*Meeting of the Aristotelian Society held at Senate House,
University of London, on 30 May 2022 at 6:00 p.m.*

XII—Deferring to Doubt

Miriam Schoenfield

In this paper I'll suggest that a certain challenge facing defeatist views about higher-order evidence cannot be met, namely, motivating principles that recommend abandoning belief in cases of higher order defeat, but do not recommend global scepticism. I'll propose that, ultimately, the question of whether to abandon belief in response to the realization that our belief can't be recovered from what I'll call 'a perspective of doubt' can't be answered through rational deliberation aimed at truth or accuracy.

> The sceptic with his whole nature adopts the doubting attitude; but which of us is the wiser, Omniscience only knows.
>
> —William James (1896)

I

Introduction. Sometimes we doubt a belief because we receive evidence that it was formed in a dubious manner.[1] Sometimes we doubt because we encounter disagreement. Sometimes we doubt in response to sceptical arguments. Sometimes we doubt because the possibility of error becomes salient. Sometimes we doubt for no apparent reason at all.

It's natural to think that we should abandon belief in some of these cases but not others. If I learn that I formed my belief when my cognitive faculties were not operating optimally, that may be a good reason to abandon it. But merely being reminded of the fact that I *could* be wrong isn't a good reason to revise my opinion. The aim of this paper is to argue against this natural thought, in a sense to be made more precise later. Very roughly, I'll argue that, from the perspective of a deliberator aiming at truth or accuracy, principles

---

[1] For the purposes of this paper, beliefs can be understood as attitudes of sufficiently high confidence, though I don't think anything essential will rest on this.

telling us to reduce confidence in response to some forms of doubt but not others will look unmotivated. I take these considerations to support a radical form of permissivism about higher-order defeat and scepticism (one can abandon belief in response to higher-order evidence or not; one can be a sceptic or not), though as we'll see, others may draw different conclusions.

<div align="center">II</div>

Doubt.[2] The aim of this section is to clarify what I mean by 'doubt'. I'll begin by illustrating the phenomenon I have in mind with the following example:

> STOVE: I'm walking to work one morning, listening to a podcast. I hear a fictional story about a house that burned down because someone left the stove on. I start worrying that I forgot to turn off the stove. I pause for a moment and think, 'I remember cleaning the stove right before I left. If the stove were on, I would have noticed and turned it off. So the stove must be off'. I maintain my belief and move on with my day.

This is a story in which I subject a belief (that the stove is off) to doubt. What we're interested in when we subject a belief to doubt is whether we can reason our way back to the belief from what I'll call 'a perspective of doubt'—a perspective that is in some sense less committal than our usual one. In this case, the belief was recoverable: I recovered my belief that the stove was off by appealing to my belief that I cleaned the stove.

Precisely which commitments are set aside when I subject my belief that P to doubt? It depends. When I subjected my belief that the stove was off to doubt, I was not of course willing to rely on that very belief. But in this particular case, there are several other beliefs I was not willing to rely on as well. For example, I was not willing to rely on the belief that I moved the stove knob in a certain direction, that the stove is off or $2 + 2 = 5$, or that there are no open flames in my kitchen. But there also plenty of beliefs that were not set aside: my belief that I own a stove, that I cleaned the stove, and that I came into existence more than five minutes ago. I could have doubted my belief in a more global way. If I'd set a lot more aside, I would not have been able to recover the belief from the perspective of doubt.

---

[2] The discussion in this section is an elaboration of some ideas in Schoenfield (forthcoming).

So, given the way I'm thinking about doubt, there is no universal characterization of what we set aside when we doubt a belief. There are indefinitely many ways to doubt, corresponding to indefinitely many perspectives of doubt, and whether we can recover the belief we've subjected to doubt depends on which perspective of doubt we're considering.

But what exactly is a perspective? A perspective, in the sense relevant to this paper, is just a set of truth- or accuracy-aimed doxastic commitments: these commitments can include beliefs, attitudes of agnosticism, credences, and rules which impose constraints on cognitive transitions. For example, the perspective I currently occupy includes a belief in the existence of California and a 0.5 credence that a fair coin lands heads. It requires that I move from a visual perception as of *P* to a belief that *P*, and it forbids transitions that commit the gambler's fallacy. What do I mean by a 'truth-aimed' commitment? We can leave the notion relatively vague, but at a minimum it implies that a perspective will forbid cognitive transitions that it regards as having low expected accuracy or are conducive to forming false beliefs, and it will permit transitions that it regards as having high expected accuracy or are conducive to forming true beliefs, or that it regards as resulting in no loss of truth or accuracy. When a perspective permits a series of transitions that form a path to a particular attitude, I'll say that the perspective 'permits' the attitude in question. When the perspective forbids all but one attitude towards *P*, I'll say that the perspective 'recommends' that attitude towards *P*.[3]

Note that it is not only beliefs that we subject to doubt. Inferences, transitions, or reliance on certain capacities can be doubted as

---

[3] A few more notes about perspectives. First, I'm using 'permitted' synonymously with 'not forbidden'. Second, unless stated otherwise, I'm assuming that the perspectives under discussion are coherent in the sense that two incompatible attitudes are never required by a given perspective. However, it's worth noting that when a belief can be recovered from doubt, that is often (but not always) because the perspective of doubt is incoherent. After all, in virtue of being a perspective of *doubt* concerning *P*, the perspective has some sort of agnostic attitude towards *P*. But if the *belief* that *P* can be recovered from that perspective, its commitments must either require or permit belief. If its commitments require belief in *P*, then when we 'recover' belief from that perspective we're in effect coming to recognize that the perspective of doubt in question is not actually available to us. Finally, these brief remarks are by no means intended to be a full account of the notion of a perspective. (Indeed, I'm using the word 'perspective' stipulatively to refer just to a set of doxastic commitments with the features above.) See Camp (2019) for an extended discussion of perspectives more generally. In so far as what I say here diverges from Camp's views about perspectives, the disagreement is terminological.

well. If I'm doubting inductive inferences, I'm wondering whether I can defend the use of induction without relying on induction. If I'm doubting my perceptual capacities, I'm wondering whether I can defend my reliance on perception in a way that doesn't rely on beliefs I've formed perceptually. So far, I'm not making any claims about whether this activity of doubting is rational; I'm just observing that we do it sometimes.

One final comment before proceeding: my use of the phrases 'doubting' and 'subjecting to doubt' is partially stipulative, and so may diverge somewhat from ordinary usage. We might ordinarily say things like 'I thought there would be a picnic today, but now I doubt that it will happen—look at those clouds!' This is not an instance of doubting in my sense. This is a case of ordinary belief revision through a respectable process like conditionalization. In this case, I received some evidence (it's cloudy) that led me to abandon my belief about the picnic. My own perspective, which has as one of its commitments Pr(Picnic | Cloudy) = low, recommends a reduction of confidence in the proposition that there will be a picnic upon learning that it is cloudy. The revisions in response to doubt that I have in mind, in contrast, don't proceed by conditionalization. To see this, suppose that, for whatever reason, I couldn't recover the belief that I turned off the stove from the perspective of doubt (perhaps I set aside too much), and as a result I abandon the belief. Doing so would not have been the result of conditionalizing on 'I heard a fictional story about a house burning down'. After all, I don't take my having heard such a story to be any evidence whatsoever about the status of my stove.

In sum: subjecting a belief to doubt amounts to engaging in an inquiry. I'm asking what is, in a way, a logical question: does a certain perspective—one that's less committal than my usual perspective—permit transitions that form a path to the belief in question? If, upon subjecting a belief to doubt and realizing that it can't be recovered from doubt, we respond by abandoning the belief, I'll say that we've 'deferred to doubt' because we've adopted the attitude that the perspective of doubt recommends.

III

Higher Order Evidence. In STOVE, I managed to recover my belief from the perspective of doubt. But what if I can't? There are cases in which it is tempting to think that we ought to abandon belief upon

realizing that it can't be recovered from a perspective of doubt. I'm going to suggest that typical higher order defeat cases are of this sort. Consider:

> SLEEPY (adapted from Horowitz 2014): You are a police detective investigating a jewel theft. There are two suspects under consideration and, before examining any evidence, you assign 0.5 credence to each one being the thief. Late one night, after hours of cracking codes and scrutinizing photographs, you conclude that the thief was Lucy. In fact, it is Lucy, and you evaluated the evidence correctly. You call your partner, Alex. 'I've gone through all the evidence', you say, 'and it points to Lucy! I've found the thief!' But Alex is unimpressed. She replies, 'I know you've been up all night working on this. Your late-night reasoning has been awful in the past. You're always very confident that you've found the culprit, but under these circumstances you do no better than chance. So I'm not convinced.' You rationally trust Alex and believe that you've done no better than chance on such occasions.

The case is a bit artificial. Still, try to imagine yourself in this situation. How confident should you be after hearing Alex's testimony? Many think that maintaining your belief under these circumstances is unreasonable and that a 0.5 credence (which was your prior) would be the appropriate attitude upon learning about your no-better-than-chance track record. We'll call this 'the defeatist verdict'.

It turns out that standard ways of thinking about belief revision, like conditionalization, don't do a good job at capturing this verdict.[4] (It's worth flagging that this fact is crucial for what follows, and not at all obvious!) Some might take this as a reason to simply deny the defeatist verdict and claim that the detective (having in fact reasoned correctly about Lucy) should stick to their guns. But for others, the force of the defeat intuition is compelling enough to motivate thinking outside the box. One common approach to explaining why we should abandon belief in cases like SLEEPY is to appeal to what are called 'independence principles'.[5] Very roughly (and we'll get less rough shortly), independence principles say that when evaluating how likely you are to be right about whether *P*, you must do so in a way that is *independent* of, or sets aside, the reasoning about

---

[4] See, for example, Schoenfield (2018, MS) and Levinstein (MS). For related points, see also Christensen (2010) and Weisberg (2015).

[5] See Elga (2007), White (2009, 2010), Christensen (2010, 2011), Lasonen Aarnio (2014), Vavova (2014, 2018), Horowitz and Sliwa (2015).

P that is in question. To see how independence principles support the defeatist verdict, consider how one might respond to a stubborn detective who responds to her situation as follows:

> LUCKY ME: 'When I'm tired, my reasoning will sometimes lead me to the wrong conclusion. But not always. So the question is: how likely is it that I got the right answer *on this particular occasion*? Well, I got things right on this occasion if and only if Lucy is the thief. So is she? Let's look at the evidence. The fingerprint evidence says … and the letter she wrote says … and if I calculate the distance between the other suspect's house and the crime scene … so it is almost certainly her! This means I probably got things right on this occasion despite being sleepy. Lucky me!'[6]

The problem with the 'lucky me' response, say advocates of independence principles, is that the response essentially relies on the very reasoning that's being questioned. And this, they think, is inappropriate.

Despite their success at blocking 'lucky me' responses, when stated explicitly these independence principles can sound a bit odd. Why, when thinking about whether *P*, would it be rational to set aside reasoning or evidence that is relevant to *P*?

What I want to suggest here is that the appeal to independence principles will seem more natural if we think of what is going on in cases like SLEEPY as instances of doubt, just as in STOVE. The effect of higher-order evidence in cases like SLEEPY is that it becomes impossible to recover our belief from a specific perspective of doubt.[7] In STOVE, I have the resources to reason my way to the belief from the perspective of doubt in question. But in SLEEPY, if I subject my reasoning to doubt, I won't be able to use it to reason my way back to the belief. I won't be able to appeal to my own reliability about such matters either, because of the presence of the defeater. This means that in SLEEPY (because of the presence of the defeater) there's simply no way to recover the belief from the perspective of doubt in which the reasoning at issue is set aside. When the independence principles are telling us to reason in a way that 'sets aside' or 'brackets' certain

---

[6] This can be dramatized by imagining a version of the case in which the detective's evidence entails that Lucy is the thief, since no matter what else you add to the evidence, it will still entail that Lucy is the thief. But nothing about what follows requires entailment.

[7] See Schoenfield (forthcoming) for more on this point.

reasoning or evidence, they are encouraging us to reason from a certain perspective of doubt and adopt the belief-state such a perspective recommends. I am not aiming to defend any particular verdict about SLEEPY right now. I'm just suggesting that in so far as we're inclined to reduce confidence in such cases, this inclination is naturally thought of as a response to the realization that the belief we formed can't be recovered from doubt.

It will be important for what follows to be clear about what people who defend independence principles are thinking: they acknowledge that without bracketing the reasoning in question, or, in my terminology, reasoning from the perspective of doubt, the 'lucky me' response would make sense. The 'non-doubtful' perspective—the one that doesn't do any 'bracketing' and simply proceeds by conditionalizing—does not recommend a 0.5 credence (see the references above in note 4). But, they claim, a 0.5 credence is what you get when you reason from the perspective of doubt (or bracket the reasoning in question) and this fact figure in the explanation of why your credence should be 0.5 in such cases.

IV

The Challenge for Defeatism. So far, the story looks something like this: sometimes we subject beliefs to doubt. When we recover them from doubt, as in STOVE, we happily maintain belief. When we can't recover them from doubt, as in SLEEPY, we give them up. But that can't be the full story. For consider:

> SCEPTICISM: When I subject my belief that the sun will rise tomorrow to doubt in a way that sets aside my commitment to induction, I can't recover my belief.

This is just the old problem of induction. Most of us, upon realizing that our belief that the sun will rise tomorrow can't be recovered from a perspective in which some of our commitments are set aside, are inclined to believe that the sun will rise tomorrow anyway. Similarly, one might think, if I set aside all of my beliefs about the external world, I won't be able to recover my belief that I'm sitting at a café right now. But I still believe that I am.

So we don't think that for *any* belief and *any* perspective of doubt, if the belief can't be recovered from the perspective of doubt, we should abandon it. This, then, is the challenge: consider the cases in which

belief cannot be recovered from doubt. Some, like SLEEPY, are cases in which this realization motivates (many of) us to abandon belief. Others, like SCEPTICISM, are cases in which we shrug our shoulders and move on with our lives. Can we give a well-motivated account of why in some cases we defer to doubt and in others we don't?[8]

Before considering some proposals, let me be clear about who does and doesn't face this challenge. 'Steadfasters' think you should not reduce confidence to 0.5 in SLEEPY and so are immune from the challenge.[9] They have available to them a nice, clean view according to which we should always revise our beliefs by conditionalizing—nothing fancy involving independence principles or doubting needs to happen. Sceptics are also immune. They think that you should reduce confidence in SCEPTICISM, or rather, they think that it was unreasonable to have formed these non-sceptical beliefs to begin with. So they're also off the hook. Lastly, people who have solutions to scepticism of the 'convince the sceptic' variety might think that, in fact, in any case in which it's plausible that we should believe *P*, *P* can be recovered from *all* perspectives of doubt—no matter how sceptical. Perhaps, for example, sceptical perspectives turn out to be self-undermining (Rinard 2018). *It's the thought that there are cases in which we cannot recover belief from a perspective of doubt, combined with the thought that we should defer to doubt in some of these cases but not others that gives rise to the challenge.* Any view which doesn't countenance such a contrast doesn't need to explain it.

A number of people have proposed ways of meeting the challenge. I discuss a proposal from Schoenfield (forthcoming) in the appendix, but below I'll focus on ideas from Christensen (2011) and Vavova (2014). I'll argue that the proposal is unsuccessful and suggest, more generally, that there's a sense in which any proposal of this sort is bound to fail. At least, a certain version of the challenge cannot be met.

V

Reasonful versus Reasonless Perspectives of Doubt. David Christensen (2011) and Katia Vavova (2014, 2018) diagnose the difference between SLEEPY and SCEPTICISM as follows. In SLEEPY, they

---

[8] Discussion of challenges to defeatism along these lines can be found in Elga (MS), White (2010), Christensen (2011), and Vavova (2014, 2018).

[9] For views along these lines, see White (2010), Lasonen Aarnio (2014), and Titelbaum (2015).

say, the perspective of doubt is one that contains *good reasons* to think you are wrong: your evidence supports the claim that your reasoning was performed in a cognitively compromised state. In contrast, they claim, while the perspective that results from bracketing your commitment to induction lacks good reason to think you are right, it doesn't have a good reason to think you are wrong. Setting aside induction, the thought goes, you don't have much in the way of reasons to believe anything about such matters as whether the sun will rise. The general thought is that if the perspective of doubt *has* good reasons for thinking a mistake was made, you should defer to it, but if it merely *lacks* good reasons to think you got things right, you should not. Vavova (2018) formulates this proposal by distinguishing two principles, the first of which she endorses, the second of which she rejects:

> *Good Independent Reason Principle* (GIRP): To the extent that you have good independent [undefeated] reason to think that you are mistaken with respect to *p*, you must revise your confidence in *p* accordingly. (Vavova 2018, p. 145)

> *No Independent Reason Principle* (NIRP): To the extent that you [lack] good independent [undefeated] reason to think that you are [correct] with respect to *p*, you must revise your confidence in *p* accordingly. (Vavova 2018, p. 148)

I'll call this proposal 'GIRP-not-NIRP'.

VI

Motivational Difficulties. Suppose that you haven't (yet) adopted a principle concerning when to defer to the perspective of doubt. You're in a situation in which your ordinary, non-doubtful perspective (one that doesn't do any bracketing) recommends being more confident than not in *P*, while the perspective of doubt (one that brackets your reasoning concerning some evidence *E*) recommends agnosticism. Let's call the perspective that doesn't do any bracketing 'the fat perspective', and the perspective that brackets your reasoning concerning *E* 'the skinny perspective'. I'll assume throughout the argument that all of the perspectives under discussion are coherent in the sense that they don't contain conflicting commitments. (It follows from this, for example, that the fat perspective, which recommends belief, doesn't also contain commitments that recommend

refraining from belief.) You find yourself wondering whether to defer to doubt.

Christensen and Vavova come forward and offer their proposal: you *should* defer to doubt if the perspective of doubt has good reasons for thinking you got things wrong, but you *should not defer* if it merely lacks good reasons for thinking you got things right.

My aim in this section is to argue that such a principle will, from your current perspective, look unmotivated. Indeed, the argument below suggests something more general: that there is no principle in the vicinity of GIRP-not-NIRP, that will look well motivated to a deliberator in the situation described above. The argument starts from the assumption that all deliberation takes place from some perspective—that is, it takes place in the context of some set of commitments (even if it's the empty set—the maximally permissive perspective). You can't deliberate, so to speak, 'from nowhere'.

With this assumption in hand, we can ask the following question: when you are deliberating about whether to defer to doubt, what is the nature of the perspective *from which* this deliberation (about whether to defer) takes place? We have two possibilities to consider: the case in which the perspective is fat—it doesn't bracket the reasoning concerning *E*—and the case in which it is skinny—it brackets the reasoning concerning *E*.[10] I'll consider each possibility in turn.

Suppose first that the perspective from which you're deliberating about whether to defer to doubt is fat. Given that no bracketing is taking place, *E*, and your reasoning concerning *E*, are available for use. Because straightforward conditionalization on your total evidence recommends being opinionated, if you're considering whether to defer to the perspective of doubt, *but you haven't yet done so*, it will look as though the thing to do (if you're interested in accuracy) is to conditionalize—to make use of all the resources at your disposal—and so not defer to doubt.[11]

This holds even if, as in the case of SLEEPY, the perspective of doubt contains good reasons for thinking you made a mistake. For although *the perspective of doubt* contains good (and undefeated) reasons for thinking you made a mistake, *you,* who have more

---

[10] I'm assuming throughout this argument that the only ways in which the set of commitments from which you're deliberating about whether to defer to doubt diverge from your pre-deliberation set of commitments involve differences in commitments concerning *E*.

[11] For arguments supporting connections between accuracy and conditionalization, see Greaves and Wallace (2006), Briggs and Pettigrew (2020), and Horowitz (2021).

epistemic resources available than the perspective of doubt, have good reasons for thinking that you did *not* make a mistake—that you got lucky. One way of putting this is that while the perspective of doubt has a defeater, you (in virtue of occupying the fat perspective) have a defeater *of* that defeater—for you can appeal to your reasoning concerning E and conclude that this is one of the occasions in which you got things right.[12]

Let's now see what happens if I'm deliberating about whether to defer to doubt from a perspective that doesn't permit me to make use of my reasoning about E—a skinny perspective. If my perspective forbids making use of the reasoning concerning E, the perspective from which I'm deliberating about whether to defer to the perspective of doubt *just is* the perspective of doubt. This perspective is going to recommend deferring to doubt, regardless of the structure of reasons. Why? We've already stipulated that we're dealing with cases in which the perspective of doubt recommends abandoning belief. If you're occupying some deliberative perspective, and you ask it, 'Should I maintain the beliefs you recommend abandoning?', the answer is going to be a resounding no. This is true regardless of what sorts of reasons the perspective has available. Suppose, for instance, that your perspective of doubt recommends a credence in P of 0.5 on the basis of the Principle of Indifference: there are no reasons (in the skinny perspective) to believe P or $\neg P$, so it recommends dividing your credence evenly between them. If you ask such a perspective, 'Should I believe P?', the answer is going to be no. Suppose you respond, 'But you don't have much to go on in recommending 0.5. You're just recommending 0.5 because you're in a very evidentially unfortunate position with respect to P. You lack reasons to believe one thing or another.' The perspective will come back with, 'Exactly. That is why I'm telling you to assign 0.5 to P.'

The general point is this: if we bring our attention to cases in which (coherent) fat perspectives recommend maintaining belief and

---

[12] Note that this is not dogmatism-paradox-reasoning, according to which, whenever you believe Q, you should regard any evidence against Q as misleading (since after all, according to you, Q is true!). For if e is evidence against Q, then so long as your credence in in Q is not 1, your perspective will contain not only a high *unconditional* credence that Q but also a low *conditional* credence in Q given e. This means that reducing confidence in Q can be motivated by the commitments in your perspective. The issue here is that conditionalizing doesn't motivate agnosticism in higher-order defeat cases (recall the fact I flagged as crucial earlier with references in note 3), and so even when we account for your conditional credences, your perspective will not recommend agnosticism.

(coherent) skinny perspectives recommending abandoning belief, the answer to the question of whether to defer to doubt will be fully determined by the nature of the perspective *from which* we're deliberating about whether to defer. If that perspective has not already bracketed the reasoning that's in question, abandoning belief will look bad (because, by stipulation, your perspective has commitments that recommend belief). If it has already bracketed that reasoning, abandoning belief will look good (because we're focusing on cases in which the bracketed perspective recommends agnosticism). The challenge that is posed by settling questions about whether to defer to doubt deliberatively is that the very act of deliberating about which of two perspectives to adopt requires that you've already adopted one of them.

VII

Objections and Responses. In this section I'll consider some objections and responses to the argument presented in the previous section.

Objection 1. When we're deliberating about whether to defer to doubt, we're deliberating from neither of the two perspectives you described. We're deliberating from a *third* perspective, one that hasn't yet made up its mind about whether to permit the reasoning about *E*.

  Response. A perspective is just a set of commitments. So it's a logical truth that every perspective is one whose commitments do or do not permit the reasoning about *E*. There is no 'third' perspective that is 'neutral' with respect to whether it permits reasoning with *E*. It may, however, be indeterminate which of two perspective you're occupying. If it's indeterminate which perspective you're occupying, then it's indeterminate whether your perspective recommends deferring to doubt. This still doesn't provide us with a consideration that favours doing one thing rather than another.

Objection 2. I suggest we reframe the role of GIRP-not-NIRP. The idea is not to tell us what to do when we *can't* recover a belief from a perspective of doubt (that is, defer to doubt or not). Rather, Christensen and Vavova are trying to show us how we can *recover* belief from sceptical perspectives of doubt.

Response. GIRP-NOT-NIRP does not offer a way to recover belief from sceptical perspectives of doubt. To see why, it will be helpful to first make note of a principle that, if included in the perspective of doubt, *would* arguably allow us to recover belief:

> BELIEF-WITHOUT-REASONS: When you have no reasons to believe either *P* or ¬*P*, it's permissible to believe *P*.

If we had BELIEF-WITHOUT-REASONS at our disposal, then somebody who has set aside their commitment to induction might recover belief as follows: 'True, I have no reason to believe that the sun will rise tomorrow (I've set aside induction). But I also don't have reasons to believe that it won't. So BELIEF-WITHOUT-REASONS tells me that it's fine for me to believe it will rise tomorrow.' But GIRP-NOT-NIRPers don't (and shouldn't) accept BELIEF-WITHOUT-REASONS. They're not claiming that it's fine to believe empirical propositions for no reason at all. (Presumably they don't think you can rationally believe that a black marble will be drawn from an urn of black and white marbles with unknown ratio because you lack a reason to believe both that it will be black and that it won't be.) They're only claiming that it can be permissible to believe some propositions with no *independent* reason—no reason that exists *in some trimmed down version of your perspective*. But if your perspective *is* the sceptical perspective, then the sceptical perspective isn't a trimming down of your perspective: it's all you've got. The issue, for the inductive sceptic, isn't a lack of independent reason, it's a lack of reason at all. That's why BELIEF-WITHOUT-REASONS might help a sceptic recover belief from doubt, but the permission to believe without independent reason will not.

I want to flag that I'm not aiming to give an argument for the claim that nobody could have commitments that allow them to recover beliefs from sceptical perspectives. People might have all sorts of commitments that allow them to make all sorts of interesting moves from a variety of different perspectives. As I mentioned at the outset, if you have a 'convince the sceptic'-type of solution to scepticism, the challenge described in this paper doesn't arise for you. But GIRP-NOT-NIRP, at least as stated, does not provide such a solution. For this reason, I think Christensen and Vavova are best interpreted as offering us a way of ignoring the sceptic, not of convincing her.

Objection 3. You claim that GIRP-NOT-NIRP will appear unmotivated to a deliberator trying to decide whether to defer to doubt. But the

principle seems intuitively compelling. Why isn't that motivation enough?

Response. I want to propose an error theory for the apparent appeal of GIRP-not-NIRP. What is arguably correct is that if you have a good undefeated reason to think you made a mistake, you should think you made a mistake. But that's not what GIRP-not-NIRP says. Rather, it says that if *the perspective of doubt* has a good undefeated reason to think you made a mistake, *you* should think you made a mistake. However, it's just misguided to think that any undefeated reason had by a perspective that's skinnier than yours is an undefeated reason of *yours*. (Note that nobody thinks, for instance, that the following is an attractive principle: if your less informed neighbour has good reason to think you made a mistake, you should think you made a mistake.) GIRP, then, should only look appealing if we're assuming that the question of whether to defer to doubt should be settled *from* the perspective of doubt. For if that were the case, the perspective of doubt's good undefeated reasons would just be *your* good undefeated reasons. But if we assume that the question of whether to defer to doubt should always be settled from the perspective of doubt, we would end up sceptics. This is because any coherent perspective of doubt (including sceptical ones) will recommend deferring to itself. And as we saw in response to the second objection, GIRP-not-NIRP won't help: it can't fish us out of the sceptical perspective if that's where we're starting from.

VIII

What About Rationality? I haven't answered the following question: when is it *rational* to defer to the perspective of doubt? For all I've said, something like GIRP-not-NIRP describes a truth about rationality. But, *as a deliberator*, I'm not satisfied by GIRP-not-NIRP, because the view can't be motivated from the perspective of somebody aiming at accuracy who is trying to decide whether to defer to doubt. All this is to say is that, given the way that I'm approaching the question (imagining a deliberator trying to decide whether to defer to doubt), principles like GIRP-not-NIRP, whether they are truths about rationality or not, have no traction.

My own view is that rationality is important because truth is important, and principles of rationality are meant to help us in our pursuit of the truth. So I'm inclined to think that if a proposed

principle of rationality can't be motivated from the truth-seeker's perspective, that is problematic. You may disagree with me on this front, and here is not the place to get into these meta-epistemological questions. Suffice it to say that if your view is right and my view is wrong, then there may just be two interesting intellectual projects worth pursuing: in addition to theorizing about what's rational, it may also be interesting to try to figure out what deliberative moves are available from our own perspective when we're seeking the truth.

IX

Why Do We Defer to Doubt When We Do? My inclination in response to the considerations above is to think that questions of whether to defer to doubt are simply not ones that can be settled deliberatively. They belong to that strange category of questions in which the very act of deliberating about how to proceed requires that you've already made up your mind. But there is still a descriptive question that puzzles me: why are we, *in fact*, inclined to defer to doubt in some cases but not others? In this section I offer some speculations about why we might have the doubt-deferring tendencies that we do.

The first concerns what I'll call 'epistemic absurdity'. Here's the thought: what the classical sceptical arguments teach us is that if we set aside too much, we won't find our way back—we'll end up in the sceptical abyss. Most of us, however, have learned to live with this fact. We know that we can't recover our ordinary beliefs if we set aside perception, induction, memory, other minds, and so forth. But we don't set all that aside; we embrace these commitments and live our lives accordingly. This phenomenon is the epistemic analogue of what Nagel (1971) dubbed 'the absurd'. What Nagel was interested in was the fact that if we take a big step back from our *practical* perspective, we find ourselves with no way of returning: our pursuits, from that perspective, look trivial and meaningless. Nonetheless, we engage in them anyway. He writes, 'We see ourselves from outside, and all the contingency and specificity of our aims and pursuits become clear. Yet when we take this view and recognize what we do as arbitrary, it does not disengage us from life, and there lies our absurdity' (Nagel 1971, p. 720). What is emphasized less in Nagel is that if we take *small* step backs, what we do doesn't look arbitrary at all. Suppose I step back from my commitment to grade papers today.

Does grading papers look arbitrary? Not at all. I value my students, my job, and following through on my promises.

Our situation in the epistemic realm is similar. Usually, if I take a small step back, I can recover my belief. The problem posed by higher-order defeaters is that they block ways of recovering belief that are usually available (specifically, ways that appeal to the fact that we're generally reliable about some domain). In the presence of a higher-order defeater, even taking a small step back forces me to confront the absurdity (in the sense above) of maintaining belief. While most of us have reconciled ourselves to some degree of epistemic absurdity—we know that if we take a huge step back, our beliefs can't be recovered—the fact that just a small step back prevents recovery of the belief is something that tends to make us uncomfortable. This discomfort can lead us to abandon belief. But agnosticism in the face of absurdity isn't the result of the chugging along of the Bayesian machinery—it's something else.

All this is to say that one possible explanation for our tendency to defer to doubt in some cases but not others may have something to do with the extent to which we're willing to tolerate epistemic absurdity. But there may be more mundane explanations as well—there may be good reasons that we've evolved to defer to doubt in some cases but not others, even if these tendencies aren't well motivated philosophically.[13]

To see why, imagine that you're programming a robot that's going to explore Mars. Suppose first that you're certain that the robot will respond to evidence in *exactly* the way you tell it to, no matter what. (I'm not claiming this is realistic.) Now you wonder: 'Should I program the robot in such a way that it doubts its capacities to respond to evidence?' No! For *I* know that its capacities are shipshape. So even if the robot were to encounter some Martians who say to it, 'You know, robots like you tend to malfunction in our environment', and provide a track record of malfunctioning robots like this one, I'll want my robot to ignore all that, since *I* know that such evidence would be misleading (my robot will not malfunction).

The less confident I am in my robot's capacities, the more I'll want the robot to take into account the possibility that it malfunctioned. Suppose, for example, that I think that if the robot's battery

---

[13] See Pinillos (2019) for an extended discussion of evolutionary explanations of this variety.

is running low, it will do no better than chance at performing certain calculations. In that case, I'll want to program the robot in such a way that if it performs a calculation concerning *P*, and discovers that its battery is low, it abandons the results of the calculation and reverts to its prior probability. If *I'm* leaving open the possibility that the robot will make a mistake, I'll sometimes want the robot to defer to the perspective of doubt.

However, even if I leave open the possibility that my robot will malfunction in certain conditions, I certainly will not want the robot to defer to any old perspective of doubt it might entertain. For suppose I program the robot with all sorts of beliefs about Earth (such beliefs may, after all, prove useful if the robot is called upon to make various interplanetary comparisons). If the robot one day starts wondering why it has all these beliefs about a planet it's never set foot on, and brackets the Earth beliefs, it won't be able to recover them from doubt. But I won't want the robot to defer to this doubtful perspective. So I'll simply program the robot to shrug its shoulders, as we do in response to sceptical worries. More generally, I'll program the robot so that it defers to doubt under all and only the circumstances in which *I* have doubts about its capacities or beliefs. If the robot tried to come up with an epistemology that justifies its doubt-deferring dispositions, it would most certainly fail. Its dispositions are simply a result of the varying degrees of confidence that I, the designer, have in various propositions.

What's the moral of this story? A very speculative proposal about our own tendences is that, in some sense, we're like this robot. We've been 'programmed' to be sensitive to the possibility of certain kinds of errors ('Did I reason about this particular matter correctly?') but not others ('Is there an external world?', 'Will the future be like the past?'). If which sorts of errors we're concerned with is explained by the fact that concern about certain sorts of errors rather than others was conducive to survival, then it's likely that the sorts of errors we're sensitive to are errors that we were, at some point, in fact prone to make. If there is an external world, it certainly won't do you any good to worry that there isn't, and if the past proceeded in a relatively patterned way for a while, creatures that made inductive inferences would have done better than sceptics. The sorts of errors we find ourselves sensitive to might be quite a hodgepodge, and there might not be much to say about what the members of the hodge-podge have in common that goes beyond the fact that concerns

about some errors are, or were, for completely contingent reasons, more useful than others.

But don't take any comfort in this hypothesis. Don't think that you can motivate your tendency to get worried about your reasoning when you're sleepy, but not be a sceptic, by appealing to the proposal that you've been programmed in ways that make you sensitive to errors you are in fact prone to make. For this proposal is only plausible from a perspective in which you're not doubting the external world or induction. Currently, I'm not occupying a perspective of doubt with respect to these matters, so I'm perfectly happy putting this proposal on the table as a possible explanation of our doubt-deferring tendencies.

<div align="center">X</div>

Conclusion. Patrick Shanley, in his play *Doubt*, describes doubt as 'a wordless Being' that 'moves just as the instant moves; it presses upward without explanation, fluid and wordless' (2005, p. vii). I agree with this characterization. We can always choose to *entertain* doubt—we can notice what would follow if we did or didn't rely on various things that we generally take for granted. But whether to actually *take up* that perspective—to form or abandon the beliefs it recommends—is not something we can decide deliberatively. Deference to doubt should be thought of as something that simply happens to us, without explanation, fluid and wordless.[14]

Department of Philosophy
University of Texas at Austin
2210 Speedway Road
Mail Stop c*3500,* wag *316*
Austin, tx *78712,* USA
mschoenfield@utexas.edu

References

Berger, Dominik, and Nilanjan Das 2020: 'Accuracy and Credal Imprecision'. *Noûs*, 54(3), pp. 666–703.

Briggs, R. A., and Richard Pettigrew 2020: 'An Accuracy-Dominance Argument for Conditionalization'. *Noûs*, 54(1), pp. 162–81.

Builes, David, Sophie Horowitz, and Miriam Schoenfield 2022: 'Dilating and Contracting Arbitrarily'. *Noûs*, 56(1), pp. 3–20.

Camp, Elisabeth 2019: 'Perspectives and Frames in Pursuit of Ultimate Understanding'. In Stephen R. Grimm (ed.), *Varieties of Understanding: New Perspectives from Philosophy, Psychology and Theology*, pp. 17–45. Oxford: Oxford University Press.

Christensen, David 2010: 'Higher Order Evidence'. *Philosophy and Phenomenological Research*, 81(1), pp. 185–215.

——2011: 'Disagreement, Question-Begging and Epistemic Self-Criticism'. *Philosopher's Imprint*, 11(6), pp. 1–21.

Elga, Adam 2007: 'Reflection and Disagreement'. *Noûs*, 41(3), pp. 478–502.

——MS: 'Lucky to be Rational'. Unpublished manuscript.

Greaves, Hilary, and David Wallace 2006: 'Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility'. *Mind*, 115(459), pp. 607–32.

Horowitz, Sophie 2014: 'Epistemic Akrasia'. *Noûs*, 48(4), pp. 718–44.

——2021: 'Shrinking Three Arguments for Conditionalization'. *Philosophical Perspectives*, 35(1), pp. 303–19.

——and Paulina Sliwa 2015: 'Respecting All the Evidence'. *Philosophical Studies*, 172(11), pp. 2835–58.

James, William 1896: 'The Will to Believe'. In his *The Will to Believe, and Other Essays in Popular Philosophy*. New York: Longmans, Green, and Co.

Lasonen-Aarnio, Maria 2014: 'Higher-Order Evidence and Limits of Defeat'. *Philosophy and Phenomenological Research*, 88(2), pp. 314–45.

Mayo-Wilson, Conor, and Gregory Wheeler 2016: 'Scoring Imprecise Credences: A Mildly Immodest Proposal'. *Philosophy and Phenomenological Research*, 93(1), pp. 55–78.

Nagel, Thomas 1971: 'The Absurd'. *Journal of Philosophy*, 68(20), pp. 716–27.

Pinillos, N. Ángel 2019: 'Skepticism and Evolution'. In Brian Kim and Matthew McGrath (eds.), *Pragmatic Encroachment in Epistemology*, pp. 141–69. New York: Routledge.

Rinard, Susanna 2018: 'Reasoning One's Way Out of Skepticism'. In Kevin McCain and Ted Poston (eds.), *The Mystery of Skepticism: New Explorations*, pp. 240–64. Leiden: Brill.

Schoenfield, Miriam 2017: 'The Accuracy and Rationality of Imprecise Credences'. *Noûs*, 51(4), pp. 667–85.

——2018: 'An Accuracy-Based Approach to Higher-Order Evidence'. *Philosophy and Phenomenological Research*, 96(3), pp. 690–715.

——forthcoming: 'Meditations on Beliefs Formed Arbitrarily'. To appear in Tamar Szabó Gendler and John Hawthorne (eds.), Oxford Studies in Epistemology, 7. Oxford: Oxford University Press.

——MS: 'Bayesianism and Higher Order Defeat'. Unpublished manuscript.

Seidenfeld, Teddy, Mark J. Schervish, and Joseph B. Kadane 2012: 'Forecasting with Imprecise Probabilities'. International Journal of Approximate Reasoning, 53(8), pp. 1248–61.

Shanley, P. 2005: *Doubt: A Parable*. New York: Theatre Communications Group.

Titelbaum, Michael G. 2015: 'Rationality's Fixed Point (or: In Defense of Right Reason)'. In Tamar Szabó Gendler and John Hawthorne (eds.), Oxford Studies in Epistemology, *5, pp.* 253–94. Oxford: Oxford University Press.

Vavova, Katia 2014: 'Debunking Evolutionary Debunking'. In Russ Shafer-Landau (ed.), Oxford Studies in Metaethics, *9, pp.* 76–101. Oxford: Oxford University Press.

——2018: 'Irrelevant Influences'. Philosophy and Phenomenological Research, 96(1), pp. 134–52.

Weisberg, Jonathan 2015: 'Updating, Undermining, and Independence'. *British Journal of Philosophy of Science*, 66(1), pp. 121–59.

White, Roger 2009: 'On Treating Oneself and Others as Thermometer'. *Episteme*, 6(3), pp. 233–50.

——2010: 'You Just Believe that Because …' *Philosophical Perspectives*, 24(1), pp. 573–615.

Appendix: Precise versus Imprecise Doubt

In this appendix I'll discuss a proposal from Schoenfield (forthcoming) about how to meet the challenge posed by higher-order defeat. A similar idea in the context of peer disagreement can be found in Elga (2007).

Here's the thought: in SLEEPY, it's plausible that the perspective of doubt contains a 0.5 credence in the proposition that Lucy committed the crime. After all, in the perspective of doubt, you can't rely on the reasoning you just did (it's been subject to doubt) and your prior was stipulated to be 0.5. Contrast this with SCEPTICISM. Suppose you were to give up your commitment to induction. How confident would you be that the sun will rise tomorrow? You might think that removing your commitment to induction wouldn't result in a 0.5

credence in the proposition that the sun will rise tomorrow. Perhaps, in such a highly impoverished perspective, you'd simply *have no idea* how likely it is that the sun will rise tomorrow. In such a case, the resulting attitude might be best represented by an imprecise credence like [0, 1]. So, in contrast to SLEEPY, where the perspective of doubt is precise, in SCEPTICISM, one might think, the perspective of doubt is imprecise.

Why would the difference between precise and imprecise perspectives of doubt be relevant to the question of whether to defer to doubt? One difference is that while perspectives of doubt containing sharp credences recommend those credences from an accuracy point of view (they regard those credences as most expectedly accurate[15]), imprecise perspectives don't recommend their imprecise credences.[16] More specifically, Builes et al. (2022) and Schoenfeld (forthcoming) argue that an attitude of, say, [0, 1] towards the proposition that the sun will rise tomorrow is not a state that recommends against moving to a more opinionated state. So one might argue that we *can* in fact recover our belief that the sun will rise tomorrow from the perspective of doubt in which induction is set aside, so long as the perspective of doubt is (sufficiently) imprecise. For while the perspective doesn't recommend believing that the sun will rise tomorrow, it doesn't forbid it either. This means that, if we choose to, we can simply transition to a more confident attitude towards the proposition from the perspective of doubt, in a way that the perspective of doubt permits.[17] (This is, in effect, a 'convince the sceptic' solution to scepticism, though it goes by way of pointing to a permission rather than a requirement.)

The problem with the proposal is that I'm not convinced it's going to do all the anti-sceptical work we might want it to. Some sceptical arguments are motivated by principles like the Principal of Indifference: that in the absence of reasons for treating different

---

[15] I'm relying here on the thought that our accuracy measures are 'strictly proper' and that the credences in question are probabilistic. Note that the dialectic of this paper doesn't get off the ground if we use the main competitor to a strictly proper rule—the absolute value score. For on the absolute value score, a 0.5 credence permits a credence of 1 or 0. So it wouldn't be true that the perspective of doubt in SLEEPY recommends 0.5.

[16] See Seidenfeld et al. (2012), Mayo-Wilson and Wheeler (2016), Schoenfeld (2017), Berger and Das (2020), and Builes et al. (2022).

[17] Here too I am relying on 'accuracy-first' epistemology. For all I've said, there could be non-accuracy-related considerations that favour remaining imprecise.

hypotheses differently we should distribute our credence evenly over
the relevant possibilities. So, if I'm considering how things look
from a perspective in which I don't rely on induction, and I wonder
whether the sun will rise tomorrow, one might argue that, rather
than being spread all over the interval, my credence will be 0.5 that it
will rise and 0.5 that it won't rise. I'm not claiming that all sceptical
arguments are motivated by this sort of reasoning, aiming to defend
the Principle of Indifference, or this particular application of it. My
point is just that I don't think the precise/imprecise distinction gives
us a *general* solution to the problem, because it seems as though
how precise or imprecise your credence is in a sceptical perspective
will depend a lot on the details of your scepticism. I don't see an
argument for the claim that sceptical worries *always* go along with
extremely imprecise probabilities.