

How to be a Bayesian Dogmatist

Brian T. Miller

September 16, 2015

Abstract

Rational agents have (more or less) consistent beliefs. Bayesianism is a theory of consistency for partial belief states. Rational agents also respond appropriately to experience. Dogmatism is a theory of how to respond appropriately to experience. Hence Dogmatism and Bayesianism are theories of two very different aspects of rationality. It's surprising, then, that in recent years it has become common to claim that Dogmatism and Bayesianism are actually inconsistent: how can two independently consistent theories with distinct subject matter be jointly inconsistent? In this essay I argue that Bayesianism and Dogmatism are inconsistent only with the addition of a specific hypothesis about how the appropriate responses to perceptual experience are to be incorporated into the formal models of the Bayesian. That hypothesis isn't essential either to Bayesianism or to Dogmatism, and so Bayesianism and Dogmatism are consistent. That leaves the matter of how experiences and consistent partial belief states are related, and so in the remainder of the essay I offer an alternative account of how perceptual justification as the Dogmatist understands it can be incorporated in the Bayesian formalism.

1 Introduction

I'm walking down the street and I have a visual experience as of a red ball lying on the grass. What's the epistemic significance of my having had that experience? One likely result is that I obtain some justification for a belief about my own experiences, something like *I've had an experience as of a red ball lying on the grass*. Another is that I obtain some justification for a belief about the world, something like *there's a red ball lying on the grass*. Yet another is that I now find myself with justification to believe further propositions inferentially

related to the first two: if I already had justification to believe that there's a bike on the grass and then I have my perceptual experience as of the ball, I obtain some justification for believing *there are at least two toys on the grass*. My justification for the last of these three propositions is unambiguously mediate, as it's at least in part my justification for believing something else that makes me justified in believing that there are at least two toys on the grass. In contrast, my justification for believing *I've had an experience as of a red ball lying on the grass* comes directly from the experience itself without the mediation of some other justification that I have, and hence that justification is immediate. That much I'll take for granted as common ground between *Inferentialist* and *Dogmatist* accounts of perceptual justification. What's contentious between the two is the status of the second proposition.

According to the Dogmatist, perceptual justification is both immediate¹ and underminable^{2,3}. Moreover, the Dogmatist thinks that while a perceptual experience may generate immediate and underminable justification for *I'm having an experience as of p* or some other proposition about the agent's mental states, it also generates immediate and underminable justification for p itself.

In contrast, the Inferentialist claims that my beliefs about the external world are never immediately justified (at least not on the basis of experience), but rather depend upon an inference from an immediately justified proposition about my own experiences together with an auxiliary proposition connecting facts about my experiences to facts about the external world, e.g. *If I have a perceptual experience as of p then, probably, p*. Hence it's my justification for believing *I've had an experience as of a red ball lying on the grass* together with

¹My justification for believing that p is *immediate* in the relevant sense unless it is in part my having justification to believe something else that *makes me* justified in believing that p. 'Makes' here expresses a relation of epistemic dependence, a variety of modal dependence. Hence Dogmatism shouldn't be confused with the much stronger thesis that having a perceptual experience in the absence of defeaters is *sufficient* for obtaining perceptual justification, as we allow that there might be other necessary conditions for obtaining perceptual justification besides my having justification to believe something else, as long as the satisfaction of that condition is not part of what makes me justified. For example, an externalist Dogmatist might insist that obtaining perceptual justification for believing that p requires that perception be a reliable belief forming process. Alternately the Dogmatist might claim with White and Wright (see fn. 4) that we possess default, a priori justification to believe that our perceptual faculties are reliable, but then deny that this justification is even in part what makes me have the perceptual justification that I have (see Silins (2007) and McGrath (2013) for discussion of this possibility). In neither case is it my justification to believe some other proposition q part of what makes me justified in believing that p, and so my justification for p is immediate.

²For the distinction between undermining/ undercutting and opposing/ rebutting defeaters see Pollock and Cruz (1999) p. 196-7.

³See Pryor (2000), (2005), and (2013).

my justification for believing some such auxiliary proposition that makes it the case that I have justification for believing *there's a red ball lying on the grass*, and so that last bit of justification is mediate.⁴

Dogmatism makes obtaining perceptual justification relatively easy: any agent capable of having a contentful experience and lacking defeaters is in a position to obtain justification for lots of beliefs about the world without first acquiring justification for beliefs about the relationship between experience and the external world. Whether this is ultimately a virtue of the theory or a shortcoming is contentious: easily acquired justification for propositions about the external world might be thought to license too-easy responses to skeptical challenges to our knowledge of the the external world and too-easy knowledge of the reliability of our perceptual faculties. If Inferentialism is correct then obtaining perceptual justification is in some sense harder, as we first need justification to believe the auxiliary proposition connecting the having of an experience with facts about the world. Making it harder to obtain perceptual justification comes with its own set of problems, as now we're faced with the difficult task of explaining where justification for believing the auxiliary propositions comes from, potentially leaving skeptical problems insoluble.⁵

In this essay I defend Dogmatism against a very different objection: that it is inconsistent with Bayesianism. The Bayesian Argument (as I'll call it) purports to show that given Bayesianism, acquiring perceptual justification for believing that *there's a red ball lying on the grass* requires that I already have justification for ruling out a wide range of skeptical scenarios on which I have a non-veridical experience as of the ball lying on the grass. If obtaining perceptual justification for believing that B requires that I already have justification for believing that A, then (the objection goes) it's plausible that my justification for A is what makes me justified in believing B, in which case my justification for believing that B

⁴Versions of Inferentialism are distinguished by their account of how we obtain justification for that auxiliary hypothesis. One example is Abductivism, the view that our justification comes from an inference to the best explanation. Russell (1997) argued that of all of the possible explanations for our having the experiences that we do in fact have, the hypothesis that an external world exists and causes our experiences is the simplest hypothesis that '[accounts] for the facts of our own life'.(23) Vogel (1990) pursues a similar strategy. More recently a variety of Inferentialism known as Rationalism has been widely discussed, on which we have a special type of a priori justification for believing the auxiliary hypothesis. Defenses of this form of Rationalism are found in White (2006) and Wright (2004).

⁵The details will of course depend upon the particular versions of Dogmatism and of Inferentialism in question. Inferentialists of a Rationalist persuasion can dismiss skeptical objections as easily as the Dogmatist, and needn't bootstrap their way to justification for believing in the reliability of their perceptual faculties, as that justification is an a priori 'entitlement'. On the other hand, the Reliabilist Dogmatist who happens to be a brain in a vat may never be in a position justifiably to believe that skeptical scenarios are false.

isn't immediate. Since this result allegedly follows from the Bayesian formalism, we thereby have some reason to believe that Dogmatism and Bayesianism are inconsistent, and since Bayesianism is an attractive theory we thereby have a reason to reject Dogmatism.

The literature contains two types of response to this argument on behalf of the Dogmatist. The first response is to accept the inconsistency of Dogmatism and Bayesianism and take that as a good reason to revise orthodox Bayesianism (see Weatherston (2007)). The second and seemingly more common response is to accept the formal result — that a necessary condition for obtaining justification for believing the content of perceptual experience is antecedent justification for believing some other proposition — but then to deny that it entails the mediacy of perceptual justification. One way to do this would be to take inspiration from Silins (2007) and argue that having justification to believe that A might be a necessary condition for obtaining justification for believing that B without A being what makes it the case that I have that justification for believing that B. Mere modal dependence just isn't what matters when it comes to questions of immediacy, and hence my justification for believing that B might nonetheless be immediate.⁶

I pursue a third response to the Bayesian Argument on behalf of the Dogmatist: I deny that the putatively problematic formal result is a commitment of the Bayesian at all. The derivation of that result requires a premise that goes beyond the core commitments of Bayesianism to specify precisely how the epistemic significance of having an experience is to be reflected in the model. This requires that I be clear about exactly what the Bayesian is and is not committed to, an issue that I discuss in §2. In §3 I lay out the formal details of the Bayesian Argument. The heart of the paper is found in §4, in which I identify the problematic premise and argue that it is neither a commitment of the Bayesian nor a claim that the Dogmatist should find plausible. I then offer an alternative account of how the epistemic impact of experience can be incorporated into Bayesian models. If that account is correct, then the Bayesian Argument is unsound and hence presents no reason to doubt the consistency of Dogmatism and Bayesianism. In §5 I consider the implications of adopting my

⁶A second way to accept the formal result of the Bayesian Argument without abandoning Dogmatism exploits the fact that Dogmatism is discussed in the idiom of reasons while Bayesianism is discussed in the idiom of credences. Translating between the two idioms is not entirely straightforward. In particular, it's not obvious that obtaining a reason to believe that p always leads to an increase in one's credence that p. If not then it might be possible to learn that p perceptually even though my prior credence in a skeptical hypothesis prevents my increasing my credence in p as a result of having an experience as of p. See Zardini (2014).

suggestion for various versions of Bayesianism, concluding that the Dogmatist should embrace a version that incorporates Richard Jeffrey's more permissive approach to conditionalization as opposed to the strict version of conditionalization associated with Classical Bayesianism.

2 Bayesianism

Bayesianism is a theory of the rationality of partial belief states. The starting assumption is that an agent's partial beliefs can be represented as a function from propositions to numbers representing that agent's subjective probability or credence that various propositions are true. Some such functions are rational and some are not. The core of Bayesianism is the postulation of two conditions necessary for a credence function to be rational. The first is Probabilism:

Probabilism: all rationally permissible credence functions are probabilistically coherent (i.e. consistent with the probability axioms)

Probabilism imposes a synchronic constraint upon rational credence functions. Constraining the rationality of revisions to those credence functions over time is the thesis of Conditionalization. Conditionalization requires that we divide our credences into two types: conditional and unconditional. Whereas unconditional credences reflect an agent's degree of confidence in the truth of a proposition, conditional credences reflect the agent's confidence in a proposition given the truth of some other proposition. For example, the agent might assign a low unconditional credence to *the street is wet* but a much higher credence to it given *it's raining*: formally, $P(\textit{the street is wet}) = .2$ and $P(\textit{the street is wet} \mid \textit{it's raining}) = .9$. The intuition motivating Conditionalization is that the credences that an agent should adopt in the future upon obtaining new information are importantly constrained by the conditional credences that he accept right now, and that those constraints are embodied in the agent's currently held conditional credences. I'll be discussing two ways of making this intuition rigorous. First:

Strict Conditionalization: if for some reason I set $P_{new}(B) = 1$, then I must conditionalize upon B by setting my posterior credence in any proposition A, $P_{new}(A)$, to the prior conditional probability of A upon B, $P_{old}(A|B)$.

It important to note that according to Strict Conditionalization, incorporating new information B by conditionalizing upon it requires changing $P(B)$ to

1.⁷ Jeffrey Conditionalization generalizes Strict Conditionalization by allowing for updates upon changes in credences to values other than 1:

Jeffrey conditionalization: if for some reason I change $P(B)$ to any value, then I'm rationally required to set $P_{new}(A) = P_{old}(A|B)P_{new}(B) + P_{old}(A|\neg B)P_{new}(\neg B)$.⁸

Bayesianism is the combination of Probabilism with either version of Conditionalization (I'll just say 'Conditionalization' when my comments apply equally to both versions), with Classical Bayesianism being the combination of Probabilism and Strict Conditionalization and Jeffrey Bayesianism being the combination of Probabilism and Jeffrey Conditionalization.⁹ Since my task in this essay is to show the Bayesian Argument exposes no great tension between Dogmatism and either version of Bayesianism, I will proceed to show that that argument exposes no great tension between Dogmatism and the combination of Probabilism and either version of Conditionalization (though later on I settle on Jeffrey Bayesianism as the better compliment to Dogmatism).

3 The Bayesian Argument Against Dogmatism

Suppose that Dogmatism is true and so perceptual justification is immediate and underminable. In that case there is an asymmetry in our treatment of potential undermining defeaters. To illustrate, consider a case in which I have an experience as of my hands and a skeptical scenario in which I'm a handless brain in a vat having hand-like experiences (I call the proposition that I am in this skeptical scenario 'BIV'). Most epistemologists, including Dogmatists, will agree that if I have high levels of justification *for* BIV then my experience as of my hands will fail to generate much justification for the proposition *I have*

⁷For reasons that I discuss in §4.1, especially fn. 21, I'll assume throughout the essay that proponents of Strict Conditionalization will prohibit 'exogenous' credence revisions (again, see §4.1) to values less than 1.

⁸See Jeffrey (1983) p. 169. Jeffrey goes on to generalize this condition to accommodate changes to the partition involving more than two propositions, a complication inessential to the present essay.

⁹Let me emphasize that by 'Jeffrey Bayesianism' I mean the combination of Probabilism and Jeffrey Conditionalization (as defined above) only. Richard Jeffrey seems to have been sympathetic to Donald Davidson's claim that perceptual experiences cause but do not justify belief (see Jeffrey's (1983) p. 184-5 and 211 and Davidson's (2001) (though Jeffrey doesn't specify which article in that anthology he had in mind, presumably it's *Actions, Reasons, and Causes* from 1963); Davidson's clearest take on the issue is found in his *A Coherence Theory of Truth and Knowledge*, but that wasn't published until 1986, three years after Jeffrey's publication). I do not share those sympathies.

hands: BIV is undermining the justificatory force of the experience.¹⁰ One distinctively Dogmatist idea is that while possessing justification for believing that an undermining defeater obtains can prevent me from acquiring perceptual justification, the absence of justification for believing that an undermining defeater does not obtain does not prevent me from acquiring that justification. In our particular case that means that I don't need justification for \neg BIV in order to acquire justification for *I have hands* from that experience. For the dogmatist, possessing an undermining defeater blocks the acquisition of perceptual justification, but lacking justification to reject an undermining defeater is perfectly consistent with the acquisition of perceptual justification.

The Bayesian argument against Dogmatism purports to show that this asymmetric treatment of undermining defeaters is in tension with Bayesianism by showing that the credence that I assign to *here's a hand* as a result of having an experience as of a hand cannot exceed the credence that I assigned to \neg BIV before I had that experience. In other words, my prior credence in \neg BIV imposes a cap or limit upon my posterior credence in *here's a hand*. If that's correct then it is impossible rationally to adopt a high credence in *here's a hand* unless I already have a high credence in \neg BIV. That provides some reason to believe that perceptual justification isn't immediate after all. This result accords nicely with Inferentialism, but purportedly creates a problem for Dogmatism.¹¹

Recall that BIV is the hypothesis that I'm a handless brain in a vat having experiences as of my hands. Hence the truth of BIV entails that I'm having an experience as of my hands. Taking e as the proposition *I'm having an experience as of my hands*, that means:

1. $P_{old}(BIV|e) \geq P_{old}(BIV)$ ¹²

If conditionalizing upon e raises or leaves unchanged my credence in BIV then it must also lower or leave unchanged my credence in \neg BIV, and so:

2. $P_{old}(\neg BIV|e) \leq P_{old}(\neg BIV)$

¹⁰BIV is also an opposing defeater in this case.

¹¹Versions of the argument have appeared in Cohen (2002), Hawthorne (2004), Schiffer (2004), and White (2006).

¹²For clarity of presentation my treatment of the tense of BIV, e , and h has been loose. Concerned readers may index each proposition to some moment t_2 , so that BIV becomes *At t_2 I will be a handless brain in a vat having an experience as of my hands*, e becomes *At t_2 I'll have an experience as of my hands*, and h becomes *At t_2 I'll have hands*. We then interpret P_{new} as the credence function that I adopt at t_2 and P_{old} as the credence function that I adopt at the previous moment t_1 .

When I have an experience as of my hands I thereby obtain some justification for believing that I'm having an experience as of my hands. According to Conditionalization I must now update upon that changed credence in e ,¹³ and so:

$$3. P_{new}(\neg BIV) = P_{old}(\neg BIV|e)^{14}$$

Combining terms from (2) and (3) we get:

$$4. P_{new}(\neg BIV) \leq P_{old}(\neg BIV)$$

BIV is the hypothesis that I'm a handless brain in a vat having an experience as of my hands. If I have hands then I'm not a handless brain in a vat having an experience as of my hands, and so in that case BIV is false. Taking h as the proposition *I have hands*, it follows that:

$$5. h \models \neg BIV$$

¹³While this is correct as far as it goes — Conditionalization does indeed require that we update upon changes to the probabilities that we assign to propositions like e — I will argue in §4 that updating upon e alone is both unwarranted and key to the argument. Nonetheless at this point I'll suppose that it is correct in order to present my opponent's argument.

¹⁴In particular, (3) is meant to follow from *Strict* Conditionalization (plus the description of the case). The argument is slightly different when Jeffrey Conditionalization is employed, as (3) now becomes:

$$P_{new}(\neg BIV) = P_{old}(\neg BIV|e)P_{new}(e) + P_{old}(\neg BIV|\neg e)P_{new}(\neg e)$$

As I note in §4.3, for our purposes BIV is equivalent to the hypothesis that $e \& \neg h$, and so this is equivalent to:

$$P_{new}(\neg(e \& \neg h)) = P_{old}(\neg(e \& \neg h)|e)P_{new}(e) + P_{old}(\neg(e \& \neg h)|\neg e)P_{new}(\neg e)$$

Equivalently:

$$P_{new}(\neg e \vee h) = P_{old}(\neg e \vee h|e)P_{new}(e) + P_{old}(\neg e \vee h|\neg e)P_{new}(\neg e)$$

Which simplifies to:

$$P_{new}(\neg e \vee h) = P_{old}(h|e)P_{new}(e) + 1(P_{new}(\neg e))$$

$e \& \neg h$ is at least possible, and so $P_{old}(h|e) < 1$. As a result, the higher the value of $P_{new}(e)$ the lower the value of $P_{new}(\neg e \vee h)$. This is most easily seen by first considering the case in which $P_{new}(e) = 0$. In that case $P_{new}(\neg e) = 1$, and so $P_{new}(\neg e \vee h) = (P_{old}(h|e))0 + 1(1) = 1$. As $P_{new}(\neg e)$ decreases, $P_{new}(e)$ increases, which in this case means that as $1(P_{new}(\neg e))$ decreases, $P_{old}(h|e)P_{new}(e)$ increases. Importantly, however, the changes aren't proportional: since $P_{old}(h|e) < 1$, when the value of $P_{new}(e)$ increases then the increase in $P_{old}(h|e)P_{new}(e)$ is smaller than the decrease in $1(P_{new}(\neg e))$. Hence if P_{new} is the credence that I ought to adopt upon increasing my confidence in e (and nothing else) and updating accordingly, then $P_{new}(\neg e \vee h) < P_{old}(\neg e \vee h)$. $h \models \neg e \vee h$, and so $P_{new}(h) \leq P_{new}(\neg e \vee h)$. Put this all together and we get:

$$P_{new}(h) \leq P_{new}(\neg e \vee h) < P_{old}(\neg e \vee h).$$

In plain English, if P_{new} is the credence function that I adopt as a result of increasing my confidence in e (and nothing else) and updating accordingly, then my new credence in h must actually be lower than my old credence in $\neg e \vee h$, i.e. in $\neg BIV$.

Hence:

$$6. P_{new}(h) \leq P_{new}(\neg BIV)$$

Combining terms from (6) and (4) we get

$$7. P_{new}(h) \leq P_{old}(\neg BIV)$$

What (7) says is that my credence in $\neg BIV$ before I had the experience as of my hands must be at least as high as my posterior credence in *I have hands*, the credence that I adopt after having an experience as of my hands and conditionalizing. Since the Dogmatist thinks that after having an experience as of my hands my credence in the proposition *I have hands* is very high, that means that my prior credence in $\neg BIV$ must have been very high as well. That's tantamount to saying that assigning a high credence to $\neg BIV$ is a necessary condition for assigning a high credence to *I have hands* on the basis of my perceptual experiences, which (it is claimed) is inconsistent with the hypothesis that my perceptual justification for *I have hands* is immediate.¹⁵ Analogous arguments show that no perceptual justification is immediate, and so Dogmatism is false.

4 Modeling Experience

(1), (2), (5) and (6) follow from Probabilism together with the logical relations that obtain between BIV, h, and e.¹⁶ (4) is a consequence of (2) and (3), and (7) is a consequence of (4) and (6). But what about (3), that $P_{new}(\neg BIV) = P_{old}(\neg BIV|e)$? Rejecting any other step in the argument requires giving up on Probabilism (and hence giving up on Bayesianism itself), but not so with (3). If (3) is false then the argument for (7) is unsound, and so the putative tension between Dogmatism and Bayesianism is resolved.

But how can we reject (3) without rejecting Conditionalization? Note that there are two importantly different ways to think about credence function $P_{new}(\cdot)$

¹⁵That the existence of such a necessary condition is inconsistent with the immediacy of perceptual justification is far from obvious, as I discussed in fn. 1. However, if we rely on that point to respond to the Bayesian Argument we are essentially denying that (7) is problematic without disputing its truth, and hence we must still concede that obtaining perceptual justification requires that we already have justification for assigning low credences to the relevant skeptical hypotheses. I find that implausible, and so in what follows I offer a response that allows the rejection of (7) without requiring the rejection of Bayesianism itself.

¹⁶Plus a very weak assumption about the relationship between conditional and unconditional probabilities: that those values satisfy the equation at the heart of the Ratio Analysis of conditional probability. Presumably any reasonable credence function would assign a credence other than 1 or 0 to BIV, e, and h, and so that assumption is not in question. See §5.

and so two importantly different ways of thinking about (3). $P_{new}(\cdot)$ might be understood simply as the credence function resulting from accepting $P_{old}(\cdot)$ and then updating on e, and in that case (3) is a trivial consequence of Conditionalization. Alternately, $P_{new}(\cdot)$ might be understood only as *the credence function that an agent who accepts P_{old} ought to adopt after having an experience as of his hands* and updating accordingly, whatever that function happens to be.

To avoid equivocation, $P_{new}(\cdot)$ must be interpreted the same at (3) and at (7). (7) asserts a condition on $P_{new}(h)$: that it can be no higher than $P_{old}(\neg BIV)$. Dogmatism is a thesis about the epistemic impact of perceptual experience. Hence in order for (7) to present even a prima face problem for Dogmatism, $P_{new}(h)$ must be interpreted as the credence function that an agent who accepts $P_{old}(\cdot)$ ought to adopt after having an experience as of his hands and updating accordingly. In other words, at both (3) and (7), $P_{new}(\cdot)$ must be interpreted in the second of the two ways considered in the previous paragraph.

If we assume that the appropriate way for an agent to respond to having an experience as of his hands is to raise his credence in e and update accordingly then the two formulations are equivalent, but that's a substantive assumption. In particular, it's a substantive assumption that it's upon e (and e alone) that the agent must update. My proposal on behalf of the Bayesian Dogmatist is that instead of updating on $e = I'm\ having\ an\ experience\ as\ of\ a\ hand$ we should update on $h = here's\ a\ hand$.

This proposal is consistent with both Classical and Jeffrey Bayesianism (see §4.1), though for independent reasons its combination with Classical Bayesianism is unappealing to the Dogmatist. Indeed, I will later argue that adopting Jeffrey Bayesianism, together with the thought that it is upon h that we should update (rather than upon e), not only allows the Dogmatist to avoid (7), but also provides a very natural way for the Dogmatist to model perceptual learning in a Bayesian framework.¹⁷

¹⁷Since I completed this essay, a somewhat similar approach has appeared in Moretti (2015). Our responses to the issue are, nevertheless, importantly different. According to Moretti, a basic problem with White's argument is that it requires updating on a belief rather than on an experience – it “presuppose[s] a notion of perceptual evidence that is not the one distinctive of dogmatism” (271). But all Bayesian models share that requirement — you can't conditionalize on an experience! — and hence if White's presupposition is inconsistent with Dogmatism then Bayesianism is inconsistent with Dogmatism too. This rests on a mistake: what's required is simply that we allow experiences to spark credence revisions that are exogenous to the model (see §4.1 below); without some such allowance it's hard to see how the epistemic significance of experience could ever make it into a Bayesian model.

On my view, White's argument is unsound not because he updates on a proposition, but because he updates on the wrong proposition. Moretti is uncomfortable with updates on the contents of experiences when the agent also introspects on that experience (though he

I'll defend that proposal in three stages: first I'll argue that the rejection of (3) is consistent with Probabilism and with both Strict Conditionalization and Jeffery Conditionalization, and hence it cannot properly be regarded as a commitment of the Bayesian. Second, I'll argue that this is precisely what the Dogmatist should have wanted all along – that (3) is implausible to the Dogmatist independently of the Bayesian Argument. Obviously if (3) is rejected then the Bayesian argument is unsound, and so we are not compelled to accept the problematic (7). We may still worry, however, that a similar problem arises from updating on h . In the third stage of my defense I'll show that this is not the case, that my revision does not give rise to analogous problems, and what's more that it provides an account of how we come to assign a high credence in \neg BIV when we were unable to do so before.

4.1 Bayesianism Does Not Entail (3)

I begin by showing that updating on h and hence rejecting (3) is perfectly consistent with Bayesianism. My comments in this section will apply equally to both the Jeffrey and the Classical versions of Bayesianism except where I specify otherwise.

As discussed in §2, the theses of Probabilism and Conditionalization allow Bayesians to construct a model of an agent's partial belief states and the revisions in those states over time. It's important to notice that this is at best a partial model of rational credence revision. Given Conditionalization, a prior credence function plus a new credence assignment to some proposition at t_1 completely determine the posterior credence function that must be adopted at the next moment t_2 . So for any propositions A and B, if at t_1 my credence function P_{old} assigns $P_{old}(A|B) = .7$ and then I set $P_{new}(B) = 1$, then Conditionalization requires that at t_2 I set $P_{new}(A) = .7$.

does note that doing so would provide the Dogmatist with a technically adequate response to White). Contrary to my approach, his position shares White's assumption that an experience together with the belief that you've had that experience have the same evidential significance as the belief alone (274). There are two reasons that the Dogmatist should reject this picture. First, it's independently problematic: suppose that S and S* each come to believe *I've had an experience as of A*. S believes it because she really did have an experience as of A, and S* believes it as a result of wishful thinking. Do their respective beliefs have the same epistemic significance? Second, if the White/ Moretti assumption is correct, then unless I assign a low prior probability to BIV and other undermining hypotheses, the only way that I can obtain perceptual justification for h is to avoid introspecting on my experiences. When we do introspect we're back on the hook for the Bayesian Argument, as we're then required to set $P_{t_3}(BIV) = P_{t_1}(BIV|e) \geq P_{t_1}(BIV)$, even if at t_2 we conditionalized on h and so decreased our confidence in BIV. (What happened to our evidence h ? The Bayesian formalism doesn't allow do-overs! See §5.) This is a defense of Dogmatism that Dogmatists might not welcome.

But what about that other factor determining my posterior credence: the input to the model, which is the revision in credence that prompted the conditional update in the first place? Clearly these can't all be the result of other conditionalizations, as the process of conditionalization only gets going with a change in credence and ends as soon as the new credence function is adopted. Hence if there are to be any rationally permitted credence revisions at all there must be some that do not proceed by conditionalization. At least some credence revisions are rational, and hence any plausible version of Bayesianism must accept the permissibility of at least some credence revisions that don't proceed via conditionalization. All of the credence revisions that are modeled by the Bayesian formalism are conditionalizations, so it follows that some credence revisions are not modeled. Call those credence revisions that are not modeled by the Bayesian *exogenous* revisions (as in *exogenous to the model*) and those occurring within the model via conditionalization *endogenous* revisions.¹⁸

Two points about exogenous credence revisions are worth emphasizing. First, the rational permissibility of an exogenous revision is largely unconstrained by the Bayesian machinery. Probabilism prohibits the adoption of any probabilistically incoherent credence, and since exogenous credence revisions are incorporated into the agent's posterior credence function¹⁹ it follows that Probabilism prohibits exogenous revisions that are themselves probabilistically incoherent. For example Probabilism prohibits the exogenous revision of $A \& \neg A$ to any value other than 0. Jeffrey Conditionalization imposes no constraints upon the appropriateness of the exogenous inputs: its sole function is to determine the appropriate response to a given revision. Hence any exogenous credence revision is consistent with Probabilism and Jeffrey Conditionalization as long as it is probabilistically coherent with itself.²⁰

Things are a bit more complicated with Strict Conditionalization, on which updating is permitted only on propositions assigned a credence of 1. Exogenous credence revisions that don't lead to updating can result in an incoherent posterior credence function,²¹ so there's good reason for the Classical Bayesian

¹⁸This terminology comes from Howson and Ubach (1993).

¹⁹So if setting my credence in A to 1 is an exogenous revision then after updating my posterior credence function P_{new} will be such that $P_{new}(A) = 1$.

²⁰That's a bit too strong, though the details will depend on the how we understand the relationship between conditional and unconditional probabilities. It's common for Bayesians to accept that $P(A|B) =_{df} \frac{P(A \& B)}{P(B)}$. If this so-called Ratio Analysis of conditional probability is accepted then once a proposition is assigned a credence of either 1 or 0 it becomes impossible to change that credence by conditionalization – see §5.

²¹If my credence function started out coherent and I exogenously revise my credence in a

to prohibit exogenous revisions that cannot be updated upon, i.e. to prohibit exogenous revisions to credences other than 1. Nonetheless, exogenous credence revision is consistent with Classical Bayesianism as long as (i) it is probabilistically coherent with itself and (ii) the credence of the proposition being exogenously revised is thereby raised all the way to 1.²²

The second point is that the process of incorporating the epistemic impact of having had a perceptual experience must begin with an exogenous credence revision. Suppose that that's false, and the credence revisions that result from having a perceptual experience are entirely endogenous and so proceed by conditionalization. As we've seen, conditionalization is a process set in motion by a very specific type of event: a change in credence or subjective probability. To illustrate, note that I update on *it's raining* not when it's actually true that it's raining (i.e. not when the objective probability of rain changes), but when my credence or subjective probability in *it's raining* changes. All instances of conditionalization begin with a change in credence and end with a change in credence. In contrast, when I revise my credences in response to a perceptual experience, the process begins with something that isn't a change in credence — the actual having of the experience — and ends with a change in credence. Hence the initial credence revision coming in response to perceptual experience can't proceed via conditionalization and hence can't be endogenous to the Bayesian machinery.

With these points in mind I return to (3), that upon having an experience as of my hand I must set $P_{new}(\neg BIV)$ equal to $P_{old}(\neg BIV|e)$. It's now clear that adopting P_{new} required two credence revisions: an exogenous revision in response to the experience, and the endogenous revision resulting from conditionalizing upon that exogenously revised credence. It's also clear that the Bayesian machinery constrains the endogenous revision but for the most part does not constrain the exogenous one, and that none of the Bayesian constraints (of either variety) require that the exogenous revision be on *I'm having an experience as of my hands* rather than on *I have hands* or on some other proposition.

single proposition then the resulting function will be incoherent. For example, if $P_{old}(A) = .7$ and P_{old} is coherent, then $P_{old}(\neg A) = .3$. If I exogenously revise my credence in A so that $P_{new}(A) = 1$ without updating then $P_{new}(\neg A) = .3$. Since A and $\neg A$ are inconsistent, $P_{new}(A \vee \neg A) = P_{new}(A) + P_{new}(\neg A)$, which in this case is 1.3.

²²As with Jeffrey Bayesianism this will depend upon how we understand the relationship between conditional and unconditional probabilities. In fn. 20 I noted that once the Ratio Analysis has been accepted, propositions assigned a credence of 1 or 0 are thereafter unrevisable. Since Strict Conditionalization requires assigning a credence of 1 to every proposition that is to be updated upon this condition is more often relevant for the Classical Bayesian than for the Jeffrey Bayesian. See §5.

If we suppose that it's my credence in e (and e alone) that I revise in light of my experience then Bayesianism ensures the truth of (3), but Bayesianism is simply silent about whether updating my credence in e is the right thing to do upon having my experience. Hence Bayesianism is silent concerning whether the credence function that I ought to adopt in light of having my experience, $P_{new}(\cdot)$ is equal to $P_{old}(\cdot|e)$. So the rejection of (3) is consistent with Bayesianism.

4.2 Dogmatists should update on h

Dogmatists claim that perceptual experience can generate immediate justification, but they also go further and specify precisely which proposition is immediately justified by an experience: the proposition constituting the content of that experience. So for the Dogmatist, when I have a perceptual experience as of A (and I lack defeaters and any other necessary conditions for obtaining perceptual justification are satisfied), I thereby obtain some immediate justification for believing A . Inferentialists deny that my justification for believing A is immediate, but that doesn't commit them to saying that no proposition is immediately justified by the experience. The Inferentialist thinks that obtaining justification for believing the content A of a perceptual experience requires justification for believing *I'm having an experience as of A* , and also justification for believing some auxiliary proposition such as *If I've had that experience as of A then, probably, A* . Though on that picture my justification for believing that A can't be immediate, presumably my justification for believing that I'm having an experience as of A is immediate. Hence the Dogmatist and the Inferentialist agree that my perceptual experience as of A generates at least some immediate justification, they just disagree about which proposition is immediately justified by that experience.

How is this talk of immediate justification to be translated into the Bayesian idiom of credences? One plausible thought is that my obtaining immediate justification for believing that A is tantamount rationally to increasing my credence in A without conditionalizing on something else in order to do so.²³ In other words, obtaining immediate justification for believing that A *just is* exogenously revising your credence upward in A in a rational way. Since the Dogmatist thinks that upon having an experience as of A I become immediately justified in believing that A , there's a strong prima facie case that a Bayesian Dogmatist should think that upon having that experience I should exogenously raise my credence

²³See Pryor (2013), especially his 'Assumption 2' on p. 105.

in A and then update upon it. Similarly, since the Inferentialist thinks that upon having an experience as of A I become immediately justified in believing that *I'm having an experience as of A*, a Bayesian Inferentialist should think that upon having that experience I should exogenously raise my confidence in *I'm having an experience as of A* and then update upon it.

As I've noted, (3) is not neutral concerning what we update upon, but requires that we update upon the proposition that happens to be what the Inferentialist thinks we're immediately justified in believing rather than the one that the Dogmatist thinks that we are immediately justified in believing. But that requirement begs the question against the Dogmatist, who should reject it even in without the putative problem that the Bayesian Argument brings to light.

Before moving on I'd like to briefly sketch an objection raised by Roger White in his (2006, p. 534-5). According to White, even if the Dogmatist is right and having an experience as of my hands provides immediate justification for believing $h = I \text{ have hands}$, it no doubt *also* provides immediate justification for believing *I'm having an experience as of my hands*, and hence I should also exogenously raise my credence in e. In that case Conditionalization requires that I update upon e. But if we're back to conditionalizing on e, then what does it matter if we also conditionalize on h? Won't updating on e raise my confidence that BIV, and hence even further limit my confidence in h? And in that case isn't the Dogmatist still stuck with the problematic conclusion at (7) after all?

No. The success of the Bayesian Argument does not depend on whether we update on e, but on whether we update on h. Allowing exogenous revisions to h means that my prior credence in BIV no longer limits my posterior credence in h, and hence the putatively problematic (7) is false. To see this point, however, it's helpful to first appreciate how updating on h solves the problem, and so I put off my full response to White's objection until §4.3.

4.3 How updating on h resolves the problem

Intuitively, the problem with learning that h by updating on e is that my posterior credence in h is limited by my prior credence in $\neg\text{BIV}$, and so if updating on e allows me to become highly confident in h then I must have started out highly confident in $\neg\text{BIV}$. In other words, when I update on e, my prior credence in $\neg\text{BIV}$ caps my posterior credence in h. This capping effect is not unique to

BIV, e, and h, or to the matter of perceptual justification. The relevant features of the case are that it's e being conditionalized upon, that $BIV \models e$, that $BIV \models \neg h$, and that $e \not\models h$ – the capping effect will be the same for any case meeting those conditions.²⁴

Updating on h instead of e changes this dynamic dramatically. I noted earlier that exogenous credence revisions are mostly unconstrained by the Bayesian machinery. Here's another point: they aren't constrained by the credence function that I accepted before my exogenous revision either.²⁵

Suppose that at t_1 I'm certain that $\neg(A\&B)$ and my credence in A is .7. My credence function at t_1 is coherent, and so $P_{t_1}(B) \leq 1 - P_{t_1}(A)$ or .3. Then at t_2 I exogenously revise my credence in B to 1, but I don't change any of my other credences. Now I'm incoherent: if $P_{t_2}(\neg(A\&B)) = 1$ and $P_{t_2}(B) = 1$, then $P_{t_2}(A)$ must be 0, but since I haven't changed that credence $P_{t_2}(A) = .7$.

At this point I could regain my coherence in at least three ways:²⁶ reduce my confidence in A to 0, reset my confidence in $\neg(A\&B)$ to .3, or reduce my confidence in B back to .3 or lower. If Bayesians were solely concerned with maintaining coherence — if they cared about satisfying Probabilism but they didn't care about satisfying Conditionalization — then any of those responses would do. But Bayesians do care about Conditionalization, and according to Conditionalization the three responses are importantly different. Conditionalization requires (i) that any incoherence introduced by an exogenous credence revision be resolved by changes elsewhere, never by reversing the exogenous revision, and (ii) that those changes be made (surprise!) by conditionalizing on the exogenously revised proposition. Hence the strategy of avoiding incoherence by reducing my exogenously revised credence in B back to .3 or less is not available to the Bayesian.

The point generalizes: in the Bayesian framework, any time there's a tension between an exogenously revised credence and an endogenously set credence, the tension is resolved by changing the endogenously set credence.²⁷

²⁴If I'm a handless brain in a vat having hand-like experiences (i.e. if BIV is true) then it's not the case that I have hands (so $\neg h$) and it is the case that I'm having hand-like experiences (so e). It's possible to have non-veridical hand-like experiences, and so $e \not\models h$.

²⁵Subject to the same qualification that I mentioned in footnotes 20 and 22, which should be assumed throughout this section.

²⁶Combinations of these three are also possible. Here I'm ignoring any incoherence arising from propositions other than those discussed.

²⁷For the Classical Bayesian who also accepts the Ratio Analysis of conditional probability, this will be true for any exogenously set credence, regardless of when it was set. By Strict Conditionalization all exogenous revisions result in a credence assignment of 1, which given the Ratio Analysis can never thereafter be revised downward (see §5). For the Jeffrey Bayesian

The relationship between BIV and h is very much like that of A and B above. Since BIV and h are logically inconsistent (remember: $BIV \models e \& \neg h$), Probabilism requires that $P(BIV \& h) = 0$. Let's stipulate that my prior credence in h is .3 and my prior credence in BIV is .7. Now suppose that I have an experience as of my hands, and as a result I exogenously revise $P_{new}(h)$ to .9, but I don't revise any of my other credences. At this point my credence function is incoherent. Resolving that incoherence requires that I revise at least one of $P_{new}(BIV \& h)$, $P_{new}(h)$, or $P_{new}(BIV)$. I can't change $P_{new}(BIV \& h)$, as every probabilistically credence function must set that value to 0. I can't endogenously revise $P_{new}(h)$ back to its previous value of .3, as its current value was set exogenously. Hence the only way that I can regain coherence is to reduce my confidence in BIV.²⁸

The lesson, then, is this: neither the formal commitments of Bayesianism nor the prior credence function of the agent meaningfully constrain exogenous credence revisions. Once an exogenous credence revision is made, coherence is preserved by changing other credences, never the exogenously set credence itself. Hence (i) there's no *formal* barrier to exogenously revising my credence in h , and (ii) my prior credence in BIV cannot constrain my exogenously set posterior credence in h : if there's a conflict between the two and a change must be made, the change must be made to my credence in BIV.

It's important not to interpret this conclusion too strongly. What I have shown is that the formal commitments of Bayesianism do not entail that my credence in h after having an experience as of my hands is limited by my prior credence in not BIV. What I have not shown, and what I do not believe to be true, is that facts about my epistemic state before I've had an experience as of my hands can never constrain the attitudes that I ought to adopt once I've had that experience. After all, if I'm very confident that I'm about to

who accept the Ratio Analysis it will be true of any credence exogenously set to 1, but past exogenous revisions with values less than 1 are subsequently revisable via conditionalization like any other credence. Current revisions, however, are not endogenously revisable, and so any incoherence that they introduce must be resolved by making adjustments elsewhere by conditionalization.

²⁸I'm not saying here that my low credence in BIV is justified by coherence alone. A classical foundationalist who obtains foundational justification for A might be required on pain of incoherence to believe that $\neg A$, but that doesn't make the latter belief justified solely by coherence considerations! Foundationalists think there's an asymmetry between foundationally justified beliefs and non-foundationally justified beliefs: when tension arises, the latter gives way to the former. Bayesians think something similar: there's an asymmetry between exogenously revised credences and prior credences that ensures that when tension arises the prior credences are changed to accommodate the exogenous revisions. Bayesianism is (among other things) a theory of probabilistic coherence, but it needn't be a coherentist theory of justification.

have a hallucinatory experience as of my hands, then it seems unreasonable to become much more confident that I do in fact have hands after having just such an experience. In such a case my prior credences seem to be defeating the evidential force of my experience, i.e. they are acting as an undercutting defeater. Sometimes prior attitudes do affect the epistemic force of experience.

So what determines whether my justification for believing a defeater at t_1 constrains my attitude at t_2 toward h , or whether at t_2 I should change my attitude toward that defeater in light of my new attitude toward h ? I'm not offering a positive account here, merely pointing out that the formal commitments of the Bayesian do not force an answer upon us. That formalism constrains only the credence revisions that it models.²⁹ No credence revision immediately resulting from experience is modeled, so a fortiori no credence revision immediately resulting from experience and affected by the agent's possession of a defeater is modeled. The point is simply that if the inputs to the Bayesian model are themselves defeasible (as they must be – see §5) then that defeat is an off-model phenomena and hence will not be constrained or explained by the Bayesian formalism. In other words, it's not that the credence that one ought to adopt in light of an experience is unconstrained by one's preexisting attitudes, but rather that the effects of those constraints are felt outside of the formal model.³⁰

We are now in a position to respond to the objection from Roger White that I sketched at the end of §4.2. White objected that even if having an experience as of my hands makes it permissible to exogenously raise my credence in h , it also makes it permissible to exogenously revise my credence in e . In that case I must then conditionalize on e , which means that (3) is true and hence the Dogmatist is still stuck with (7).

It should now be clear that this objection is misguided. Suppose that White is correct and that after having an experience as of my hands I should update on both h and on e , and for simplicity suppose that that's equivalent to updating on $h \& e$. (3) says that $P_{new}(\neg BIV) = P_{old}(\neg BIV|e)$, but in the case described that's false, as now $P_{new}(\neg BIV) = P_{old}(\neg BIV|e \& h)$. As we've seen, $e \& h \models \neg BIV$, and so $P_{new}(e \& h) \leq P_{new}(\neg BIV)$. And since $P_{new}(e \& h)$ was revised exogenously, if satisfying that inequality requires any further revision it will have to be to $P_{new}(\neg BIV)$.³¹

²⁹See fn. 20 and 22 for qualifications of this claim.

³⁰David Christensen has argued that permitting off-model defeat constitutes an unacceptable limitation on the explanatory ambitions of the Bayesian who accepts defeasible inputs to the model and hence poses a serious problem for Jeffrey Bayesianism; see his (1992).

³¹Raising my confidence in both e and h will not in every case decrease my confidence in

5 Varieties of Bayesianism

Let's take stock. Dogmatists think that upon having an experience with content A, I obtain immediate, underminable justification for believing that A. I've argued that for Bayesians, all credence revisions involve both an endogenous and an exogenous component, and hence that credence revisions due to perceptual experience also involve both an exogenous and an endogenous component. The most natural way to combine Bayesianism and Dogmatism is for the immediate effect of having a perceptual experience to be understood as the exogenous component of that process, i.e. that upon having an experience as of A, I exogenously revise my credence in A and update accordingly. This suggestion stands in contrast to premise (3) of the Bayesian Argument, which states that when I have a perceptual experience as of A the exogenous revision is to my credence in *I'm having an experience as of A*, with any revision in my credence in A coming as a result of updating on (i.e. endogenously revising on) that proposition. Hence given the modeling approach most natural to the Dogmatist, the Bayesian Argument includes a false premise and so is unsound.

In this section I consider an objection to the compatibility of my approach with Classical Bayesianism. As noted, Classical Bayesianism requires that all conditionalizations begin with an exogenous revision of a credence to 1, and hence if we are to make an exogenous revision to h then the Classical Bayesian will demand that $P_{new}(I \text{ have hands}) = 1$. An initial objection is that it's just implausible that we should assign the highest possible credence to that proposition. I think that this objection is correct, but only because it's implausible to think that we should assign the highest possible credence to any contingent proposition. If that's correct then the combination of my proposal with Classical Bayesianism is implausible simply because Classical Bayesianism on its own is implausible.³² Ultimately this point is moot, however, as there's a

\neg BIV. If my prior credence in h is much higher than my credence in e, and then I increase my credence in h only slightly while my credence in e increases dramatically, then my credence in \neg BIV can increase. A conjunction can't be any more probable than its least probable conjunct, so if $P(A) = .7$ and $P(B) = .1$, then $P(A\&B) \leq .1$. In that case it might be the low value for P(B) that's limiting P(A&B), and so if P(B) increases to .3, then even if P(A) decreases slightly to .6 then P(A&B) can now be as high as .3. As I noted at the beginning of this section, the relevance of BIV to the Bayesian argument is that it entails the conjunction $e\&\neg h$. Hence we see that it's possible for a decrease in $P(\neg h)$ together with an increase in $P(e)$ to lead to an increase in $P(e\&\neg h)$, i.e. in $P(\text{BIV})$. But this is all beside the point, which is that my prior credence in \neg BIV does not constrain the credence that I exogenously assign to h, regardless of any revision to e.

³²See chapter 11 of Richard Jeffrey's (1983) for more on the implausible strictness of Strict Conditionalization.

much more serious objection to the combination of my proposal with Classical Bayesianism.

In addition to Probabilism and some version of Conditionalization, it is common for Bayesians to accept the Ratio Analysis³³ of conditional probability:

Ratio Analysis: $P(A | B) = \frac{P(A \& B)}{P(B)}$

Two implications of accepting the Ratio Analysis will be important for the present discussion. First, it becomes impossible to conditionalize upon propositions assigned a credence of zero: if $P(B) = 0$ then $\frac{P(A \& B)}{P(B)}$ is undefined, and hence by the Ratio Analysis $P(A|B)$ is undefined as well. Second, if a proposition is once assigned a credence of either 1 or 0 then it's impossible to revise that credence endogenously. If $P(A) = 1$, then $\frac{P(A \& B)}{P(B)} = \frac{P(B)}{P(B)} = 1$, meaning that for any proposition B such that $P(B) > 0$, $P(A|B) = 1$. Similarly, if $P(A) = 0$ then $P(A \& B) = 0$, and so for any proposition B such that $P(B) > 0$, $P(A|B) = 0$.

Now consider what I'll call Classical Bayesianism*, which in addition to Probabilism and Strict Conditionalization also accepts the Ratio Analysis of conditional probability. Since Classical Bayesianism* accepts Strict Conditionalization, in order to update on h I must first exogenously revise $P(h)$ to 1. The Dogmatist is committed to the underminability of h by new information and so it must be possible to decrease my credence in h, but given Strict Conditionalization and the Ratio Analysis that's impossible. The lesson is that the Classical Bayesian* can't simultaneously hold that (i) we should update on what we're immediately justified in believing, (ii) upon having a perceptual experience as of h I obtain some immediate justification for believing that h, and (iii) my justification for h is underminable. The Dogmatist is committed to (ii) and (iii), and my suggestion is that we accept (i), so my response to the Bayesian Argument is unavailable to the Dogmatist who is also a Classical Bayesian*.

There's a sense in which this shouldn't much bother the Dogmatist. The conflict arises from the fact that that the inputs to Classical Bayesian* models must have a credence of 1, which by the Ratio Analysis can't ever be reduced endogenously from 1 to any other value. But what possible inputs to the model are immune to rational revision? The suggestion of the Inferentialist is that we

³³For our purposes, whether the Ratio Analysis really is an *analysis* — whether one side of the equation is more basic than the other — is irrelevant: what matters is whether the equation holds without exception. Of course if the equation is understood as an analysis of conditional probability in terms of the more basic unconditional probability, or if it is taken as a stipulative definition of '|', then the equation will hold without exception.

update on facts about our own experiences. Grant for the sake of argument that we have perfect access to the content of our own perceptual experiences³⁴ and hence that it's reasonable to assign a credence of 1 to the proposition *I'm having an experience as of my hands* upon having an experience as of your hands. Nonetheless, there are two fairly obvious reasons that you might revise that credence over time without thereby becoming irrational. First, you might forget what experiences you've had. Does anyone really remember each and every experience that they've ever had and updated upon, and remember them with absolute certainty? And if not, is such forgetting a failure of rationality, as the Classical Bayesian* is committed to saying?

The second reason that you might rationally decrease a credence from 1 is that you acquire a good reason to believe that that credence is inappropriate, i.e. you obtain a defeater. Suppose that the philosophy oracle tells you that although you seem to remember having a visual experience as of your hands at 11:21PM on June the 4th, 1986, you're misremembering – you actually had an experience as of a photograph of your identical twin's hands on that date. The lesser philosophy oracles all agree. For the first time ever the philosophy oracles are all mistaken at the same time: you are remembering correctly. Still, are you really really to disregard their testimony, or in light of that testimony should you lower your credence that you had that experience?³⁵

Classical Bayesianism* suffers from what I'll call the Problem of Invincible Certainty: once a proposition is assigned a superlative credence (1 or 0) it becomes impossible to revise that credence via conditionalization. Strict Conditionalization requires that inputs to the formal model — the propositions updated upon — must have a credence of 1, and hence Classical Bayesianism* also suffers from what I'll call Problem of Invincible Evidence: once a proposition is updated upon it's impossible to revise your credence in that proposition.

Evidence is not invincible: revising downward our confidence in some proposition that we once updated upon, whether due to forgetting the acquisition of a defeater or for some other reason, is both common and rationally permissible (and sometimes even obligatory). It follows that for any such proposition A, either $P(A)$ was less than 1 when I updated upon it, in which case Strict

³⁴A highly questionable assumption: if you concede even the mere possibility that the content of perceptual experience is determined by something external to our own minds then you should think that there's a nonzero probability that you are wrong about the content of at least some of your experiences.

³⁵Timothy Williamson raises similar objections to the unrevisable inputs to Classical Bayesian models in his (2000) p. 203-7.

Conditionalization is false, or its credence when I conditionalized really was 1 but we allow that credence to be subsequently revised downward, in which case either the Ratio Analysis is false or Strict Conditionalization is false. Either way one of the commitments of Classical Bayesianism* is false, and so Classical Bayesianism* is false.

Let's consider what happens if we retain Strict Conditionalization and give up on the Ratio Analysis of conditional probability, an approach associated with Alan Hájek.³⁶ The idea here is to accept the equation of $P(A|B)$ with $\frac{P(A\&B)}{P(B)}$ in all instances in which $P(B) > 0$ and to reject it otherwise, so that when $P(B) = 0$, $P(A|B)$ is independent of the degree of correlation between A and B (since when $P(B) = 0$ there isn't any correlation).

Giving up on the Ratio Analysis of conditional probability offers a very limited defense to the Problem of Invincible Evidence as it is now possible reduce my credence in A after once having updated upon it. Consider some proposition A that I've updated upon at some point in the past. Since we're supposing Strict conditionalization, I must have assigned a credence of 1 to A when I updated upon it. As we've seen, that means that for any proposition B such that $P(B) \geq 0$, $P(A|B) = \frac{P(A\&B)}{P(B)} = 1$,³⁷ and so it's impossible to reduce my credence in A by updating on B. But if we update on some proposition C such that $P(C) = 0$, we are freed from the constraints of the Ratio Analysis and so there's no formal barrier to assigning $P(A|C)$ a value less than 1. Hence for any proposition A that we've previously updated upon and hence become certain is true, we can back away from that certainty only by becoming certain of the truth of some proposition C, which we formerly regarded as being certainly false, and updating accordingly.

Independent of one's views on Dogmatism, this is not an appealing response to the Problem of Invincible Evidence. Nonetheless, things are even worse for the Dogmatist who accepts my modeling proposal due to the logical relations between h and BIV. Intuitively, even after having an experience as of my hands and updating accordingly it should be possible to increase my confidence in BIV and on those grounds decrease my confidence in h. But if I've updated on h and hence set $P(h) = 1$ then by conditionalization I will also have set my credence

³⁶See Hájek (2003). Hájek actually wants to give up on the Ratio Analysis altogether, as he thinks that conditional credences are primitive and so cannot be analyzed into unconditional probabilities or anything else. For our purposes the important issue is not which side of the equation is considered primitive, but instead whether that equation actually holds.

³⁷Because we've only rejected the Ratio Analysis in cases where my credence in the proposition being updated upon is greater than zero.

in every proposition inconsistent with h — including BIV — to 0. But now how can I increase my credence in BIV from 0? As with all credence revisions, that revision will be either exogenous or endogenous. Once I've set $P(\text{BIV})$ to 0, the only way to revise that value endogenously is to update on some other proposition with a credence of 0.³⁸ So for example if a very reliable source were to tell me that I'm a brain in a vat after all then I should at least slightly raise my confidence in BIV, but on the current proposal that's only possible if I assign a credence of zero to my obtaining that testimony.³⁹

Can $P(\text{BIV})$ be revised exogenously? That would be consistent with the formalism, though I won't comment on its plausibility. However, the problem that I will mention is that given Strict Conditionalization, exogenously revising my credence in BIV means assigning it a credence of 1, which means that upon conditionalizing I must now revise my credence in h all the way back down to 0, which is not what's wanted in many cases of undermining.⁴⁰

For these reasons the Classical Bayesian, whether or not they accept the Ratio Analysis of conditional probability, has no plausible response to the Problem of Invincible Evidence. This problem is if anything exacerbated by my proposal that we update upon the contents of our experiences. Hence Dogmatism, Classical Bayesianism, and updating on the contents of experience is an unappealing combination, but that's mostly because Classical Bayesianism is unappealing.

Though Classical Bayesianism and Jeffrey Bayesianism are equally susceptible to the Problem of Invincible Certainty,⁴¹ Jeffrey Bayesianism fares much better with the Problem of Invincible Evidence. For the Bayesian, evidence is only invincible when it's certain. Given Strict Conditionalization, all evidence is certain and so all evidence is invincible (ignoring the possibility of updating on $P(\cdot) = 0$ propositions). Jeffrey Conditionalization allows updates on

³⁸Because we've only rejected the Ratio Analysis in cases in which the proposition being updated upon is assigned a credence of 0. If $P(\text{BIV}) = 0$ and $P(A) > 0$ then we're still committed to saying that $P(\text{BIV}|A) = \frac{P(\text{BIV} \& A)}{P(A)} = 0$. Hence $P(\text{BIV}|A) \geq 0$ is only possible if $P(A) = 0$.

³⁹Though I've been considering whether Hájek's proposal of abandoning the Ratio Analysis offers a solution to the Problem of Invincible Evidence, Hájek himself was not motivated by that problem. Hájek's objection to the Ratio Analysis is that it makes it impossible to update on propositions assigned a credence of 0. I'm sympathetic — I too “hold this truth to be self-evident: the conditional probability of any (non-empty) proposition, given itself, is 1” (Hájek, 2003, p. 286) — so the criticisms in this section should not be interpreted as criticisms of Hájek's proposal.

⁴⁰Plausibly when my justification for believing that h is undermined my credence in h returns to whatever it was before I obtained my now-defeated justification. That value will frequently be greater than 0.

⁴¹As are Classical Bayesianism* and Jeffrey Bayesianism* (= Probabilism + Jeffrey Conditionalization + the Ratio Analysis of conditional probability).

propositions that aren't certain, and so evidence needn't be invincible.

The combination of Jeffrey Bayesianism and my proposal that we update on the contents of our experience is very appealing. It allows the Dogmatist to retain the core commitments of Bayesianism (Probabilism, a version of Conditionalization, and the Ratio Analysis (if desired)) while avoiding the problematic conclusion of the Bayesian Argument.

6 Conclusion

The conclusion of the Bayesian argument has always been somewhat surprising. Typically when two theories conflict it's because they offer inconsistent accounts of the same explanandum. Bayesianism and Dogmatism seek to account for different aspects of rationality: respectively, the coherence of partial belief states and the appropriate response to perceptual experience. Hence there is no single explanandum common to both theories. As I've argued, bringing those views into conflict requires an auxiliary account of how they come into contact in the first place.

The success of the Bayesian argument requires a very specific thesis about this point of contact between Dogmatism and Bayesianism: that upon having an experience as of A the agent should exogenously revise their credence in *I've had an experience as of A*, with any revision to their credence in A itself proceeding via conditionalization. For Inferentialists that's a very natural way to model perceptual learning, as it makes explicit their view that perceptual justification for A is inferentially dependent on agent's having justification for believing propositions about their own mental states. But Dogmatists reject that inferential picture of perceptual justification, claiming instead that an experience as of A can provide immediate justification for A. Hence the Dogmatist should view the Inferentialist's modeling proposal as both inaccurate and prejudicial.

In short, the Bayesian Argument together with the Inferentialist's approach to modeling begs the question the against the Dogmatist, and the Bayesian Argument without that approach to modeling is unsound. Either way, the argument provides no reason to reject Dogmatism. The upshot of these considerations is an attractive view combining Dogmatism and Jeffrey Bayesianism, on which the epistemic impact of a perceptual experience is incorporated into the model by making rational an exogenous credence revision to the content of that experience.

References

- Christensen, D. (1992). Confirmational Holism and Bayesian Epistemology. *Philosophy of Science*, 59(4):540–557.
- Cohen, S. (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, 65(2):309–329.
- Davidson, D. (1986). A Coherence Theory of Truth and Knowledge. In LePore, E., editor, *Truth and Interpretation. Perspectives on the Philosophy of Donald Davidson*, pages 307–319. Basil Blackwell.
- Davidson, D. (2001). *Essays on Actions and Events: Philosophical Essays Volume 1*. Clarendon Press.
- Hájek, A. (2003). What Conditional Probability Could Not Be. *Synthese*, 137(3):273–323.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford University Press.
- Howson, C. and Urbach, P. (1993). *Scientific Reasoning: The Bayesian Approach*. Open Court.
- Jeffrey, R. (1983). *The Logic of Decision*. University of Chicago Press.
- McGrath, M. (2013). Dogmatism, Underminers and Skepticism. *Philosophy and Phenomenological Research*, 86(3):533–562.
- Moretti, L. (2015). In Defence of Dogmatism. *Philosophical Studies*, 172(1):261–282.
- Pollock, J. and Cruz, J. (1999). *Contemporary Theories of Knowledge*. Rowman & Littlefield.
- Pryor, J. (2000). The Skeptic and the Dogmatist. *Noûs*, 34(4):517–549.
- Pryor, J. (2005). There is Immediate Justification. In Steup, M. and Sosa, E., editors, *Contemporary Debates in Epistemology*, pages 181–202. Blackwell.
- Pryor, J. (2013). Problems for Credulism. In Tucker, C., editor, *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford University Press.
- Russell, B. (1997). *The Problems of Philosophy*. Oxford University Press.

- Schiffer, S. (2004). Skepticism and the Vagaries of Justified Belief. *Philosophical Studies*, 119(1-2):161–184.
- Silins, N. (2007). Basic Justification and the Moorean Response to the Skeptic. *Oxford Studies in Epistemology: Volume 2*, 2:108.
- Vogel, J. (1990). Cartesian Skepticism and Inference to the Best Explanation. *Journal of Philosophy*, 87(11):658–666.
- Weatherson, B. (2007). The Bayesian and the Dogmatist. *Proceedings of the Aristotelian Society*, 107(1pt2):169–185.
- White, R. (2006). Problems for Dogmatism. *Philosophical Studies*, 131(3):525–57.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Wright, C. (2004). Warrant for Nothing (and Foundations for Free)? *Aristotelian Society Supplementary Volume*, 78(1):167–212.
- Zardini, E. (2014). Confrming the Less Likely, Discovering the Unknown: Dogmatisms - Surd and Doubly Surd, Natural, Flat and Doubly Flat. In Dylan Dodd, E. Z., editor, *Scepticism & Perceptual Justification*, pages 33–70. Oxford University Press.