inquisitive Reasons. *Philosophical Studies*, 180(4):1343–1371, 2023.
- C. Flores and E. Woodard. Epistemic Norms on Evidence Gathering. *Philosophical Studies*.
- J. Friedman. Question-Directed Attitudes. *Philosophical Perspectives*, 27(1):145–174, Dec. 2013.
- J. Friedman. Why Suspend Judging? *Noûs*, 51(2):302–326, 2017.
- J. Friedman. Checking again. *Philosophical Issues*, 29(1):84–96, Oct. 2019a.
- J. Friedman. Inquiry and Belief. *Noûs*, 53(2):296–315, June 2019b.
- J. Friedman. Teleological epistemology. *Philosophical Studies*, 176(3):673–691, Mar. 2019c. ISSN 1573-0883. doi: 10.1007/s11098-018-1033-7.
- J. Friedman. The Epistemic and the Zetetic. *Philosophical Review*, 129:501–536, 2020.
- J. Friedman. The Aim of Inquiry? *Philosophy and Phenomenological Research*, 2023.
- S. C. Goldberg. On the epistemic significance of practical reasons to inquire. *Synthese*, 2020.
- H. P. Grice. Logic and Conversation. Dec. 1975.

- P. Hieronymi. The Wrong Kind of Reason. *The Journal of Philosophy*, 102(9):437–457, 2005.
- H. Joshi. The Epistemic Significance of Social Pressure. *Canadian Journal of Philosophy*, 52(4):396–410, May 2022.
- C. Kelp. Two for the Knowledge Goal of Inquiry. *American Philosophical Quarterly*, 51(3):227–232, 2014.
- C. Kelp. *Inquiry, Knowledge, and Understanding*. Oxford University Press, Oxford, New York, June 2021.
- B. Kiesewetter. *The Normativity of Rationality*. Oxford University Press, Oxford, New York, Nov. 2017.
- J. Lackey. Assertion and Isolated Second-Hand Knowledge. In *Assertion: New Philosophical Essays*. Oxford University Press, Jan. 2011.
- M. Lasonen-Aarnio. Unreasonable Knowledge. *Philosophical Perspectives*, 24(1):1–21, Dec. 2010.
- E. Lord. *The Importance of Being Rational*. Oxford University Press, Oxford, New York, Nov. 2022.
- J. C. Lyons. What we talk about when we talk about epistemic justification. *Inquiry*, 59(7-8):867–888, Nov. 2016.
- G. Matthews. Fit-Related Reasons to Inquire. *Australasian Journal of Philosophy*.
- M. McGrath. Cohen on ‘Epistemic’. *Inquiry*, 59(7-8):889–905, Nov. 2016.
- A. Millar. Why Knowledge Matters. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 85:63–81, 2011.
- J. A. Millson. Seeking confirmation: A puzzle for norms of inquiry, 2020.

- L. Miracchi. When Evidence Isn't Enough: Suspension, Evidentialism, and Knowledge-first Virtue Epistemology. *Episteme*, 16(4):413–437, Dec. 2019.
- D. Parfit. *Reasons and Persons*. Oxford University Press, Oxford, New York, Feb. 1986.
- S. Paulson. First-Class and Coach-Class Knowledge. *Episteme*, 20(3):736–756, Sept. 2023. ISSN 1742-3600, 1750-0117. doi: 10.1017/epi.2023.5.
- A. Piovarchy and S. Siskind. Epistemic Health, Epistemic Immunity and Epistemic Inoculation. *Philosophical Studies*, 180(8):2329–2354, Aug. 2023. ISSN 1573-0883.
- A. Podgorski. Rational Delay. *Philosopher's Imprint*, 17(5), Jan. 2017.
- B. Reed. A Defense of Stable Invariantism. *Noûs*, 44(2):224–244, 2010.
- B. Reed. Fallibilism. *Philosophy Compass*, 7(9):585–596, 2012.
- B. Reed. Fallibilism, Epistemic Possibility, and Epistemic Agency. *Philosophical Issues*, 23(1):40–69, Oct. 2013. doi: 10.1111/phils.12003.
- B. Roeber. How to Argue for Pragmatic Encroachment. *Synthese*, (6):2649–2664, 2018.
- W. Ross, editor. *The Right and the Good*. British Moral Philosophers. Oxford University Press, Oxford, New York, new edition, new edition edition, Jan. 2003.
- T. M. Scanlon. *What We Owe to Each Other*. Belknap Press, Cambridge, MA, Nov. 2000.
- M. Schroeder. Holism, Weight, and Undercutting. *Noûs*, 45(2):328–344, 2011.
- M. Schroeder. When Beliefs Wrong. *Philosophical Topics*, 46(1):115–127, 2018.
- M. Schroeder. *Reasons First*. Oxford University Press, Oxford, New York, Aug. 2021.
- M. Simion. Assertion: knowledge is enough. *Synthese*, 193(10):3041–3056, Oct. 2016.
- M. Simion. No epistemic norm for action. *American Philosophical Quarterly*, 55(3):231–238, July 2018.

- M. Simion. *Shifty Speech and Independent Thought: Epistemic Normativity in Context*. Oxford University Press, 2021.
- D. Thorstad. Inquiry and the epistemic. *Philosophical Studies*, 178(9):2913–2928, Sept. 2021. ISSN 1573-0883.
- D. Thorstad. There are no epistemic norms of inquiry. *Synthese*, 200(5):410, Sept. 2022.
- P. van Elswyk and Y. Sapir. Hedging and the ignorance norm on inquiry. *Synthese*, 199(3): 5837–5859, Dec. 2021.
- W. Wang and Y. Yan. Suboptimal health: a new health dimension for translational medicine. *Clinical and Translational Medicine*, 1:28, Nov. 2012. ISSN 2001-1326.
- T. Williamson. *Knowledge and Its Limits*. Oxford University Press, Oxford, 1st edition edition, 2000.
- T. Williamson. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, 55(219):213–235, 2005.
- E. Woodard. Why Double-Check? *Episteme*, pages 1–24.
- E. Woodard. The Ignorance Norm and Paradoxical Assertions. *Philosophical Topics*, 49(2): 321–332, 2022.