

Deferring to Doubt

The sceptic with his whole nature adopts the doubting attitude; but which of us is the wiser, Omniscience only knows.

- *William*

James

Abstract: When we doubt a belief, we examine how things look from a perspective in which that belief is set aside. Sometimes we care about what that perspective recommends and, as a result, we abandon the belief we've been doubting. Other times we don't: we recognize that a perspective in which a certain belief is set aside recommends abandoning it, but we go on believing it anyway. Why is this? In this paper, I'll consider and then reject some proposals concerning when to defer to the perspective of doubt. I'll argue that ultimately the question of whether to defer to doubt on any given occasion can't be answered through rational deliberation aimed at truth or accuracy. If I'm right, this means that a certain challenge facing defeatist views about higher order evidence cannot be met: namely, providing a motivation for abandoning belief in cases of higher order evidence, but not becoming a global skeptic.

1. Introduction

Sometimes we doubt a belief because we receive evidence that it was formed in a dubious manner. Sometimes we doubt because we encounter disagreement. Sometimes we doubt in response to skeptical 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, there will be no way to motivate the judgment that we should reduce confidence in response to some forms of doubt but not others. I take these considerations to motivate a radical form of permissivism about higher order defeat and skepticism (one can abandon belief in response to higher order evidence, or not; one can be a skeptic, or not) though, as we'll see, others may draw different conclusions.

2. Doubt

Let's start with a simple example:

Stove: You're walking to work one morning, listening to a podcast. You hear a fictional story about a house that burned down because someone left the stove on. You start worrying that you forgot to turn off your stove. You 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." You maintain your belief and move on with your day.

This is a story in which you subject a belief (that your stove is off) to doubt. What we're interested in, when we subject a belief to doubt, is whether we can reason our way 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 you could: you recover your belief that the stove was off by appealing to your belief that you cleaned the stove.

Precisely which commitments are set aside when I subject my belief¹ that P to doubt? When I subjected my belief that the stove was off to doubt, I of course wasn't willing to rely on that very belief. But there are several other beliefs I wasn't willing to rely on as well. For example: the belief that I moved the stove knob in a certain direction, that the stove is off or $2+2 = 5$, and that there are no open flames in my kitchen. There are, however, plenty of commitments that I did not set aside in doubting 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. So on the way I'm thinking about doubt, there is no universal characterization of what we set aside when we doubt a particular 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 that permit certain cognitive transitions and forbid

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.

others. 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 head. It permits the transition from a visual perception as of P to a belief that P, and forbids transitions that commit the gamblers fallacy. What do I mean by “truth-aimed”? 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/are conducive to forming false beliefs, and it will permit transitions that it regards as having high expected accuracy/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.

Note that it’s not only beliefs that we subject to doubt. Inferences, transitions, or reliance on certain capacities can be doubted as 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 usage of the phrases “doubting” and “subjecting to doubt” are 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 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 conditionalizing. In this case, I received some evidence (it’s cloudy) that led me to abandon my belief about the picnic. This is not what goes on in the kinds of doubt I have in mind. 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.” I don’t take my having heard such a story to be any evidence whatsoever about the status of my stove. So the crucial thing about doubt in my sense is that modifying one’s beliefs in response to doubt is *not* an instance of standard conditionalization (more on that soon).

In sum: subjecting a belief to doubt, in my sense, 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.

2. 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 argue that typical "higher order evidence" cases are of this sort. Consider:

SLEEPY (adapted slightly 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 all points to one person! I've found the thief!" But Alex is unimpressed. She replies: "I can tell 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." But how to motivate it? In particular, how can we explain what's wrong with the following response to the case:

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!”²

In the higher order evidence literature, the “lucky me” response is standardly blocked by appeal to what are called “independence principles.”³ Very roughly, independence principles say that when evaluating how likely you are to be right about whether P in such cases, you should do so in a way that is *independent* of, or sets aside, the reasoning in question. The problem with the “lucky me” response, the thought goes, is that it essentially relies on the very reasoning that’s being questioned, and, according to advocates of independence principles, this is inappropriate.

These sorts of independence principles, when stated explicitly, can sound odd. Why on earth, when thinking about whether P, could it be rational to set aside reasoning or evidence that is relevant to P? I’ll come back to this question in the moment, but for now the point is just that something like an independence principle is required to get the defeatist verdict. One way to see this is to note that you can’t get the defeatist verdict by ordinary conditionalization. While conditionalizing might allow for some reduction of confidence, it won’t take us all the way to agnosticism. I argue for this in detail in (ms.) so I won’t rehearse the argument for it here.⁴ (But I want to flag that this claim is both crucial for what follows, and not at all obvious!). Because higher order defeat isn’t well modeled as an instance of conditionalization, we need some other belief revision principle – one that recommends a *violation* of conditionalization in these cases – and independence principles, which tell us to bracket some of our commitments in the face of defeat, do just that work.

² This can be dramatized by imagining a version of the case in which E entails P, since no matter what else you add to E, it will still entail P. But nothing about what follows requires entailment.

³ Elga (2007), White (2009, 2010), Christensen (2010, 2011), Lasonen Aarnio (2014), Vavova (2014, 2018), Horowitz and Sliwa (2015).

⁴ For related points see also Schoenfield (2018), Christensen (2010), White (2010) (specifically his comparison of “Lucky Russel” and “Lucky Moore”) and Weisberg (2015).

What I want to suggest here is that appealing to independence principles will seem a lot more natural if we think of what is going on in these cases as instance of doubt, just like 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. 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 the reasoning I just did to doubt, then I won't be able to use it to reason my way to the belief. I won't be able to appeal to my own reliability about such matters either, because I know I'm sleepy. So there's no way to recover the belief from this perspective of doubt. When the independence principles are telling us to reason in a way that "sets aside" or "brackets" certain reasoning, 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*. I'm just suggesting that insofar 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 setting aside/bracketing/deferring to the perspective of doubt, the "lucky me" response would make sense. The "non-doubtful" perspective - the one that doesn't do any "bracketing," does not recommend a 0.5 credence. But, they claim, a 0.5 credence is what you get when you set aside/bracket/reason-from-the-perspective-of-doubt, and this fact figures in the explanation of why your credence should be 0.5 in such cases.

4. The Challenge for Defeatism

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

SKEPTICISM: 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 have hands. But I still believe that I do.

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 SKEPTICISM, 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?⁵

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 SLEEPY⁶ and so are immune from the challenge. 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. Skeptics are also immune. They think that you should reduce confidence in SKEPTICISM, or rather, they think it was unreasonable to have formed these non-skeptical beliefs to begin with. So they're also off the hook. Lastly, people who have solutions to skepticism of the "convince-the-skeptic" 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 skeptical. Perhaps, for example, skeptical perspectives turn out to be self-undermining.⁷ 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, 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.

⁵ Discussion of challenges to defeatism along these lines can be found in Elga (ms.), White (2010), Christensen (2011), and Vavova (2014, 2018).

⁶ For views along these lines see White (2010), Lasonen Aarnio (2014) and Titelbaum (2015).

⁷ Rinard (2018).

6. The Second Proposal: Reasonful versus Reasonless Perspectives of Doubt

David Christensen (2011) and Katia Vavova (2014, 2018) offer a different sort of proposal. Their thought is that, in SLEEPY, the perspective of doubt is one that contains *good reasons* to think you got things wrong: your evidence suggests 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 got things right, it doesn't *have* a good reason to think you got things 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” (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” (148).

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

6. Motivating The Proposal

I'll now argue for the following claim: if you're a deliberator (aiming at accuracy), and you find yourself wondering whether to defer to the perspective of doubt in some particular case, the idea that you *should* defer to doubt if you have good reasons for thinking you got it wrong, but *should not defer* if you merely *lack* good reasons for thinking you got it right, will look unmotivated. Indeed, I'll argue, there is no principle in the vicinity of GIRP-not-NIRP, that will look well motivated to a deliberator trying to decide whether to defer to doubt in a given case.

Here's the idea: Suppose I find myself deliberating about whether to defer to doubt in some particular case in which I've come to recognize that a perspective of doubt (one that brackets that my reasoning concerning some evidence E), recommends agnosticism about P, but the non-doubtful perspective (which simply proceeds by conditionalization, and doesn't do any bracketing) recommends being more confident than not in P. Because straightforward conditionalization recommends being opinionated, if I'm considering whether to defer to the perspective of doubt, *but I haven't yet done so*, it will look like the thing to do (if I'm interested I accuracy) is to conditionalize - and so not defer to doubt.

This holds even if, as in the case of SLEEPY, the perspective of doubt contains good reasons for thinking I made a mistake. This is because, although *the perspective of doubt* contains good (and undefeated) reasons for thinking I made a mistake, *I*, who have more epistemic resources available to me than the perspective of doubt (since I have not done any bracketing), have good reasons for thinking that I did *not* make a mistake - that I got lucky. One way of putting this is that while the perspective of doubt just has a defeater, I have a defeater *of* that defeater - for I (not occupying the perspective of doubt) can appeal to my reasoning about E and conclude that this is one of the occasions in which I got things right.⁸

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. If I'm not willing to rely on the reasoning about E in this perspective (and I'm not adding or removing additional commitments) the perspective from which I'm deliberating about whether to defer to the perspective of doubt *just is* the perspective of doubt. That 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 adopt 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 to go on. Suppose that your perspective of doubt recommends a credence in P of

⁸ 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 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 evidence defeat cases (see earlier references), and so even when we account for your conditional credences, your perspective will not recommend agnosticism.

0.5 on the basis of the Principle of Indifference: there aren't reasons to believe P or $\sim P$ 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. If you say: "But you don't have much to go on in recommending 0.5. You're just assigning 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 respond - "exactly. That is why I'm telling you to assign 0.5 to P ."

The problem then with GIRP-not-NIRP is that it has no deliberative grip. If you're an ordinary conditionalizer and not doing any bracketing, following the recommendations of GIRP and deferring to a bracketed perspective in cases like SLEEPY will look like a bad idea (as bad of an idea as deferring to somebody who has less evidence or cognitive resources than you do). If, however, you find yourself deliberating about whether to defer to doubt from a perspective that *doesn't* include the commitments in question - that is, you've already done some bracketing when you're trying to decide whether to defer to doubt - then that perspective, *by stipulation*, is going to tell you to abandon belief. And this will hold true regardless of the structure of reasons.

The more general point is this: when we're deliberating about whether to defer to doubt, we're always deliberating from some perspective. The question of whether to defer to doubt, is fully determined by whether the perspective from which we're deliberating has already bracketed the reasoning that's in doubt. If it hasn't, abandoning belief looks bad (because conditionalizing looks better). If it has, abandoning belief looks good (because we're focusing on cases in which the bracketed perspective recommends abandoning belief). No other considerations will be relevant, and so no criteria appealing to such considerations will be well motivated from a deliberative standpoint.

6.3. Objections

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 or not 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 permit the reasoning about E or not. 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 favors 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 (i.e. defer to doubt or not). Rather, Christensen and Vavova are trying to help us *recover* belief from certain perspectives of doubt.

Reply:

GIRP-not-NIRP does not actually offer a way to recover belief from 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. This would be something like:

BELIEF-WITHOUT-REASONS: When you have no reasons to believe either P or \sim 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 just 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 skeptical perspective, then the skeptical perspective isn't a trimming down of your perspective - it's all you've got. The issue, for the inductive skeptic, isn't a lack of independent reason - it's a lack of reason at all. That's why BELIEF-WITHOUT-REASONS might help a skeptic 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 skeptical perspectives. People might have all sorts of commitments that allow them to make all sorts of interesting moves from a variety of different

perspectives. I might think that seeing a black cat makes it likely that the sun will rise tomorrow. Nothing I've said here tells against this commitment. But GIRP-not-NIRP, at least as stated, is not a commitment that allows for recovery from doubt. For this reason, I think Christensen and Vavova are most charitably interpreted as offering us a way of ignoring the skeptic not of convincing her.

8. What about Rationality?

I haven't answered the following question: when is it *rational* to defer to the perspective of doubt? I don't know whether this is an answerable question and if it is, what the answer is. 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. For GIRP-not-NIRP is telling us to be sensitive to considerations which will seem irrelevant from the truth-seeking deliberative perspective, no matter which perspective we're deliberating from. 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 should cast suspicion on the principle. You may disagree with me on this front, and here is not the place to get into these metaepistemological 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.

9. Why Do We Defer to Doubt When We Do?

What the classical skeptical arguments teach us is that if we set aside too much, we won't find our way back: we'll end up in the skeptical abyss. Most of us however have learned to live with this fact. We know that we can't get our ordinary beliefs back 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 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 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."

What is emphasized less in Nagel is that if we take *small* step backs what we do doesn't look arbitrary at all. I step back from my commitment to grade papers today. Does that look arbitrary? Not at all. I value my students, my job, and following through on my promises. Similarly, for most of my beliefs: Usually, if I take a small step back, I can recover my belief. The mischief 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). 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. That 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 too – 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.⁹

To see this, imagine you're programming a robot that's going to explore Mars. Suppose 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 ship-shape. So even if the robot were to encounter some Martians who say to it: "you know, robots like you tend

⁹ See Pinillos (2019) for an extended discussion of such ideas.

to malfunction in our environment” and provide a track record of malfunctioning robots like this one, I’ll want my robot 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 I think that if the robot’s battery is running low, then it will do no better than chance at performing certain calculations. Then, 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 that I know all sorts of things about Earth that I want the robot to take into account when making comparisons about how things are on Earth versus Mars. If the robot sets aside all of its Earth beliefs it won’t be able to recover them from doubt. But I won’t want the robot to defer to this perspective. So 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. If the robot tried to come up with an epistemology that justified its dispositions it would most certainly fail. Its dispositions are simply a result of the varying degrees of confidence that I, the designer, have.

What’s the moral of this story? A very speculative proposal about our own tendencies is that, in some sense, we’re like this robot. We’ve been “programmed” to be sensitive to the possibility of certain sorts of errors (Did I reason about this particular matter correctly?) but not the possibility of other sorts of errors (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, in fact, 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 skeptics. 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 hodgepodge 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 skeptic, by appealing to the proposal that you've been programmed in ways that make you sensitive to errors you in fact are 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.

10. 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...(viii)”. 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 deliberately. Deference to doubt should be thought of as something that simply happens to us, without explanation, fluid and wordless.

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 cases. 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 SKEPTICISM. 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, unlike in SLEEPY, where the perspective of doubt is precise, in SKEPTICISM, 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 accurate¹⁰), imprecise perspectives don't recommend their imprecise credences.¹¹ More specifically, Schoenfield (forthcoming) argues 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 in fact we *can* 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.

The problem with the proposal is that I'm not convinced that it's going to do all the anti-skeptical work we might want it to. Some skeptical arguments are motivated by principles like the Principle of Indifference: that in the absence of reasons for treating different 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 skeptical 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 like how precise or imprecise your credence is in a skeptical perspective will depend a lot on the details of your skepticism. I don't see an argument for the claim that skeptical worries *always* go along with extremely imprecise probabilities.

¹⁰ I'm relying here on the thought that our accuracy measures are "strictly proper." 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.

¹¹ Seidenfeld et al. (2012), Mayo-Wilson and Wheeler (2016), Schoenfield (2017), Berger and Das (forthcoming), Builes et al ()

Bibliography

- Berger, D. and Das, N. (2020). "Accuracy and Credal Imprecision." *Noûs* 54 (3):666-703
- Builes, D., Horowitz, S. and Schoenfield, M. "Dilating and Contracting Arbitrarily." (2022). *Noûs* 2022 56(1): 3-20
- Briggs, R.A. and Pettigrew, R. (forthcoming). "An Accuracy-Dominance Argument for Conditionalization." *Noûs*.
- Christensen, D. (1994). "Conservatism in Epistemology." *Noûs* 28(1): 69-89.
- Christensen, D. (2010). "Higher Order Evidence." *Philosophy and Phenomenological Research* 81(1): 185-215.
- Christensen, D. (2011). "Disagreement, Question-Begging, and Epistemic Self-Criticism." *Philosopher's Imprint* 11(6): 1-21.
- Elga, A. (2007). "Reflection and Disagreement." *Noûs* 41 (3):478-502.
- Elga, A. (ms.). "Lucky to be Rational."
- Enoch, D. (2010). "The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope with It." *Philosophical Studies* 148(3):413-438.
- Greaves, H. and Wallace, D. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115(459): 607-632.
- Horowitz, S. (2014). "Epistemic Akrasia" *Noûs* 48(4): 718-744.
- Horowitz, S. and Sliwa, P. (2015). "Respecting all the Evidence." *Philosophical Studies* 17 (11):2835-2858
- James, W. (1896). "The Will to Believe" in S.M. Cahn (ed.), *The Will to Believe: And Other Essays in Popular Philosophy*. Longmans, Green, and Co.
- Lasonen-Aarnio, M. (2014). "Higher Order Evidence and Limits of Defeat." *Philosophy and Phenomenological Research* 88 (2):314-345.
- Mayo-Wilson, C. and Wheeler, G. (2016). "Scoring Imprecise Credences: A Mildly Immodest Proposal." *Philosophy and Phenomenological Research* 93(1): 55-78.
- Pinillos, N.A. (2019). "Skepticism and Evolution." in B. Kim & M. McGrath (eds.), *Pragmatic Encroachment in Epistemology*. Routledge

- Rinard, S. (2018). "Reasoning One's Way Out of Skepticism." *Brill Studies in Skepticism, Special Issue: The Mystery of Skepticism: New Explorations*
- Schoenfield, M. (2017). "The Accuracy and Rationality of Imprecise Credences." *Noûs* 51(4): 667-685.
- Schoenfield, M. (2018). "An Accuracy Based Approach to Higher Order Evidence." *Philosophy and Phenomenological Research* 96(3): 690-715
- Shanley, P. (2005). *Doubt*. Theatre Communications Group.
- Seidenfeld, T., Schervish, M.L, and Kadane, J.B. (2012). "Forecasting with Imprecise Probabilities." *International Journal of Approximate Reasoning* 53: 1248-61.
- Street, S. (2006). "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127 (1):109-166
- Titelbaum, M. (2015). "Rationality's Fixed Point (Or: In Defense of Right Reason)" in T. Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology* 5. Oxford University Press.
- Vavova, K. (2014). "Debunking Evolutionary Debunking" in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics* 9. Oxford University Press.
- Vavova, K. (2018). "Irrelevant Influences." *Philosophy and Phenomenological Research* 96(1): 134-152.
- White, R. (2009). "On Treating Oneself and Others as Thermometer" *Episteme* 6(3):233-250
- White, R. (2010). "You just believe that because..." *Philosophical Perspectives* 24(1):573-615.
- Weisberg, J. (2015). "Updating, Undermining and Independence." *British Journal of Philosophy of Science* 66(1): 121-159.\