

Meditations on Beliefs Formed Arbitrarily¹

For to say under such circumstances, “Do not decide, but leave the question open,” is itself a passional decision- just like deciding yes or no, and is attended with the same risk of losing the truth.

William James, “The Will to Believe”

Abstract: Had we grown up elsewhere or been educated differently, our view of the world would likely be radically different. What to make of this? This paper takes an accuracy-centered first-personal approach to the question of how to respond to the arbitrary nature in which many of our beliefs are formed. I show how considerations of accuracy motivate different responses to this sort of information depending on the type of attitude we take towards the belief in question upon subjecting the belief to doubt.

This paper is about how to respond to the realization that many of our beliefs are formed, in a sense, arbitrarily. Especially when it comes to matters that play a fundamental role in structuring our lives (religion, morality, politics), people adopt remarkably similar beliefs to their parents and peer groups.² This suggests that social influences are largely responsible for the fact that we hold the beliefs that we do. Had we grown up in a different city, or attended a different school, or been raised with a different religious outlook, we would almost certainly see the world very differently. The question is: what to do about this?

I will be addressing the concern about beliefs formed arbitrarily in a somewhat untraditional way. Rather than providing arguments about which way of responding to such etiological information is *rational* or how you *should* respond to this information, or whether you can have *knowledge* in such cases, I am going to

¹ For helpful discussion, and for reading and commenting on earlier drafts I am thankful to Brian Cutter, Tom Donaldson, Jane Friedman, Daniel Greco, Sophie Horowitz, Sarah Moss, Will Fleisher and especially Susanna Rinard. Thanks also to audiences at the PeRFECT Conference at University of Pennsylvania, the Ranch Metaphysics Workshop, the University of Pittsburgh conference on Formal Representations of Ignorance, Tel-Hai University, Australian National University, The University of Sydney and the Epistemic Rationality: Conceptions and Challenges Conference in Barcelona.

² Data from a Pew study on religious cross-generational retention rates as of 2007 can be found here: http://rationalwiki.org/wiki/Pew_Forum%27s_U.S._Religious_Landscape_Survey. See also Glass et al. (1986) and Argyle and Beit Hallahmi (2014), p.98.

simply describe how I've come to think about the problem in my own case.³ So before I begin, I'd like to say a few words about what motivated this choice, and why I think a piece of this form can be philosophically illuminating.

There are two reasons that I have chosen to use the first personal form in addressing the problem of arbitrarily formed belief: First, many people who regard their beliefs as arbitrarily formed (more on exactly what that means later) find themselves in a state in which they are *doubting* their beliefs. There are various thoughts that have been appealed to in the literature on beliefs formed arbitrarily that won't be of much help for someone in such a state. One example is the thought that *if you actually got things right* (in some sense or another of right), it can be rational for you to maintain your belief.⁴ It is not my purpose in this paper to argue against such views. My point is just that, for somebody engaged in a certain kind of doubt, these accounts won't be satisfying. This is because, in these contexts of doubt, one is wondering precisely about whether one got things right. And so it is, at very least, also worth thinking about this predicament from the perspective of somebody who is, oneself, in the state of doubt. One reason, then, that I am writing this piece in the first person is that my aim is to *demonstrate* how someone experiencing what is sometimes called "genealogical anxiety,"⁵ might navigate these concerns. (I take this to be at least part of what Descartes was doing in his *Meditations*: giving a demonstration of how someone who finds themselves beset with doubt might fish themselves out of the skeptical quicksand).

The second reason for using the first personal form is that I think theorists with different background epistemological views might wish to draw different conclusions from the considerations I raise here. So rather than take a stand on such large issues as internalism versus externalism or coherentism versus foundationalism, I will simply demonstrate a way of thinking about arbitrarily formed belief and let the theorist choose her own adventure on the basis of her other philosophical commitments. Along the way I'll point out some of these choice points.

So now, without further ado:

First Meditation: Why Avoid Beliefs Formed Arbitrarily?

³ The details, however, are not autobiographically accurate.

⁴ For discussion of views in this spirit see Lasonen Aarnio (2010, 2014), White (2010), Srinivasan (2015, §3.1), Titelbaum (2015), and Weatherson (ms.).

⁵ Srinivasan (2015).

I'll begin with two preliminary remarks.

First, my interest in beliefs formed arbitrarily isn't primarily with on/off belief states. I'm interested in any doxastic state in which we're more confident in one of P or $\sim P$, but we realize that this asymmetrical favoring of the proposition in question came about as a result of the sort of social influences described above. So in what follows I'll use the term "belief" in a very weak way so that an agent has a belief that P as long as she is more confident in P than in $\sim P$. This is purely terminological: it will allow me to discuss under the heading of "beliefs formed arbitrarily" not only cases of certainty, or binary belief, but also cases in which one has, say, a 0.6 credence in P , or a state in which one regards P as more likely than not.

Second, it will be helpful to be a bit more precise about what it is to regard a belief as formed arbitrarily. Here's how I'll think of things: To regard a belief as formed arbitrarily is to regard which belief one ends up adopting with respect to P as independent of whether P . (Formally, we can think of this as regarding $\Pr(I \text{ form the belief that } P|P) = \Pr(I \text{ form the belief that } P|\sim P)$, and $\Pr(I \text{ form the belief that } \sim P|P) = \Pr(I \text{ form the belief that } \sim P|\sim P)$.⁶

I'll illustrate this notion of arbitrarily formed belief by considering two toy cases inspired by White (2010). These cases are very artificial, but they'll be useful for getting some of the basic ideas on the table. (We'll get to the cases that initially concerned us – religious, moral, and political belief – in the Fourth Meditation, after some other warm-up cases).

⁶ A few notes about this definition: First, "Pr" refers to an agent's subjective probabilities. Second, the definition works most straightforwardly when thinking about cases in which I'm regarding some *future* belief of mine as one that will be arbitrarily formed (for instance, a case in which I know that I'll get some evidence later, but I don't know which belief I will form on the basis of the evidence, if I form one at all). Later, I'll talk about cases in which we're considering currently held beliefs and what's involved in regarding such beliefs as arbitrarily formed. Third, note that this is a definition of what it is *to regard* a belief as formed arbitrarily. At no point will I define what it *is* for a belief to be formed arbitrarily. One may be able to provide such a definition, but I'm primarily interested in what to think given an agent's perspective on things. So, for my purposes, it's enough to talk about what attitude the agent has that elicits the relevant concern. Finally, I'm using the term "regarding a belief as arbitrarily formed" stipulatively, to capture the sorts of cases that I'm interested in. There are many uses of the word "arbitrary" and one might think that some cases that meet my definition don't count, intuitively, as "arbitrarily" formed belief (for example, perhaps the beliefs are based on reasons and arguments). That's fine. My goal isn't to provide an analysis of our intuitive conception of arbitrariness, but rather to home in on cases in which we regard which opinion we form as independent of the truth as a result of learning about the belief's etiology.

Perceptual Coin Flip: One fair coin will determine whether the wall will be painted red or blue. Another fair coin will determine whether it will appear to me that the wall is red or it will appear to me that the wall is blue.⁷

If I thought that I'd find myself in Perceptual Coin Flip, and I expected to form a belief about the color of the wall on the basis of how things appear to me, then I'd regard my future belief as arbitrarily formed. This is because I'd think that the color of the wall, and my belief about its color, would be determined by two independent coin flips, and so I'd regard which belief I form as independent of the truth.

Logic Coin Flip: One fair coin will determine whether I'll be given a logic problem whose premises entail H or a logic problem whose premises entail $\sim H$. The flip of a second fair coin will determine whether I come up with a proof that seems to me to show that the premises entail H or I come up with a proof that seems to me to show that the premises entail $\sim H$. (Whichever answer I come up with, checking and double-checking will yield the same answer).

A similar line of reasoning applies to Logic Coin Flip. If I were to learn that I will find myself in such a situation in the future, and I expected to form beliefs on the basis of my reasoning, I would now regard my future belief state as arbitrarily formed. For I'll regard the facts about which belief I'll form as independent of whether I'm given an H-entailing problem or a $\sim H$ -entailing problem.

When I contemplate these toy cases I feel strongly that I'd much prefer maintaining a 0.5 credence in the relevant proposition to forming a belief arbitrarily. But why do I have this preference?

In trying to explain why I'm averse to forming beliefs arbitrarily in cases like the ones above, I started thinking about what it is, in general, that I'm after when I'm inquiring. And when I reflect on this question (things might go differently for you) the answer that comes to me is this: I'm trying to get at the truth. What I want out of my beliefs when I'm inquiring into some matter is that they provide me with an accurate representation of reality. There might be other goodies that would be nice to have: for example, it might be nice if my beliefs were not only true, but also

⁷ There are different ways to fill in the case, and the differences won't matter for my purposes. But note that it is perfectly consistent with the description of the case that if the color of the wall matches the color that it appears to me to be, then I have an ordinary veridical visual experience.

couldn't have easily been false (and so could constitute knowledge),⁸ or it might be nice if my beliefs contributed to my general well-being.⁹ However, I want to set these other lovely features of belief aside for the moment. I'm interested for now in whether an aversion to arbitrarily formed beliefs can be made sense of given what I take my most immediate goal to be: the truth.¹⁰

So can an aversion to arbitrarily formed belief be explained by a concern with truth? Answer: Yes, at least some of the time, but not in an obvious way. For note that, at first glance, it's not clear why, given a concern with truth, I'd be averse to forming a belief arbitrarily. It's true that if I expect to form a belief about the color of the wall in Perceptual Coin Flip, I'll think that I have a 50% chance of forming a false belief. That is, indeed, unfortunate. But, on the plus side, I'll also have a 50% shot at a true belief! If I adopt a 0.5 credence, on the other hand, I'm playing it safe – I'm not risking any falsehoods, but at the cost of not gaining any truths either. So why does 0.5 seem preferable? In the practical domain, I don't think that there's anything objectionable, given my concern with money, about taking a gamble that gives me a 50% shot at earning ten dollars and a 50% shot at losing ten dollars. I don't have a strong preference for maintaining my current monetary state. Given that I'm willing to take a monetary gamble, why am I so averse to a belief gamble?

What these considerations illustrate is that not any way of caring about the truth will vindicate an aversion to belief gambles. However, some ways of caring about the truth do vindicate such an aversion. What are these "different ways of caring about truth?" As William James long ago pointed out, there are many ways of valuing accuracy – many ways of trading off the value of truth against the disvalue of falsehood. Different ways of valuing accuracy can be encoded by different accuracy measures, sometimes called "scoring rules." An accuracy measure gives a numerical accuracy score to a credence in a proposition, given the proposition's truth value. So if a proposition is true, the higher the credence, the better the score, and if a proposition is false, the lower the credence, the better the score. While all scoring rules will agree on that much, they will differ with respect to *how much* better or

⁸ Friedman (ms.) assumes (but "mostly for expository convenience") that the goal of inquiry is knowledge. This is also suggested in Srinivasan's (2015) discussion of arbitrarily formed belief.

⁹ Rinard (ms.) defends a view according to which all reasons for belief are practical.

¹⁰ Despite the fact that my concern here is with truth, I think that what follows should still be interesting to those whose concern is, say, with knowledge, or rationality, rather than truth. For in many cases in which we're worried that our beliefs are not true, we're also worried that they don't constitute knowledge, or are not rational. So I'm going to stay focused on truth and accuracy, and you may draw your own connections between what I say and concerns about knowledge and rationality based on how you think concern with knowledge or rationality is related to concern with truth.

how much worse certain increases of decreases in credence will be. So, for example, if I'm more concerned about getting close to truths than I am at staying far from falsehood, a scoring rule that does good job at representing my concern with accuracy may assign a bigger accuracy boost to the move from 0.5 to 0.6 in a truth than to the move from 0.5 to 0.4 in a falsehood.

Now, our concern with the accuracy of our credences is not nearly precise enough to determine a unique scoring rule that represents the way we trade off the value of truth against the disvalue of falsehood. But I do think that there is good evidence for the claim that our concern with the accuracy of our credences has the feature that credences are *self-recommending*: for an agent with credence c in P and credence $1-c$ in $\sim P$, her own credences will have higher *expected* accuracy than any alternative.¹¹ Accuracy measures according to which credences are self-recommending in this way are sometimes called "strictly proper" or "immodest"¹² and I will argue that immodest ways of caring about accuracy do vindicate an aversion to belief gambling in the cases discussed above.

But before presenting the argument, why think that we care about accuracy immodestly? Two points: First, I am sympathetic to Joyce's (ms.) claim that we discover the particular shape that our concern for accuracy takes in part by looking at the ways of forming belief we endorse. As it turns out, many of the fundamental ways of forming belief we endorse¹³ would not be licensed by the aim of getting at the truth if our concern with truth were immodest. So one reason to think that we care about accuracy in immodest ways is that the claim that our concern for accuracy is immodest does an excellent job at explaining why, when we're aiming to get things right, we like to form beliefs in the particular ways that we do.

The second point I want to make in this regard is teleological: it makes sense, given the sorts of creatures that we are, that we'd care about accuracy in an immodest way. This is because of results in Schervish (1989), which have been elaborated upon by Gibbard (2008) and Levinstein (2017). These results show that belief-forming methods aimed at accuracy, *when accuracy is valued immodestly*, are exactly what we'd hope for given the prominent role that our beliefs play in guiding action. The rough idea behind these results is this: because we don't know which

¹¹ The expected accuracy of c is just the average of the accuracy scores c might get in different worlds, weighted by the probability the agent assigns to those worlds obtaining.

¹² See, e.g. Oddie (1997), Greaves and Wallace (2006), Gibbard (2008), Joyce (2009), Horowitz (2013) and Pettigrew (2016) for discussion of immodesty.

¹³ For example, being coherent, updating by conditionalization, conforming one's credences to the chances when they're known, and, as I will show in a moment, avoidance of certain belief gambles.

choices our future selves will face, if we want our future selves to make good decisions, the best thing we can do in the absence of additional evidence is “give” our future selves our actual credences. So, for the purpose of guiding action, valuing accuracy in a way that makes credal states self-recommending (in other words: immodestly) is exactly what we’d want.

Here’s an illustration (by no means a proof) of the main idea: Suppose I currently have a 0.5 credence that there’s a post office half a mile away. (Perhaps I know that there was one there a month ago, but I think it may have closed). There are many possible reasons it might matter practically to me whether this post office exists. One possibility is that I discover a job that I want to apply to at 4pm this afternoon whose deadline is tomorrow. In that case, I’d need to get to a post office before 5pm, when the post offices close. (The job is at an old-fashioned institution that requires mail-in applications). Given that now I’m only 0.5 confident that there’s a post office half a mile away, I wouldn’t want my future self in these circumstances to take a stroll to the possible post office on the assumption that it’s still there. In such a case, I’d much prefer that my future self drive to some further post office that is definitely open, than take a chance on the one that might be half a mile away. On the other hand, if my future self wants to mail a wedding gift for a wedding that’s three weeks away, and it’s a beautiful afternoon, I wouldn’t recommend against walking half a mile east and scoping things out. Worst case scenario, I mail the gift on some later date. These are just two examples, but there are countless situations my future self might face, and which action I’d want my future self to take will depend on the details. Given that how I want my future self to act is a function of what my credences are, the best thing I can do for my future self so that she’ll make good decisions (again, absent getting new evidence), is give her my actual credences. So, instead of gambling on what my credences will be, I’ll want to keep the credences I have, and let my future self do the gambling on which *actions* to perform.

But wait – didn’t I start out assuming that my goal was an accurate portrayal of the world and not an efficient arrival at the post office or a successful job application? I did. But as I mentioned earlier, there are many ways to care about accuracy: many ways to trade off the value of truth against the disvalue of falsehood. Given the role that our opinions play in governing action, it makes sense that the particular *way* in which we care about accuracy is immodest. This is not inconsistent with the idea that in an inquiry in which our sole concern is with accuracy, we are motivated to form beliefs in ways that are licensed by an immodest concern with

accuracy. (Analogy: perhaps we came to find sweet things delicious because sugar is high in calories. Still, sometimes all we care about is a thing's deliciousness, and in those cases we can favor sweet things on purely deliciousness grounds).

In sum: there are two reasons to think that our concern with the accuracy of our credences is of the immodest variety: first, the claim that we're concerned immodestly provides a good explanation of why we endorse the belief forming methods that we do when we're inquiring, and second, given the role that beliefs play in guiding action, it would make sense that we'd come to value accuracy immodestly.

Let me now explain why caring about accuracy immodestly can explain our aversion to belief gambles of the sort described above: if I assign a 0.5 credence to a proposition, and I value accuracy immodestly, then I'll prefer to be at 0.5 than to be anywhere else. So I'll prefer to be at 0.5 than to be at, say, 0.8 or 0.2. But if I don't want 0.8, and I don't want 0.2, I'm also not going to want to go through a procedure that gives me a 50% shot at arriving at 0.8 (one thing I don't want) and a 50% shot at arriving at 0.2 (another thing I don't want) in a way that I regard as independent of the truth.¹⁴ If I expect to form a belief arbitrarily, say by forming a perceptual belief in Perceptual Coin Flip, then I'll regard the process of belief formation as involving a procedure which gives me a 50% shot at a higher credence, and a 50% shot at a lower credence in a way that I regard as independent of the actual color of the wall (this follows from the fact that I expect the belief to be arbitrarily formed). This is exactly the sort of procedure that an immodest way of caring about accuracy will recommend against. If I care about accuracy immodestly, I'll prefer sticking to 0.5 to undergoing a procedure of this sort. So if we care about the accuracy of our credences immodestly, we have an explanation as to why we don't like taking belief gambles.

In sum: My aversion to arbitrarily formed belief in the toy cases can be explained by my concern with accuracy, but *only* if my concern for accuracy is immodest. Non-immodest ways of caring about accuracy will license shifts from one credence to another (even in the absence of new evidence), and, as a result, they will

¹⁴ Immodesty is consistent with the idea that I'm happy to revise my credences if I think that the way in which I'll revise them is correlated with the truth. See Schoenfield (forthcoming) for a more detailed argument explicating why immodesty prohibits belief gambles. See also Carter (forthcoming) and Eder (ms.) for a defense of the claim that the way in which we trade off the value of truth against the disvalue of falsehood favors the avoidance of falsehood over the gaining of truth.

also license certain belief gambles.¹⁵ I gave some reasons for thinking that my concern with accuracy is, in fact, immodest and so there is indeed an accuracy-based vindication for my desire to avoid beliefs formed arbitrarily in such cases.¹⁶

Bottom Line: Assuming that my concern with accuracy is immodest, there is an accuracy based vindication of my aversion to forming beliefs arbitrarily in cases like Perceptual and Logic Coin Flip.

Second Meditation: Graduate School

Big news: I've decided to pursue a PhD in neuroscience! I studied neuroscience when I was in college and I remember that around the time that I graduated there was a lively debate going on about whether olfactory information was encoded by the spatial arrangement of the neurons that fire, or in some other way (such as the temporal sequence of firing). In preparation for graduate school, I've been reading through some recent articles on the topic. But it's so complicated! I really have no idea what to think about the issue.

I had lunch with my neuroscience professor from college earlier today, Professor Katz, and I was asking him for advice about which school to attend. I've been considering two programs: Columbia and University of Arizona. He remembered my interest in the debate about olfactory coding and he said: "Well, I can tell you right now, if you go to Columbia, next time I see you you'll be favoring the spatial view, and if you end up at Arizona, you'll think that the spatial view is probably wrong. That's how things work in graduate school: everybody reads the same articles and journals but what you end up thinking about the matter depends on which social influences you are subject to."¹⁷

¹⁵ For example, on what's called the "absolute value score," a belief gamble which gives me a 50% shot at ending up at 0.8 and a 50% shot at ending up at 0.2 will look fine from the perspective in which I have a 0.5 credence.

¹⁶ For those interested in thinking about rationality, the considerations in this section could have been presented as claims about the rationality of having certain belief-forming preferences. Although epistemologists rarely talk about the rationality of belief-forming preferences, here is how such an argument would go if one were to make one: the reason that it's rationally permitted/required to have a preference for maintaining a 0.5 credence over taking a belief gamble is that we are rationally permitted/required to care about accuracy in an immodest way and immodest ways of caring about accuracy recommend maintaining 0.5 over taking a belief gamble. Thus, at least in cases in which all that one wants out of one's future opinions is that they be accurate (and in which this is a rationally permissible/required attitude to take), it's rationally permissible/required to prefer a 0.5 credence to a belief-gamble. Later in the paper I'll focus on beliefs themselves, rather than belief-forming preferences.

¹⁷ This case is inspired by G.A. Cohen (2000), and by going to graduate school.

“Actually”, I told Professor Katz, “I think that when I get to graduate school I won’t form any opinion on the matter at all given what you’ve just told me. You see, I think that forming an opinion once I get to graduate school amounts to a belief gamble, and I don’t like gambling on my beliefs.” “Well, we’ll see,” he said, and chuckled in a way that seemed mildly condescending.

But this evening, as I’ve been pondering the matter further, I started rethinking my commitment to agnosticism. This thought occurred to me when I was reflecting on which of S (the spatial view) or $\sim S$ (its negation) I *currently* think is more likely to be true. When I was reading through these neuroscience papers over the past few days, I found myself moving back and forth between which I thought was more likely, and when I sit back now and think through all of the evidence I’ve collected – well, I really have no idea. I wouldn’t say that I regard S as more likely than $\sim S$, and I also wouldn’t say that I regard $\sim S$ as more likely than S . But I also don’t have a 0.5 credence in S . One way to see that my attitude towards S is different from a precise 0.5 credence is to note that getting a teeny bit of evidence in favor of S (e.g. learning that one of the studies I read favoring S had a slightly larger sample size than I’d thought) wouldn’t make me more confident in S than $\sim S$. In contrast, when I have a precise 0.5 credence in a proposition, *any* evidence in favor of that proposition will break the tie (learning that the coin is weighted 0.5000001 towards Heads, rather than being fair, will make me more confident in Heads than Tails).

Why does it matter whether my credence is 0.5 or not? The reason it matters is that earlier I described some reasons for thinking that if I have a credence in a proposition, then I won’t want to take a belief gamble. This followed from the fact that, given the way I care about accuracy, credences are self-recommending. But if my attitude towards S can’t be represented by a credence, then the considerations I appealed to above don’t, at least in a straightforward way, provide accuracy based motivations for maintaining my current state over taking a belief gamble. So I started wondering: are there any accuracy based grounds for avoiding a belief gamble of the sort I’d be subject to by going to graduate school given my actual attitude towards S ? After some contemplation, I concluded that’s it’s very hard to see what sorts of accuracy based grounds there might be for avoiding such a gamble. In fact, I’m not convinced that there are any.

To explain why it’s difficult to provide an accuracy based motivation for avoiding belief gambles in cases like the one above, it will be helpful to get clearer

on the nature of my attitude towards olfactory coding in this case. I'm going to use the term "lacking an opinion about P" as follows:

S lacks an opinion about P if it's not the case that *S* is more confident in *P* than in $\sim P$, it's not the case that *S* is more confident in $\sim P$ than in *P*, and it's not the case that *S* has a precise 0.5 credence in *P*.¹⁸

An agent who lacks an opinion about *P* cannot be represented by a precise credence function. But some people think that such agents can be represented by a set of credence functions, called "a representor."¹⁹ On this picture, if we want to describe an agent's level of confidence towards a particular proposition *P* that she lacks an opinion about, rather than representing that attitude by a single number that represents the agent's confidence in the proposition, we can represent the agent's confidence level by an interval, like, for example, [0.1-0.9].

There are many unanswered questions about these "imprecise" or "mushy" credences and now is not the time to delve into the details. But since I think it's important to have in mind some psychological interpretation of this formalism, I'd like to offer what I take to be a promising way of thinking about what it is for an agent to be such that credence *c* is in an interval that represents her confidence-level towards *P* (I'll call such an interval a "P-representor"). It's worth noting, though, that nothing essential in what follows rests on this psychological interpretation of imprecise credences. If you have your own favorite interpretation you can use that one.

Here's how I'll understand the formalism. I'll say:

Credence *c* is a member of *S*'s P-representor if both of the following conditions are met:

- (a) It's not the case that *S* is more than *c*-confident that *P*.
- (b) It's not the case *S* is less than *c*-confident that *P*.²⁰

¹⁸ I intend the locution: "it's not the case that *S* is more confident in *P* than in $\sim P$ " to be consistent with it being indeterminate whether *S* is more confident in *P* than in $\sim P$. So the sentence "it's not the case that *S* is more confident in *P* than in $\sim P$ " could be restated as: "it's not the case that, *determinately*, *S* is more confident in *P* than in $\sim P$."

¹⁹ For instance, Kyburg (1961) Levi (1974), Jeffrey (1983), van Fraassen (1990) and Joyce (2005, 2010).

²⁰ Note that there are plausibly cases in which it is indeterminate whether *c* is a member of *S*'s P-representor. Indeed, I am sympathetic with Rinard's (2017) claim that, in many cases, there is no maximally specific and fully accurate description of an agent's confidence level. Still, we can talk

Since I'm reflecting on my own attitudes in this case, it's worth mentioning how I reflect on the question of whether some credence c is in my P -representor.²¹ First, I note that I assign credence c to a c -weighted coin landing Heads. Next, I imagine someone presenting me with a c -weighted coin and asking: "what are you more confident in: that this coin will land Heads, or P ?" Suppose it's not the case that I'd answer: "I'm more confident that the coin will land Heads than I am in P " and it's not the case that I'd answer: "I'm more confident in P than that the coin will land Heads." Perhaps I'd say: "I'm equally confident in both", or perhaps I'd shrug my shoulders, or say: "I don't know" or "I'm not sure" or maybe there is simply no fact of the matter about what I would say if asked this question. As long as I think that it's not the case that I'd answer: " P is more likely" and it's not the case that I'd answer: "Heads is more likely", I'll think that c is in my P -representor. If c is the only credence with this feature, then I'll think that I have a precise credence of c in P , since c will be the only element in the P -representor. But if there is more than one c with this feature, I'll judge my credence to be imprecise.²²

Back to my contemplations about graduate school: I find myself, prior to going to graduate school, in a state in which I lack an opinion about S : it's not the case that I'm more confident in S than in $\sim S$, it's not the case that I'm more confident in $\sim S$ than in S , and it's not the case that I have a 0.5 credence in S (so there is more than one member in my S -representor). Let's call my state \mathcal{L} (for "lacking an opinion"). The question is: are there accuracy based motivations for maintaining \mathcal{L} once I go to graduate school as opposed to letting my opinions be swayed by the influences around me?

If \mathcal{L} were a self-recommending state, then we'd have an argument for trying to maintain \mathcal{L} : if \mathcal{L} is a state that recommends itself (from an accuracy perspective) over every other state, it will also recommend itself over a gamble between two states that it disprefers. But a combination of results in Seidenfeld et al. (2012), Mayo-Wilson and Wheeler (2016) and Schoenfield (2017) show that, given some plausible constraints on the way in which we value accuracy, there is no accuracy

about some set of credences as being members of S 's P -representor so long as every member of the set in question, c , is such that it's not the case that (determinately) the agent is more than c -confident that P and it's not the case that (determinately) the agent is less than c -confident that P .

²¹ This is not meant to imply that we're always able to tell, for every credence, whether or not it is in our representor.

²² See Fishburn (1986) for a lovely representation theorem that delivers a unique set of credences on the basis of comparative confidence levels.

measure that has the feature that all imprecise credal states are self-recommending. I won't summarize these results here. Instead, I want to argue for something more specific: that \mathcal{L} doesn't recommend itself over every state in which I'm more confident in one of S or $\sim S$. In other words: \mathcal{L} doesn't recommend against every opinionated state. I'll argue for this by arguing for:

(*) If I lack an opinion about P , and c is a number in my P -representor that is not equal to 0.5, then it's not the case that I'm in a state that recommends itself over having a credence of c in P .

If c is a number in my P -representor that is not equal to 0.5, adopting credence c in P amounts to becoming more confident in one of P or $\sim P$. Thus, if I can show that my state \mathcal{L} doesn't recommend itself over having credence c , I'll have shown that it doesn't recommend against every opinionated state.

The argument I'll provide for (*) is an argument by elimination: I'll consider a number of different ways one might try to motivate a preference for \mathcal{L} over c when one is in \mathcal{L} , on the basis of accuracy considerations, and show that none of them succeed. This strategy has the weakness that I can't claim to have exhausted all of the possible accuracy based motivations for maintaining \mathcal{L} . But I will have shown (a) that the accuracy based motivations for avoiding belief gambles in the case of credences don't motivate avoiding belief gambles in cases in which I'm in a state of lacking an opinion, and (b) there is, at very least, no straightforward reason for preferring \mathcal{L} to c on accuracy based grounds. If there are accuracy based reasons for preferring \mathcal{L} to c , they are not the sorts of reasons that are based in a familiar decision theory.

There are three assumptions that I'll make in the course of arguing for (*) that are worth flagging. The first is that we value the accuracy of precise credal states in an immodest way. I make this assumption because, as I mentioned earlier, I think that our concern with the accuracy of credences does have this feature, and also because, if we weren't concerned about the accuracy of credences in an immodest way, there would be little hope of motivating an aversion to belief gambles even in cases in which we have precise credences, let alone cases of lacking an opinion. But I'll be assuming a rather weak form of immodesty when it comes to comparisons between c and \mathcal{L} . I won't assume that c must recommend itself *over* \mathcal{L} .

I'll just assume that c recommends itself over any other sharp credence, and that accuracy considerations don't *require* a move from c to \mathcal{L} .²³

The second assumption I'll make for the purposes of this argument is that \mathcal{L} isn't an accuracy self-undermining state: it doesn't, in every case, recommend against itself. One reason for this assumption is that if \mathcal{L} were always self-undermining, an agent interested in accuracy would never enter state \mathcal{L} to begin with, and so figuring out what \mathcal{L} recommends becomes a much less interesting project.

The final assumption I'll make is that \mathcal{L} is a state that it makes sense to evaluate for accuracy. The reason for this assumption is that if \mathcal{L} is not evaluable for accuracy, then there is *definitely* no accuracy based motivation for preferring \mathcal{L} to c . Thus, if there is any hope of motivating a preference for \mathcal{L} over c on the basis of accuracy considerations, \mathcal{L} must be the kind of state whose accuracy it makes sense to evaluate.

Here's how I'll proceed with the argument for (*): First, I'll argue that one can't motivate a preference for \mathcal{L} over c by claiming that \mathcal{L} is more accurate than c no matter how the world is. Second, I'll argue that \mathcal{L} can't be favored over c on the basis of thinking that probably \mathcal{L} is more accurate than c . Third, I'll argue that one can't prefer \mathcal{L} over c on the basis of expected accuracy, or on the basis of what I'll call "generalized expected accuracy." And finally, I'll argue that one can't prefer \mathcal{L} over c on the basis of other familiar decision rules like Minimax, Maximin or Hurwicz criteria more generally.

To start, note that \mathcal{L} can't be more accurate than c in every world. For if \mathcal{L} is more accurate than c no matter what, then accuracy considerations would tell us that, no matter what our current opinion is, we should never have credence c . But since we're assuming that credences are self-recommending (we're maintaining immodesty for credences), it must be the case that credence c doesn't accuracy-undermine itself.

Can an agent in \mathcal{L} prefer \mathcal{L} to c on the basis of thinking that *probably* \mathcal{L} will be more accurate than c ? No, for the accuracy of \mathcal{L} and c depend only on the truth of the proposition in question: call it P . If you were in \mathcal{L} and thought that \mathcal{L} was probably more accurate than c , then you'd have to think that, in either the P world, or the $\sim P$ world (but not both), \mathcal{L} is more accurate than c . Without loss of generality,

²³ Konek (forthcoming)'s accuracy-based argument in favor of states like \mathcal{L} violates this immodesty condition on credences. See Schoenfield (2017) note 14 for discussion.

suppose you think \mathcal{L} is more accurate than c if P is true, but not if P is false. In that case, thinking that \mathcal{L} is probably more accurate than c amounts to thinking that P is more likely than $\sim P$ (since \mathcal{L} is more accurate than c if and only if P is true). However, by stipulation, it's not the case that you regard P as more likely than $\sim P$.

But let's not give up too quickly. We know from decision theory that there are cases in which one doesn't think Option A is more likely to bring about a better outcome than Option B, but one still ought to choose Option A: these are cases in which Option A has greater *expected value* than Option B. Is it possible then, that, although it's not the case that an agent with \mathcal{L} thinks \mathcal{L} is likely to be more accurate than c , that she can assign \mathcal{L} greater *expected* accuracy than c ? Not straightforwardly. Since expected accuracy judgments are always relativized to a credence function, and our agent with \mathcal{L} lacks a credence in P , the notion of "expected accuracy" is simply not defined for an agent with \mathcal{L} .

Is there some way to generalize the notion of expected accuracy so that we can sensibly talk about the expected accuracy judgments of an agent in \mathcal{L} ? If we follow the kind of supervaluationist approach that has been prominent in the literature on imprecise credences²⁴ we can say something like this: If, for every credence function c in an agent's representor, c assigns greater expected accuracy to b_1 than to b_2 , then the agent can be said to assign greater expected accuracy to b_1 than to b_2 . Still, this way of proceeding won't motivate a preference for \mathcal{L} over c . By stipulation, c is a credence in the agent's P -representor. This means that some credence function in the agent's representor, call it c , assigns c to P . Since credences are self-recommending, it won't be the case that every credence function in the representor assigns greater expected accuracy to \mathcal{L} than to c , for this would require that c assigns greater expected accuracy to \mathcal{L} than to c , and, if this were so, c wouldn't be self-recommending. Thus, this generalization of the notion of expected accuracy won't yield the result that an agent with \mathcal{L} assigns greater expected accuracy to \mathcal{L} than to c .²⁵

²⁴ See, e.g. van Fraassen (1990, 2005, 2006), Hajek (2003), Joyce (2005, 2010) and Rinard (2015).

²⁵ Another expectation based decision rule for imprecise probabilities (more well known in the economics literature) is the "GS decision theory" (Gilboa and Schmeidler (1986)). In order to determine whether this rule could issue a recommendation for \mathcal{L} over c we need to assign an accuracy profile to \mathcal{L} . I can think of three principled ways of doing this: We can let \mathcal{L} have the same accuracy profile as the midpoint of the range of credences for P , we can let \mathcal{L} have the average of the accuracy scores of all the points in the range, or we can let \mathcal{L} have a score that is itself a range – plausibly corresponding to the accuracy scores of the credences in the P -representor. The first two interpretations yield the result that \mathcal{L} is always self-undermining, which conflicts with one of the assumptions I'm making for the purposes of this argument. On the third interpretation (which to my

What about other decision rules? Since we want to maintain immodesty for credences, we need to consider whether any rules that make credences self-recommending yield a preference for \mathcal{L} over c when one is in \mathcal{L} . But it's not clear that there are plausible decision rules, other than expectation related ones, that can yield the result that credences are self-recommending. For note that other familiar decision rules like Maximin, Minimax, and Hurwicz rules don't take an agent's doxastic state into account when issuing a recommendation. But any rule that doesn't take the agent's doxastic state into account won't make credences self-recommending. Why? Because for credences to be self-recommending, what's recommended for an agent with a 0.6 credence must be different from what's recommended for an agent with a 0.5 credence. If, however, what's recommended doesn't depend on the agent's credences, this won't be the case.

In sum, it's hard to see what there is about the state of lacking an opinion which would privilege itself, from an accuracy perspective, over every state in which I have an opinion. Since I currently lack an opinion about how olfactory information is encoded, I don't think that I'm in a state that recommends itself over every state in which I am more confident in S than $\sim S$. A similar argument would show that I'm not in a state that recommends itself over one in which I'm more confident in $\sim S$ than S . Having reflected on this, I find myself much less averse to taking a belief-gamble: letting myself become opinionated as a result of the school that I choose to attend.²⁶

Bottom Line: It's difficult to find an accuracy-based motivation for maintaining my state of lack of opinion over taking a belief gamble: allowing my opinions to be formed by whichever graduate school I choose to attend.

mind is the most promising), the accuracy of \mathcal{L} is represented by a set of numbers, and so plausibly it will be indeterminate which of \mathcal{L} or c is more accurate, no matter how the world turns out to be. Thus, \mathcal{L} won't recommend itself over c . An interesting question: if we regard the comparative accuracy of \mathcal{L} and c to be indeterminate no matter how the world is, will c accuracy-recommend itself over \mathcal{L} ? I suspect that it won't, but there are some subtle issues here that deserve further investigation.

²⁶ Once again, these arguments could be reformulated as claims about the rationality of belief-forming preferences. Here's how such an argument would go: it's not the case that if one's aim is accuracy, and one is in \mathcal{L} , there is a rational requirement to prefer maintaining \mathcal{L} over becoming opinionated in the graduate school case. Why? Because it is rationally permissible for one's belief forming preferences to be determined by accuracy considerations, it is rationally permissible to be in \mathcal{L} , and it's not the case that, for an agent in \mathcal{L} , there are accuracy based reasons for preferring \mathcal{L} to every opinionated state. (I'm not defending these claims about rationality here. I'm merely describing which premises concerning rationality would need to be accepted for the considerations here to be turned into such an argument).

Third Meditation: Higher Order Evidence and the Perspective of Doubt

All of this meditating has been taxing, and yesterday my friend Jane suggested that we go out for a drink. “I really shouldn’t,” I said, “I have to finish an answer key for my logic class.” But Jane can be very convincing, and before I knew it I was at the bar, sipping Merlot, as my concerns about beliefs formed arbitrarily melted away. When I arrived home, I was tired and inebriated, but I quickly got to work. I had just finished what seemed to me a very satisfying proof that the set of premises given by the problem entailed H when my spouse popped in and said: “Please don’t tell me you’re doing logic problems. You know what happens when you do logic problems in this state. Your answers are complete nonsense! Remember last time? You checked in the morning and only half of your answers were correct!”²⁷

I started to get worried. Did those premises actually entail H? At first I cheered myself with the thought that I could just double or triple check my answers, but then I remembered that, last time, when I was doing logic problems while tired and drunk, I did just that and *still* only half of the problems were correctly answered.

It occurred to me that I am currently in a state that is in some respects similar to that of the hypothetical subject I had imagined in Logic Coin Flip. When I’m drunk, and am reasoning about these logic problems in a way that’s no better than chance, the answers I get are only 50% likely to be correct. Looking back through my notes, I remembered that I had concluded that it’s better to adopt a 0.5 credence than to form a belief that’s only 50% likely to be true. Indeed, I planned that if I ever find myself in a situation like this one I’ll adopt a 0.5 credence. But now I find myself with the belief that the premises entail H, and it is only after having formed this belief, that I realized what kind of situation I’m in. If the accuracy-based motivations for *avoiding* forming a belief are to motivate *abandoning* a belief that I already formed I must *now* think that the belief is only 50% likely to be true. But is that what I think? It’s not so clear. If I were thinking about this matter from a perspective that includes all of the beliefs that I formed, then I don’t think the belief is only 50% likely to be correct. For I formed the belief that the premises entail H. In fact, I was certain or nearly certain that the premises entail H. This means that the perspective that includes the belief that I formed is one in which it’s certain, or

²⁷ This case is inspired by Christensen’s (2010, p.187) “Drugs” case.

nearly certain, that the belief I formed is correct (since the belief is correct if and only if the premises entail H). If I think it's certain or nearly certain that the belief I formed is correct, then I *don't* think the belief is only 50% likely to be correct. So how do the considerations I raised prior to being in such a situation carry over to the case in which I now am in that situation?

Here's what I realized: it's true that, from the perspective in which I am certain or nearly certain that the premises entail H (let's call this proposition "EH"), I'll think that my belief is highly likely to be true. But when I started wondering: "should I give up my belief that EH?" upon being reminded of my track record, I wasn't asking this question from a perspective that takes my belief that EH for granted.²⁸ Why? In general, if I have some belief, and I start wondering whether to give it up, then I'm engaged in *doubt*. When I doubt a belief that I currently have, I am considering whether to give up that belief, but I am considering whether to give it up from a perspective that doesn't take the belief in question for granted. After all, if I were taking it for granted, then it would be obvious, assuming my goal is accuracy, that I wouldn't want to give it up. Why would I want to give up a true belief? (In credal talk: if I have a high credence, I will regard it as more expectedly accurate than a middling credence, so why would I want to give it up?)²⁹

The perspective of doubt that I occupy in this case is also one in which I'm not willing to take for granted that the inferences I made in deriving H from the premises are good ones. After all, if I took the inferences I made in deriving H from the premises for granted, then it would also be clear that I wouldn't want to give up

²⁸ When I say that I take some proposition P for granted I mean that I'm willing to deliberate on the basis of my belief that P. Because I'm including credal states favoring P as beliefs in P, it's worth pointing out that when I say that an agent is taking P for granted, this should not be taken to mean that the agent is ignoring all possibilities in which P is false. It merely means that whatever her asymmetrical attitude favoring P is, she is willing to reason with it.

²⁹ The dogmatism paradox raises the question of why we don't dismiss or avoid evidence that disconfirms our beliefs. Why not think: P is true, so any disconfirming evidence must be misleading? This is an interesting puzzle, but not the one that I'm concerned with here. First, the dogmatist reasoning doesn't apply straightforwardly in cases in which we're less than certain that P (if I'm 0.6 in P, I can't reasonably assert "any evidence against P must be misleading"), but it is compatible with the cases I'm focusing on here that the agent is less than certain in the proposition in question. Second, the dogmatism paradox concerns cases in which one gets evidence that disconfirms P, but the cases I'll be focusing on are cases in which we subject a belief to doubt either in the absence of new evidence, or, if there is new evidence, it's such that the prior probability of P conditional on that evidence is the same as the prior probability of P. The reason for this focus is that reduction of confidence in higher order evidence cases, of the sort described here, can't be accommodated by ordinary conditionalization (Christensen (2010) p.200, Schoenfield (forthcoming)). However, as I'll argue, we can explain a reduction of confidence in such cases by appealing to the fact that the beliefs become subject to doubt.

my belief: for if all of the inferences are good ones, then EH must be true! So the fact that I'm wondering whether to give up my belief and wondering whether to rely on these inferences tells me that the perspective from which I want to deliberate isn't one that takes the belief and associated inferences for granted.

There are many interesting questions about what's going on when we doubt beliefs or inferences and I won't delve into a discussion of the psychology of doubt here. But I do want to address two questions that might arise about what sort of perspective I have in mind when I talk about "the perspective of doubt."

First question: if we believe P, and then subject this belief to doubt, is P the *only* belief that we aren't taking for granted in the deliberation? What about beliefs like "P or $2+2=5$ "? Answer: there is no univocal answer to the question "what are we setting aside when we doubt our belief that P?" A perspective of doubt is one in which we're not willing to rely on certain kinds of reasoning that we are willing to rely on when we're not engaged in doubt. For example, if we're doubting P, we're not willing to engage in a pattern of reasoning like this: "Since P is true, and P entails Q, Q is true." Almost always, when we're doubting P, we're also not willing to engage in a pattern of reasoning like this: "Since P or $2+2=5$ is true and $2+2$ doesn't equal five, then P is true." It's not that it's impossible to engage in a deliberation that doesn't take it for granted that P, but does take it for granted that either P or $2+2=5$. It's just that, generally, when we doubt our belief that P, the perspective we wish to be deliberating from is one that *also* won't rely on a number of other closely related beliefs. Which beliefs exactly are the ones that we don't take for granted when we doubt that P? It depends. Doubting that P doesn't *entail* some particular set of beliefs that one is not willing to take for granted. In any given case in which an agent is doubting, there *just are* the beliefs that she's willing to rely on and the ones that she isn't.

Second question: How confident is one in P in the perspective of doubt? The perspectives of doubt that I'm particularly interested in are ones in which, relative to the perspective of doubt, it's not the case that one is more confident in P than in $\sim P$ and it's not the case that one is more confident in $\sim P$ than in P. This is because I'm interested in doubt that's elicited when we worry about what caused us to asymmetrically favor P over $\sim P$ or $\sim P$ over P, so the perspective of doubt is one that removes this asymmetric favoring. But, as I mentioned earlier, there are many doxastic attitudes one might take towards P that are consistent with it not being the case that one is more confident in P than in $\sim P$ and it's not being the case that one is more confident in $\sim P$ than in P. So, here too, there is no univocal answer to the

question “what attitude does one take towards P once one sets aside one’s belief that P?” It might be that when we set aside our asymmetrical favoring of P over \sim P we find ourselves in a perspective in which our credence in P is 0.5. Or it might be that when we set aside our asymmetrical favoring of P over \sim P we find ourselves in a perspective that is represented by some wide interval of credences.

With these preliminaries about doubt on the table, I now want to think about how things look from the perspective in which I doubt my belief that the premises entail H. In this case, my degree of confidence in EH once I start doubting EH is 0.5. (This is because, in the book that I’m working with, half of the answers have premises that entail H and half have premises that entail \sim H). So I’m now interested in thinking about how to proceed from the perspective in which I’m doubting EH, and in which my degree of confidence in EH, in the perspective of doubt, is 0.5. Can I recover my belief that EH from the perspective in which I doubt it?

There are two primary ways that we recover beliefs that we doubt. The most straightforward way involves engaging in first order reasoning. For example, suppose, having left my apartment, I begin to doubt my belief that I turned off the stove. I may be able to recover that belief by engaging in reasoning like this: “I remember cleaning the stovetop before leaving the apartment, I would have noticed if the stove were on when I was cleaning it, at which point I would have turned it off. So I must have turned off the stove.” But this way of recovering a belief from doubt doesn’t always work: sometimes, once we set aside what is in doubt, we don’t have the resources left to recover the belief in this way. In these cases, we can sometimes appeal to higher order considerations. For example, consider my belief that my grandmother grew up in Massachusetts. Suppose that I subject this belief to doubt. I don’t have many other beliefs about my grandmother’s upbringing and so I can’t find any biographical information in my stock of remaining beliefs from which I could infer that she grew up in Massachusetts. I also don’t remember an occasion on which I was told that she grew up in Massachusetts. I just find myself believing it. Still, I can recover my belief that my grandmother grew up in Massachusetts in the perspective of doubt by thinking: “I find myself with the belief that my grandmother grew up in Massachusetts. The best explanation for why I have this belief is that a family member told me that she grew up in Massachusetts, and if a family member told me that she grew up in Massachusetts, she probably did grow up in Massachusetts.” In this way of recovering a belief from doubt, I use the very fact that I formed a certain belief as evidence for its truth.

Unfortunately, however, when I doubt my belief that EH, I can't recover my belief in either of these ways. I can't recover the belief using first order reasoning because the perspective I'm occupying when I doubt my belief that EH in this case doesn't license the very reasoning that I would need to derive the answer – my reasoning about this logic problem is *itself* part of what I am doubting. I also can't recover the belief in the higher order way because, given what I know about the circumstances, I don't take the fact that I formed the belief that EH as evidence for its truth.

Given that I can't inferentially recover my belief that EH in this perspective of doubt, and that, in this perspective, my credence in EH is 0.5, if I deliberate about whether to give up my belief that EH from the perspective of doubt, the answer will be yes. This is because, from this perspective, maintaining the belief that I formed will look like a belief gamble: it will involve a 50% chance of having a true belief and a 50% chance of having a false belief.³⁰ Since, when I'm at 0.5, I prefer a 0.5 credence to a belief gamble, the perspective of doubt will recommend that I abandon the belief that I formed and adopt a 0.5 credence instead.³¹

³⁰ Returning to the probabilistic interpretation I gave of “regarding a belief as arbitrarily formed”: the sense in which we can regard a currently held belief as independent of the truth is that *from the perspective in which we subject the belief in question to doubt* we regard the fact that we believe P as probabilistically independent of the truth of P. One might worry about talk of independence in such cases given that, once I've formed the belief, I may already be certain about which belief I formed. This is just the problem of old evidence, so I will not address this issue here. It's worth noting, however, that since, as I mentioned, there are different ways of subjecting a belief to doubt, there won't be a univocal answer to the question of whether some subject *S* regards a currently held belief as arbitrarily formed: the answer will be relativized to some particular way of doubting that belief.

³¹ In the literature, it has been common to explain judgments about the rationality of abandoning belief in such cases by appealing to a rational requirement along the following lines: if, *independently of one's reasoning about P*, one has good reason to think that one is unreliable with respect to P, one is required to abandon one's belief that P. (See, e.g. Elga (2007), Christensen (2007, 2010), Horowitz and Sliwa (2015) and Vavova (forthcoming)). This requirement is sometimes described as a requirement to “bracket” part of one's evidence or to not “give all of one's evidence its' due” (Christensen). But this raises two questions: First, *why* is one rationally required to bracket part of one's evidence or not give all of one's evidence its due? Second, exactly what does one need to bracket in order to satisfy this requirement? Note that on the approach I'm taking in this paper, these questions don't arise. For I'm not claiming that one is rationally required to deliberate from some perspective that doesn't take all of one's evidence into account. Rather, so far, all I'm claiming is that when we engage in doubt, we *just do* take up a perspective in which we're not willing to rely on certain beliefs or inferences that normally we're inclined to rely on. And indeed, there are many different perspectives one could take up that are compatible with doubting, and, in any given instance of doubt, we *just do* take up one of them. Recall, my primary aim is to illustrate a way of navigating one's doubts – not to comment on the merits or demerits of taking up the perspective of doubt to begin with. I'll say more at the end of the paper about what implications these considerations have concerning the rationality of taking up a doubtful perspective, and thereby the rationality of maintaining belief in cases like this one.

I set my alarm for 6:00am. I'll finish the answer key tomorrow morning.

Bottom Line: When I doubt a belief P, and the reasoning that I would need to infer P, the perspective of doubt will recommend that I abandon my belief if, from the perspective of doubt, I have a 0.5 credence in P, and I regard which belief I formed as independent of the truth.³²

Fourth Meditation: Religious Belief

It is time to turn to the cases that initially worried me: the cases of religious, moral, and political beliefs. I'm going to focus on a particular religious belief that I have: my belief that individuals don't come back to earth after they die as other life forms. Call this proposition "NR" (for "no reincarnation"). I think that I believe NR because I was subject to certain social influences rather than others. Upon realizing this, I begin subjecting my belief that NR to doubt. As I mentioned in the previous section, there are many different ways to subject a belief to doubt, and right now I'll consider two:

Way 1: I subject NR to doubt, but I don't subject various related beliefs to doubt. For example, I maintain my beliefs about the reliability of various religious texts, my beliefs about what it takes for me to persist, and my beliefs about what happens to my body and mind when I die.

If I doubt in Way 1, then I can easily recover NR from the perspective of doubt using the various related beliefs that I haven't subjected to doubt.

Way 2: I subject NR doubt, and also my beliefs about the reliability of religious texts, my persistence conditions, and my beliefs about what happens to my body and mind after I die. In other words, I subject to doubt *a cluster* of beliefs surrounding my belief in NR. I hope you have a feel for the cluster of beliefs I have in mind. Unfortunately, I can't write them all down on a list, (in part because the cluster includes an infinite number of beliefs such as NR, NR or $2+2=5$, NR or $2+2=6\dots$). The perspective I have in mind can be very roughly described as one that doesn't take for granted the beliefs that I have but that I wouldn't have had if I were raised in a

³² Here and in what follows I'm considering worst-case scenarios: complete independence. My claims can be generalized to cases in which we expect to be compromised in some way, but still do better than chance. See (Schoenfield (forthcoming)).

reincarnation-believing community. But even if I can't fully articulate the perspective, *I* have a sense of the perspective that I want to be deliberating from when I subject my belief that NR to doubt in Way 2 (and I have dispositions concerning which deliberative moves I'm willing to make when engaged in this kind of doubt and which I'm not).

And indeed, when I find myself doubting NR upon realizing that I believe NR because of the environment I grew up in, the kind of doubt I'm engaged in is of this latter variety. I suspect that the reason that, not just NR, but a cluster of related beliefs is being subject to doubt is that what elicited the doubt to begin with was a realization about the causal origins of this belief, and the causal origins of my belief that NR are the same as the causal origins of the cluster of beliefs from which NR could be inferred.

So the question is: when I engage in doubt in Way 2, can my belief in NR be recovered? It cannot be recovered using either of the two strategies I mentioned previously. I can't recover the belief using first order reasoning because all of the beliefs from which I could infer that I won't be reincarnated are being subject to doubt in Way 2. What about higher order considerations appealing to the fact that I formed the belief that NR? To determine whether this will work, I need to consider the following question: When I occupy the perspective of doubt, do I think that the fact that I formed the belief that NR makes it likely that NR is true? Answer: I do not. This is because I think that what determined whether I'd believe NR or R are facts about which community I grew up in, and I don't take the fact that I grew up amongst NR-believers as any evidence for NR.

But perhaps this is too quick: for in the perspective of doubt, I don't know *just* that I grew up amongst NR-believers. I know all sorts of things about the people I grew up with and I regard these people as a reliable source of information. So perhaps I can appeal to the fact that my belief was caused by growing up in *this* community of reliable people to recover my belief that NR. I don't think this will work. It is true that the people I was raised by are generally reliable about a host of mundane issues, but so are the people who believe in reincarnation. So the question is, in this perspective of doubt, do I think that the people I was raised by are more likely to be right than those I would have been raised by if I were raised, say, Hindu, *about reincarnation*? The problem is that the considerations I would ordinarily appeal to in defending the claim that the people I was raised by are particularly likely to be right about reincarnation rely on the various beliefs that I'm subjecting to doubt. For in this perspective of doubt, I can't appeal to considerations like "the

people who I was raised by have true beliefs about a number of related issues concerning the persistence conditions of human beings, the reliability of religious texts, and so on. Thus, being raised by these particular people made it very likely that I'd end up with a true belief about reincarnation."

Here's one final attempt to use higher order considerations to recover my belief that NR in the perspective of doubt: Perhaps I can think: "my community's beliefs are better aligned with a naturalistic or scientific world view than religious communities, and views that are better aligned with a naturalistic or scientific world view are more likely to be correct." Whether this strategy will work will depend on whether part of what I'm subjecting to doubt are the very beliefs that this chain of reasoning relies on. Am I subjecting to doubt my belief that my community's views are better aligned with science than the views of a Hindu community? Am I subjecting to doubt my belief that views that are better aligned with science are more likely to be true? If I am subjecting either of these to doubt, then I won't be able to use this sort of reasoning as a basis for thinking that I was more likely to arrive at the truth if I was raised in my community than in an alternative community. So I now face the question of whether, in the deliberation that I'm engaged in, I'm willing to rely on these beliefs. I find that I am not. This may be because I think that these beliefs are socially influenced in much the same way that my belief that NR was. So when worries about the social influences on belief lead me to take up the perspective of doubt, they lead me to doubt these beliefs as well.³³

I've argued that the two primary ways we recover beliefs in the perspective of doubt won't help me in this case: first order reasoning won't help because the beliefs from which I can infer NR are subject to doubt, and higher order reasoning won't help because of what I know about the way in which my belief was brought about. But it occurs to me that there is third way to recover a belief in the perspective of doubt, at least in one sense of the word "recover". This involves thinking about whether the perspective of doubt recommends that I abandon the belief that I formed, given that my aim is truth.

To figure out whether the perspective of doubt recommends abandoning my belief in NR, I need to think about how confident I am that NR when I adopt the perspective of doubt. As with other features of the perspective of doubt, I don't think that there is one unique attitude towards NR that anyone subjecting NR to doubt

³³ You may find that things go differently for you: perhaps, if you are an NR believer like me, beliefs along these lines are not ones that you are subjecting to doubt when you subject NR to doubt. In this case, congratulations! Your belief can be recovered.

will take. But in my own case, I don't find that, setting aside my belief in NR, I have a sharp credence, say, of 0.5 in NR. Rather, I find myself in a state in which I lack an opinion about NR – a state represented by quite a wide interval of credences.

I already established that the state in which I lack an opinion doesn't recommend itself over all states in which I'm more confident in NR than not. So, while in the case of certain beliefs formed arbitrarily, I can't recover my beliefs from the perspective of doubt inferentially, I can recognize that the perspective I'm occupying when I doubt isn't one that favors itself over a more opinionated state. Having realized this, I feel disinclined to abandon my belief that NR.³⁴

Let me end with a cautionary note: the fact that the perspective of doubt in which I lack an opinion doesn't recommend against *every* state in which I have a belief, doesn't mean that it doesn't recommend against *some* such states. For example, it might be true that in the perspective of doubt I am (determinately) less than .99 confident in the proposition that I won't be reincarnated. In this case, 0.99 won't be in my perspective-of-doubt NR-representor. If .99 is not in my NR-representor, then the arguments I've given do not support the claim that my state of lacking an opinion in the perspective of doubt permits being more than 0.99 confident that I won't be reincarnated. The arguments only show that the perspective doesn't recommend against degrees of confidence that are in my representor.³⁵ So if, in the non-doubting perspective (my ordinary one), I am more than 0.99 confident that I won't be reincarnated, then the perspective of doubt may recommend a significant reduction of confidence.

Bottom Line: Many of the beliefs that we're concerned are arbitrarily formed can be recovered from the perspective of doubt, if the attitude we take towards these propositions when we doubt them is the state of lacking an opinion. However, sometimes the perspective of doubt will recommend that we be less confident in these beliefs than we were previously.

³⁴ An interesting feature of this way of recovering a belief from doubt is that the perspective of doubt not only permits believing that I won't be reincarnated; it also permits believing I *will* be reincarnated. It's just that, in my own case, believing NR comes much more naturally to me than believing R or lacking an opinion about NR. This explains why, having subject my belief that NR to doubt, I return to a state in which I believe NR, rather than adopt one of the many other attitudes that are left open by the perspective of doubt.

³⁵ For example, according to the generalized expected accuracy rule, if every credence in my representor is below 0.99, then the state I'm in will recommend against being more than 0.99 confident.

Fifth Meditation: Disagreement

Jane and I went out for dinner tonight and at the end of the meal we each calculated our share of the restaurant bill.³⁶ We shared our answers: I concluded that we each owed 48.30 and Jane concluded that we each owed 46.50. As Jane started counting out her change, blatantly ignoring my opinion, I said to her: “Jane, don’t you think you should reconsider, perhaps redo your math, or use a calculator? After all, we’ve been keeping track of our arithmetical successes and failures during our nights out together, and when we’ve disagreed, I’ve been right 50% of the time.” “True,” Jane said, “but I must be right this time. After all, 1.2% of the total divided by two is 46.50, you say it’s not 46.50, so this must be one of the times in which I’m right and you’re wrong.”³⁷

I started wondering: is there anything I could say to Jane that would lead her to abandon her opinion? If Jane were to doubt her belief, then I could show her that, given the circumstances, the perspective of doubt recommends abandoning it. But so long as she isn’t inclined to doubt her belief, I don’t think there is any accuracy-based *argument* I could give her that would sway her. For as long as Jane is taking it for granted that 1.2% of the total is 46.50, abandoning this belief will look like a bad idea.

This frustrating experience with Jane led me to think that there isn’t much in the way of accuracy-aimed deliberation that we can engage in about whether to take up the perspective of doubt. For any such deliberation must take place from some perspective, and every perspective takes certain things for granted. Since a coherent perspective that takes P for granted will recommend continuing to take P for granted,³⁸ no amount of accuracy-aimed deliberation from such a perspective will motivate a shift to a perspective that doesn’t take P for granted. It follows that deliberating about which perspective to adopt can’t be done in a non-trivial way on the basis of accuracy considerations.

³⁶ This case is from Christensen (2007).

³⁷ This is the sort of reasoning that would be encouraged by, e.g., Kelly (2005), White (2010) and Titelbaum (2015).

³⁸ For recall that taking P for granted means that one is willing to reason on the basis of one’s belief that P. And if one is more than 0.5 confident that P, and willing to reason on the basis of this attitude, then one will think that maintaining a greater than 0.5 credence in P is advisable. (This will hold if one has a sharp credence in P, because credences are self-recommending, but also if one has a representor all of whose members are above 0.5, and one appeals to the notion of “generalized expected accuracy” discussed earlier). Since a belief that P will be recommended, reasoning on the basis of this attitude (in accuracy-approved ways, like conditionalization) will also be recommended. Thus, a perspective in one which one takes P for granted will recommend that one continue to take P for granted.

So what determines what we take for granted and what we subject to doubt? I suspect that it is largely arational processes: we try on different perspectives, some of them stick, and some of them don't. Doubt, in my view, is something that happens *to* us, not the outcome of a reasoning process that originates in our non-doubtful stance. Perhaps, in an epistemology class, I'm led to take up the perspective of doubt concerning all external world propositions. But this perspective doesn't stick. Once I leave the seminar, I find myself once again occupying a perspective that takes all sorts of external world propositions for granted. In contrast, in cases like Perceptual Coin Flip, once I realize that I can't recover my belief about the wall's color from the perspective in which I doubt it, I'm inclined to give it up. The doubtful perspective sticks. Those are clear cases in which I expect that there is a great deal of agreement. But some cases are murkier. When Jane and I disagreed, I was led to occupy a perspective in which I doubted my belief about the bill. But Jane wasn't. There are likely a variety of factors that contribute to whether we take up the perspective of doubt in any given case and whether it sticks once it's taken up. They might include the degree to which the belief is embedded in our overall web of beliefs, the degree to which error is made salient, the practical costs of abandoning the belief, our personality, which beliefs our friends and families are inclined to doubt, and so forth.

Bottom Line: Accuracy-based considerations don't privilege the perspective of doubt. Doubt is something that happens to us – not the output of accuracy-aimed deliberation.

Interlude: Normative Upshots

So far, I've simply noted that sometimes we're bothered when we can't recover a belief or inference from doubt, and other times we're not. But if you're someone that theorizes about rationality, this may be the juncture at which you wish to intercede. You may claim that there are substantive constraints along the following lines: If a belief of sort *B*, or an inference of sort *I*, can't be recovered from a perspective of doubt of sort *D*, then it is irrational to hold the belief that *B*. For instance, some internalists, like Descartes, may think that any belief other than a belief about one's mental states must be recoverable from a perspective in which all of one's external world beliefs are subject to doubt. Other internalists may claim that perceptual beliefs needn't be recoverable from a perspective of doubt unless there are special circumstances (e.g. "defeaters," or "positive reasons to think one is

unreliable”). Externalists may claim that so long as the belief or inference is reliably produced, it doesn’t matter whether it can be recovered from a perspective in which it’s doubted. Coherentists might claim that beliefs at the periphery of our web need to be recoverable from doubt, but not beliefs at the center. I won’t take a stand on these issues here. I will, however, register that I’m somewhat skeptical of the possibility of providing a well-motivated account that explains why in some cases it’s irrational to maintain a belief that is not recoverable from a perspective of doubt and in other cases it isn’t. My reasons for skepticism are similar to the reasons provided by Greco (2017), who is skeptical of the possibility of providing a well-motivated account of which propositions are “foundational.”³⁹

However, if you do have a theory on hand that tells you when a particular doubtful stance is required, then you can add the arguments I’ve given here to your theory, and derive verdicts about what’s rationally required in any given case, so long as you think that there is a reasonably tight connection between thinking rationally and thinking in a way that’s aimed at being accurate.⁴⁰ It’s worth noting two points though: the first is that, depending on your theory, these verdicts may not be ones that can be arrived at from a first personal deliberative stance. The second is that, depending on your theory, the parts that tell you when you’re rationally required to take up the doubtful stance may go beyond what would be licensed by an interest in accuracy alone. These are not criticisms: just points to be aware of in thinking about the relationship between the rationality theorist’s project and my primary project: demonstrating ways of navigating the doubt that arises when one regards one’s belief as arbitrarily formed, from an accuracy-aimed, first-person point of view.

Sixth Meditation: Conclusion

What’s disturbing about beliefs that have been heavily socially influenced is that these beliefs aren’t easily recoverable from a perspective in which we subject those beliefs to doubt. The beliefs aren’t recoverable in the first order manner because, in these cases, we’re not usually doubting just one particular belief, but a cluster of beliefs, a certain “picture” of how the world is. We can’t recover the belief in a higher order way either because, from the perspective of doubt, we regard the

³⁹ A belief is foundational in this sense if it is justified without the support of other beliefs. The questions of whether a belief is foundational, and of whether it’s rational to maintain a belief that can’t be recovered from a certain perspective of doubt are related but not identical.

⁴⁰ Thanks to Will Fleisher for discussion on this point.

belief as arbitrarily formed: we regard which belief we ended up with as independent of the truth. This blocks off two of the primary ways we can recover a belief from doubt. What I've suggested here, however, is that there is a third way to recover belief: so long as the state we're in when we doubt is the state of lacking an opinion, we can recognize that the perspective of doubt isn't one that recommends abandoning belief if our aim is to be accurate.⁴¹ This consideration may not compel us to maintain belief, but it allows us to realize that, if we are inclined to maintain an opinionated state, there is nothing accuracy-wise to be said against doing so.

References

- Argyle, M. and Beit Hallahmi, B. (1997). *The Psychology of Religious Belief, Behavior And Experience*. Routledge.
- Callahan, L. (ms.). "Voluntarist Permissivism."
- Carter, J.A. (forthcoming). "Sosa on Knowledge, Judgment and Guessing" *Synthese*.
- Christensen, D. (2007). "Epistemology of Disagreement: The Good News" *Philosophical Review* 116(2): 187-217.
- Christensen, D. (2010). "Higher Order Evidence" *Philosophy and Phenomenological Research* 81(1): 185-215.
- Cohen, G.A. (2000). *If You're an Egalitarian, How Come You're so Rich?* Harvard University Press.
- Eder, A.M. (ms.). "No Commitment to the Truth."
- Elga, A. (2007). "Reflection and Disagreement."

⁴¹ If we have some other aim, there may well be considerations that tell against arbitrarily formed belief in such cases. For example, we might regard it as desirable that we have a sort of epistemic agency over our beliefs, and think that if our beliefs are not going to be brought about in the appropriately agential sort of way, it's better not to have them at all (see Callahan (ms.)). I won't, in this paper, take a stand on whether there are good non-accuracy based motivations for preferring \mathcal{L} to an arbitrarily formed opinion, but I certainly don't want to rule out such a possibility.

- Fishburn, P.C. (1986). "The Axioms of Subjective Probability." *Statistical Science* 1(3):335-345.
- Friedman, J. (ms.). "Inquiry and Belief."
- Gibbard, A. (2008). "Rational Credence and the Value of Truth" in T. Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology Volume 2*, Oxford University Press.
- Gilboa, I. And Schmeidler, D. (1989). "Maxmin expected utility with non-unique prior." *Journal of Mathematical Economics* 18(2): 141-153.
- Glass, J., Bengston, V.L., and Dunham, C.C. (1986). "Attitude Similarity in Three Generation Families: Socialization, Status Inheritance, or Reciprocal Influence?" *American Sociological Review* 51(1): 685-698.
- Greaves, H. And Wallace, D. (2006). "Justifying conditionalisation: conditionalization maximizes expected epistemic utility" *Mind* 115(459): 607-632.
- Greco, D. (2017). "Cognitive Mobile Homes" *Mind* 126(501): 93-121.
- Hájek, A. (2003). "What Conditional Probabilities Could Not Be" *Synthese* 137(3): 273—323.
- Horowitz, S. (2013). "Immoderately Rational." *Philosophical Studies* 167(1):1-16.
- Horowitz, S. and Sliwa, P. (2015). "Respecting All the Evidence" *Philosophical Studies* 172(11): 2835-2858.
- James, W. (1896). "The Will to Believe" *The New World* 5: 327-347.
- Jeffrey, R. (1983). "Bayesianism with a Human Face" in John Earman (ed.), *Testing Scientific Theories*. University of Minnesota Press. pp. 133--156.
- Joyce, J. (2005). "How Probabilities Reflect Evidence." *Philosophical Perspectives* 19 (1):153-178.
- Joyce, J. (2009). "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief" in F. Huber & C. Schmidt-Petri (eds.), *Degrees of Belief*. Synthese Library.

- Joyce, J. (2010). "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives* 24 (1):281-323.
- Joyce, J. (ms.). "Evidence and the Accuracy of Credences."
- Kelly, T. (2005). "The Epistemic Significance of Disagreement" in J. Hawthorne and T. Gendler (eds.), *Oxford Studies in Epistemology, Volume 1*. Oxford University Press.
- Konek, J. (forthcoming). "Epistemic Conservativity and Imprecise Credence." *Philosophy and Phenomenological Research*.
- Kyburg Jr, H.E. (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press.
- Lasonen Aarnio, M. (2010). "Unreasonable Knowledge" *Philosophical Perspectives* 24 (1):1-21.
- Lasonen Aarnio, M. (2014). "Higher Order Evidence and the Limits of Defeat" *Philosophy and Phenomenological Research* 88 (2):314-345.
- Levi, I. (1974). "On Indeterminate Probabilities." *Journal of Philosophy* 71(13): 391-418.
- Levinstein, B. (2017). "A Pragmatist's Guide to Epistemic Utility" *Philosophy of Science* 84(4): 613-638.
- Mayo-Wilson, C. and Wheeler, G. (2016). "Scoring Imprecise Credences: A Mildly Immodest Proposal." *Philosophy and Phenomenological Research* 93(1): 55-78.
- Oddie, J. (1997). "Conditionalization, Cogency and Cognitive Value" *British Journal of the Philosophy of Science* 48: 533-541.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press.
- Rinard, S. (2015). "A Decision Theory for Imprecise Probabilities." *Philosophers' Imprint* 15(7): 1-16.
- Rinard, S. (2017). "Imprecise Probability and Higher Order Vagueness" *Res Philosophica* 94 (2):257-273.

- Rinard, S. (ms.). "Pragmatic Skepticism."
- Seidenfeld, T., Schervish, M.L, and Kadane, J.B. (2012). "Forecasting with Imprecise Probabilities" *International Journal of Approximate Reasoning* 53:1248-61.
- Schervish, M. (1989) "A General Method for Comparing Probability Assessors" *Annals of Statistics* 17(4):1856-1879.
- Schoenfield, M. (2017). "The Accuracy and Rationality of Imprecise Credences." *Noûs* 51(4): 667-685.
- Schoenfield, M. (forthcoming). "An Accuracy Based Approach to Higher Order Evidence." *Philosophy and Phenomenological Research*
- Srinivasan, A. (2015). "The Archimedean Urge" *Philosophical Perspectives* 29 (1): 325-362.
- Titelbaum, M. (2015). "Rationality's Fixed Point." *Oxford Studies in Epistemology* 5 (2015): 253-294
- van Fraassen, B. (1990). "Figures in a Probability Landscape" in J. Dunn & A. Gupta (eds.), *Truth or Consequences*. Kluwer Academic Publishers. pp. 345-356.
- van Fraassen, B. (2005). "Conditionalizing on Violated Bell's Inequalities" *Analysis* 65 (1):27-32.
- van Fraassen, B. (2006). "Vague Expectation Value Loss." *Philosophical Studies* 127(3): 483-491.
- Vavova, K. (forthcoming). "Irrelevant Influences." *Philosophy and Phenomenological Research*.
- Weatherson, B. (ms.). "The Evidentialist Theory of Disagreement."
- White, R. (2010). "You Just Believe That Because..." *Philosophical Perspectives* 24 (1):573-615