



How to Be a Bayesian Dogmatist

Brian T. Miller

University of Texas, Austin

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 jointly 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 jointly consistent. That leaves the matter of how experiences and credences 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 into the Bayesian formalism.

ARTICLE HISTORY Received 20 March 2015; Revised article received 14 December 2015

KEYWORDS Bayesianism; Dogmatism; epistemology; justification; perception; scepticism

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 if there's 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 partly 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 if there's a red ball lying on the grass* comes directly from the experience itself without the mediation of some other justification that I have,

CONTACT Brian T. Miller  briantmiller@utexas.edu

and hence that justification is immediate. That much is 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 underminable² [Pryor 2000, 2005, 2013]. Moreover, the Dogmatist thinks that, while a perceptual experience may generate immediate and underminable justification for *I'm having an experience as if A* or some other proposition about the agent's mental states, it also generates immediate and underminable justification for *A* 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: for example, *if I have a perceptual experience as if A then, probably, A*. Hence, it's my justification for believing *I've had an experience as if there's a red ball lying on the grass* together with 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.

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, or instead a shortcoming, of the theory is contentious: easily acquired justification for propositions about the external world might be thought to license too-easy responses to sceptical challenges to our knowledge of 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 sceptical problems insoluble.

In this essay, I defend Dogmatism against a very different objection—namely, 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 *there's a red ball lying on the grass* requires that I already have justification for ruling out a wide range of sceptical scenarios on which my experience as of the ball lying on the grass is non-veridical. 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

¹My justification for believing that *A* is *immediate* unless it is in part my having justification to believe something else that *makes me* justified in believing that *A*. '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.

²For the distinction between undermining/undercutting and opposing/rebutting defeaters, see Pollock and Cruz [1999: 196–7].

justification for believing that *B* isn't immediate. Since this result allegedly follows from the Bayesian formalism, we thereby have some reason to believe that Dogmatism and Bayesianism are jointly 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 joint inconsistency of Dogmatism and Bayesianism and to take that as a good reason to revise orthodox Bayesianism [Weatherson 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 having 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 to 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.³

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 experience should be reflected in the formal model. This requires that I be clear about exactly what the Bayesian is and is not committed to, an issue that I discuss in section 2. In section 3, I lay out the formal details of the Bayesian Argument. The heart of the paper is found in section 4, where I identify the problematic premise and argue both that it is optional for the Bayesian and that it is prejudicial against the Dogmatist. I then offer an alternative account of how the epistemic impact of experience should be incorporated into Bayesian models. In section 5, I consider the implications of adopting my suggestion for various versions of Bayesianism, concluding that the Dogmatist should embrace a version that incorporates Richard Jeffrey's permissive approach to conditionalization over the strict approach of the Classical Bayesian.

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. The core of Bayesianism is the postulation of two necessary conditions on the rationality of a credence. The first is Probabilism:

Probabilism. All rationally permissible credence functions are probabilistically coherent (that is, they satisfy the probability axioms).

³A 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 *A* always leads to an increased credence *A*. See Kung [2010] and Zardini [2014].

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 considered on its own, 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, $P(\text{the street is wet}) < P(\text{the street is wet} \mid \text{it's raining})$. 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 or she holds right now, and that those constraints are encoded in the agent's currently held conditional credences. I'll be discussing two ways of making this intuition rigorous. First:

Strict Conditionalization. If I revise my credence in B to 1, then I must set my new credence in A equal to my old credence in A conditional on B :

$$P_{\text{new}}(A) = P_{\text{old}}(A \mid B)$$

It is important to note that, according to Strict Conditionalization, incorporating new information B by conditionalizing upon it requires changing $P(B)$ to 1.⁴ Jeffrey Conditionalization generalizes Strict Conditionalization by allowing updates upon changes in credences to values other than 1:

Jeffrey Conditionalization. If I revise my credence in B to any value, then I must set my new credence in A equal to the weighted sum of A conditional on B and A conditional on $\neg B$.⁵

$$P_{\text{new}}(A) = P_{\text{old}}(A \mid B)P_{\text{new}}(B) + P_{\text{old}}(A \mid \neg B)P_{\text{new}}(\neg B)$$

For our purposes, Classical Bayesianism is the combination of Probabilism and Strict Conditionalization, and Jeffrey Bayesianism is the combination of Probabilism and Jeffrey Conditionalization⁶. Since my task in this essay is to show that the Bayesian Argument exposes no great tension between Dogmatism and either version of Bayesianism, I will proceed to show that the argument exposes no great tension between Dogmatism and the combination of Probabilism and either version of Conditionalization (although later on I settle on Jeffrey Bayesianism as the better complement to Dogmatism).

3. The Bayesian Argument against Dogmatism

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. To illustrate, consider *BIV*, the hypothesis that I'm a handless brain in a vat having experiences

⁴For reasons that I discuss in section 4.1, especially note 13, I'll assume throughout the essay that proponents of Strict Conditionalization will prohibit 'exogenous' credence revisions (again, see section 4.1) to values less than 1.

⁵See Jeffrey [1983: 169]. He goes on to generalize this condition to accommodate changes to the partition involving more than two propositions, a complication inessential to the present essay.

⁶Let me emphasize that, by 'Jeffrey Bayesianism', I mean only the combination of Probabilism and Jeffrey Conditionalization (as defined above).

as of my hands (that is, the sorts of experience that we expect to have when we look down to and past the ends of our wrists). Dogmatists hold that if I have high levels of justification for believing that *BIV* is true then my experience as of my hands will fail to generate much justification for the proposition *I have hands*, as the justificatory force of the experience is undermined. But Dogmatists also hold that I don't need justification for believing that *BIV* is false in order to acquire justification for *I have hands* from my experience. It's this last point that is the target of the Bayesian Argument.

BIV is the hypothesis that I'm a handless brain in a vat having experiences as of my hands, and so *BIV* implies that I'm having an experience as of my hands. Taking *e* as the proposition *I'm having an experience as if I have hands*, that means the following⁷:

$$(1) P_{old}(\neg BIV|e) \leq P_{old}(\neg BIV)$$

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*,⁸ and so we have this:

$$(2) P_{new}(\neg BIV) = P_{old}(\neg BIV|e)^9$$

⁷Proof: $BIV \models e$, so $P(e|BIV) = 1 \geq P(e)$, so $P(BIV|e) \geq P(BIV)$ (by Bayes's Theorem), so $P(\neg BIV|e) \leq P(\neg BIV)$.

⁸While this is correct as far as it goes—Conditionalization does indeed require that we update upon changes to our credences in propositions like *e*—I will argue in section 4 that updating upon *e* alone is both unwarranted and vital to the argument. Nonetheless, at this point I'll suppose that it is correct, in order to present my opponent's argument.

⁹In particular, (2) is meant to follow from *Strict* Conditionalization (plus the description of the case). The argument is slightly different when Jeffrey Conditionalization is employed, as (2) now becomes this:

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

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

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

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 this:

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

$e \& \neg h$ 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 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 all of that together, and we get this:

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

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 $\neg BIV$.

Combining terms from (1) and (2), we get (3):

$$(3) P_{new}(\neg BIV) \leq P_{old}(\neg BIV)$$

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*, this follows:

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

Finally, (3) and (4) imply (5):

$$(5) P_{new}(h) \leq P_{old}(\neg BIV)$$

What (5) 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 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.¹⁰ Analogous arguments show that no perceptual justification is immediate, and so Dogmatism is false.

4. Modelling Experience

The Bayesian Argument is valid. (1) and (4) follow from Probabilism, the Ratio Analysis of conditional probability (see [section 5](#)), and the logical relations that obtain between BIV , h , and e . (3) is a consequence of (1) and (2). But what about (2)—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 (2). If (2) is false then the argument for (5) is unsound, and so the putative tension between Dogmatism and Bayesianism is resolved.

But how can we reject (2) without rejecting Conditionalization? Note that there are two importantly different ways of thinking about credence function $P_{new}(\cdot)$, and so two importantly different ways of thinking about (2). $P_{new}(\cdot)$ might be understood simply as the credence function resulting from having $P_{old}(\cdot)$ and then updating on e , and in that case (2) is a trivial consequence of Conditionalization. Alternately, $P_{new}(\cdot)$ might be understood only as *the credence function that an agent who holds P_{old} ought to adopt after having an experience as of his hands and updating accordingly*, whatever that function happens to be.

¹⁰That the existence of such a necessary condition is inconsistent with the immediacy of perceptual justification is far from obvious, as I discussed in note 1. However, if we rely on that point to respond to the Bayesian Argument then we are essentially denying that (5) 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 sceptical hypotheses. I find that implausible, and so in what follows I offer a response that allows the rejection of (5) without requiring the rejection of Bayesianism itself.

This is important because, in order to avoid equivocation, $P_{new}(\cdot)$ must be interpreted in the same way within (2) and within (5). If $P_{new}(\cdot)$ is interpreted in the first of these two ways, then all that (5) says is that updating on e (and e alone) can't raise $P_{new}(h)$ any higher than $P_{old}(\neg BIV)$. But Dogmatism isn't a theory of how an agent with P_{old} who updates on e ought to revise his beliefs; it's a theory of how an agent with P_{old} who has an experience as if he has hands ought to revise his beliefs. Hence, in order for (5) to pose a challenge to Dogmatism, $P_{new}(\cdot)$ must be understood as the credence function that an agent who holds P_{old} ought to adopt after having an experience as of his hands and updating accordingly.

The two interpretations of $P_{new}(\cdot)$ aren't necessarily jointly inconsistent: it may be that the credence function that an agent ought to adopt upon having an experience as if h just is the one that results from updating on e and on e alone. If so, then the two interpretations are equivalent; but that's a substantive assumption, and as I argue below the Dogmatist has independent reasons to reject it. After all, Dogmatists think that my experience as if h provides immediate justification not just for e , but for h as well. In that case, it's just false that $P_{new}(\cdot) = P_{old}(\cdot|e)$, and so the interpretation of $P_{new}(\cdot)$ on which (2) is a trivial consequence of Conditionalization is distinct from the interpretation that makes (5) problematic for Dogmatism.

My proposal, then, is that, instead of (or in addition to) updating on $e = I'm$ *having an experience as if I have hands*, Bayesian Dogmatists should update on $h = I$ *have hands*. This is consistent with both Classical and Jeffrey Bayesianism (see [section 4.1](#)), although 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 (and not merely upon e), not only allows the Dogmatist to avoid (5), but also provides a very natural way for the Dogmatist to model perceptual learning in a Bayesian framework.¹¹

4.1 Bayesianism Does Not Entail (2)

I begin by showing that updating on h and hence rejecting (2) 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.

I described in [section 2](#) how Bayesians construct formal models of agents' partial belief states and of revisions to those states over time. It's important to note that these are at best *partial* models of rational credence revision. Given Conditionalization, a prior credence function plus a revised credence in some proposition completely determine the posterior credence function that must be adopted: if I revise my credence in B

¹¹Compare Pryor [2013: sec. 6]. Since 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 (i.e. the Bayesian 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 outside the model (see [section 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.

to 1, then the model determines that I should revise my credence in A to my prior credence in A conditional on B . What the model does not determine is the rational status of my initial revision to my credence in B .

What does determine the rational status of the credence revisions that spark conditionalization? 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 it 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 modelled by the Bayesian formalism are conditionalizations, so it follows that some rationally permissible credence revisions are not modelled. Call those credence revisions that are not modelled by the Bayesian *exogenous* revisions (as in: *exogenous to the model*) and those occurring within the model via conditionalization *endogenous* revisions.

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 so it prohibits exogenous revisions that are themselves probabilistically incoherent. For example, I cannot revise my credence in $A \& \neg A$ to any value other than 0. Jeffrey Conditionalization imposes no additional 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 (and assuming that the credence was strictly between 0 and 1 before the revision).

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,¹² so there's good reason for the Classical Bayesian to prohibit exogenous revisions that cannot be updated upon—that is, to prohibit exogenous revisions to credences other than 1. Nonetheless, any 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 (again assuming that the credence in that credence was strictly between 0 and 1 before the revision).

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 that the credence revisions that result from having a perceptual experience are entirely endogenous and so proceed by conditionalization. As we've seen, conditionalization results from a change in credence or subjective probability: I update on *it's raining*, not when it's actually true that

¹²If my credence function started out as coherent and I exogenously revise my credence in a 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 jointly inconsistent, $P_{new}(A \vee \neg A) = P_{new}(A) + P_{new}(\neg A)$, which in this case is 1.3.

it's raining, 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 (2)—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 an 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 neither Probabilism nor Conditionalization requires that the exogenous revision be on *I'm having an experience as if I have hands* rather than on *I have hands* or on some other proposition.

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 (2); but Bayesianism is simply silent about whether updating my credence in e is the right way to respond to 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 (2) 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—namely, the proposition constituting the content of that experience. So, for the Dogmatist, a perceptual experience as of A typically generates immediate justification for believing A . Inferentialists deny that my justification for believing A is immediate, but this 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 if A* , and also justification for believing some auxiliary proposition such as *if I've had an experience as if A then, probably, A* . Although 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 it immediately justifies.

How is this talk of immediate justification to be translated into the Bayesian idiom of credences? One thought is that my obtaining immediate justification for believing that A is tantamount to rationally increasing my credence in A without conditionalizing on something else in order to do so.¹³ In other words, obtaining immediate justification

¹³See Pryor's 'Assumption 2' [2013: 105].

for believing that *A just is* exogenously raising one's credence 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 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, (2) is not neutral concerning what we update upon: it requires that I update on *facts about* my experience rather than on the *content* of my experience. But that requirement begs the question against the Dogmatist, who should reject it even 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 [2006: 534–5]. According to him, even if 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 $e = I'm \text{ having an experience as if I have hands}$. Hence, if my proposal is accepted and we update on whatever we become immediately justified in believing, we'll update on both h and e . But if I'm also conditionalizing upon e , then won't I still end up increasing my confidence in *BIV* and thereby limiting even further my posterior confidence in h ? And in that case isn't the Dogmatist still stuck with the problematic conclusion at (5), after all?

No. The success of the Bayesian Argument depends not on whether we update on e , but instead 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 (5) is false. To see this point, however, it's helpful to first appreciate exactly how updating on h solves the problem, and so I delay my full response to White's objection until the end of the next section.

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 $\neg BIV$, and so if updating on e allows me to become highly confident in h then I must have started out as highly confident in $\neg BIV$. In other words, when I update on e , my prior credence in $\neg BIV$ 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 alone that's being conditionalized upon, that $BIV \models e$, that $BIV \models \neg h$, and that $\neg(e \models h)$ —the capping effect will be the same for any case meeting those conditions.

The situation changes dramatically when we also update upon h . Since h and *BIV* are jointly inconsistent, Probabilism requires that $P_{old}(BIV|h) = 0$. If, upon having an experience as of my hands, I strictly conditionalize only on h , then $P_{new}(\neg BIV) = P_{old}(BIV|h)$, and so $P_{new}(BIV)$ must be 0 as well. In other words, if my experience makes it rational to exogenously revise my credence in h to 1, then I'm forced to become maximally confident that I'm not a handless brain in a vat having hand-like experiences. Hence, even though I can't be any more confident in h than I am in $\neg BIV$ —after all, (4) is a theorem of the probability calculus—my credence

in $\neg BIV$ is already as high as it can go, and so my new credence in h is not in any way constrained.

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 $\neg BIV$. 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.

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, and no credence revision immediately resulting from experience is modelled. Hence, no credence revision immediately resulting from experience and affected by the agent's possession of a defeater is modelled. The point is simply that if the inputs to the Bayesian model are themselves defeasible then that defeat is an off-model phenomenon 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; rather, it's that the effects of those constraints are felt outside the formal model.¹⁴

We are now in a position to respond to the objection from White that I sketched at the end of section 4.2. He 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 which case e must be among the propositions that I conditionalize upon. The idea seems to be that this implies (2)—that $P_{new}(\neg BIV) = P_{old}(\neg BIV|e)$ —and hence the Dogmatist is still stuck with (5).

White is no doubt correct that having an experience as of my hands provides me with immediate justification for believing that I've had an experience as if I have hands, and so e should be *among* the propositions that I update upon. It's also true that my posterior credence in h will be capped by my prior credence in $\neg BIV$ if e is the *only* proposition that I conditionalize upon in response to my experience. I've proposed that agents should update on whatever propositions are immediately justified by their experience, and Dogmatists think that h is one of those propositions. Hence, the view to which White is objecting—Bayesian Dogmatism, with my proposal for what to update upon—is one on which the agent ought to update both on e and on h . For the Classical Bayesian, updating on e and also on h is equivalent to updating on $e\&h$,¹⁵ and so on

¹⁴David Christensen [1992] 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.

¹⁵Due to the commutativity of Strict Conditionalization, the order of update does not affect the credence function that's ultimately adopted [Weisberg 2009].

this proposal $P_{new}(\neg BIV) = P_{old}(\neg BIV|e\&h)$. Since $e\&h$ implies $\neg BIV$, $P_{old}(\neg BIV|e\&h)$ and $P_{new}(\neg BIV)$ must both be 1, and so updating on $e\&h$ solves the problem for the Classical Bayesian, for exactly the same reason that updating on h alone solves the problem.¹⁶

So where, exactly, does the objection go wrong? The assumption seems to be that updating on e is a sufficient condition for producing the capping effect that I discussed above—namely, that it commits the Dogmatist to premise (2) and hence to conclusion (5). But that’s just wrong: updating on h in addition to e implies that (2) is actually false (assuming that $P_{old}(\neg BIV) > 0$), as now $P_{new}(\neg BIV)$ should be set to $P_{old}(\neg BIV|e\&h)$ rather than to $P_{old}(\neg BIV|e)$.

5. Varieties of Bayesianism

Let’s take stock. Dogmatists are committed to two theses about perceptual justification: that it’s immediate, and that it’s underminable. Bayesians are committed to two theses of their own: Probabilism and Conditionalization. The point of the Bayesian Argument is to show that Bayesianism is inconsistent with the Dogmatist’s immediacy thesis by showing that an experience as if h only generates a high posterior credence in h when the agent has a high prior credence in $e\&\neg h$. I’ve argued that this conclusion follows only with the additional thesis that, upon having an experience as of h , Bayesian agents ought to update on e and e alone. But updating on h instead of e (or in addition to e) is both consistent with Bayesianism and much more natural for the Dogmatist, and hence the Bayesian Argument is unsound.

One potential objection is that, although adopting my proposal resolves the apparent conflict between Bayesianism and the *immediacy* of perceptual justification, it appears to create a new conflict between Bayesianism and the *underminability* of perceptual justification.¹⁷ To see the problem, first consider Classical Bayesianism*, the conjunction of Probabilism, Strict Conditionalization, and the Ratio Analysis of conditional probability (we’ll back off that assumption in a moment):

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

The Ratio Analysis is important to the present discussion because it commits the Bayesian to invincible certainty: 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 $P(A|B) = \frac{P(A\&B)}{P(B)} = \frac{P(B)}{P(B)}$, meaning that, for any proposition B such that $P(B) > 0$, $P(A|B) = 1$. Similarly, if $P(A) = 0$ then $P(A|B) = \frac{P(A\&B)}{P(B)} = \frac{0}{P(B)}$, and so, for any proposition B such that $P(B) > 0$, $P(A|B) = 0$.

Now for the problem. Since Classical Bayesianism* accepts Strict Conditionalization, if that view is correct then in order to update on h I must first exogenously revise $P(h)$

¹⁶For the Jeffery Bayesian, updating on e and on h needn’t be equivalent to updating on $e\&h$, but that conjunction will be an element of the partition updated upon, as will $e\&\neg h$, $\neg e\&h$, and $\neg e\&\neg h$. Conditionalizing on this partition involves assigning credences to each of its elements, and $P(h) = P(h|e\&h) + P(h|\neg e\&h)$, so weighting this partition determines the posterior credence of h . Finally, since h is inconsistent with BIV , it follows that $P_{new}(BIV)$ can’t be any higher than $1 - P_{new}(h)$, and so a high posterior credence in h results in a low posterior credence in BIV , regardless of the prior credence in BIV . For general remarks on this approach, see Jeffrey [1983: 173].

¹⁷Thanks to Miriam Schoenfeld for pressing this objection.

to 1. The Dogmatist is committed to the underminability of my credence in h , and so it must be possible to decrease that credence; 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*.

What options are available to the Bayesian Dogmatist at this point? Any Bayesian who accepts the Ratio Analysis is committed to the invincibility of certainty—that boundary credences (0 or 1) can never be revised endogenously. The Classical Bayesian*'s further commitment to *Strict* Conditionalization forces them to assign a credence of 1 to any evidence proposition, and hence the Classical Bayesian* is also committed to the invincibility of evidence. Bayesianism is consistent with the rejection of either thesis, and each holds the promise of yielding a version of Bayesianism that's consistent with defeasible updates.¹⁸

Consider, first, what happens if we retain Strict Conditionalization and Probabilism, yet give up on the Ratio Analysis of conditional probability. 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 [Hájek 2003].

Without the Ratio Analysis, it's possible to reduce some maximal credences, although only in a limited set of circumstances. 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 credence 1 to A when I updated upon it. As we've seen, that means that, for any proposition B such that $P(B) \geq 0$, $P(A|B) = 1$, 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 thus 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 by updating accordingly.

This is not an appealing way to accommodate undercutting defeat in a Bayesian framework. 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 to 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 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 $P(BIV)$ to 0, the only way to revise that value endogenously is to update on some other proposition with a credence

¹⁸A thorough discussion of the independent reasons to prefer Jeffrey Conditionalization over Strict Conditionalization or to reject the Ratio Analysis is beyond the scope of this essay. My purpose in this section is instead to identify the version of Bayesianism most amenable to Dogmatism. For broader criticism of the strictness of Strict Conditionalization, see Jeffrey [1983: ch. 11] and Williamson [2000: 203–7]. For a critique of the Ratio Analysis, see Hájek [2003].

of 0.¹⁹ 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.²⁰

Can $P(BIV)$ be revised exogenously? That would be consistent with the formalism, although I won't comment on its plausibility. The problem 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.

Hence, the combination of Dogmatism's commitment to underminable perceptual justification, and Strict Conditionalization's requirement that all propositions being updated upon be assigned a credence of 1, is a poor match for my proposal that the Bayesian Dogmatists should update on the contents of experience.

Classical Bayesianism and Jeffrey Bayesianism treat boundary credences in the same way, and so Jeffrey Bayesians have no great advantage when it comes to the defeasibility of credence 1 propositions. Nonetheless, Jeffrey Bayesianism proves far more amenable to my proposal. For either type of Bayesian, evidence is invincible only 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 propositions that aren't certain, and so evidence needn't be invincible. Hence, the (Jeffrey) Bayesian Dogmatist is free to respond to hand-like experience by exogenously revising their credence in *h* to a value just below 1, thereby preserving its defeasibility.

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 jointly 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.

¹⁹This is because we've rejected the Ratio Analysis only in cases in which the proposition being updated upon is assigned a credence of 0. If $P(BIV) = 0$ and $P(A) > 0$, then we're still committed to saying that $P(BIV|A) = \frac{P(A \& B)}{P(A)} = 0$.

²⁰Although 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. His 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: 286]—and so the criticisms in this section should not be interpreted as criticisms of Hájek's proposal.

The success of the Bayesian argument requires a very specific thesis about this point of contact between Dogmatism and Bayesianism—which is that, upon having an experience as of *A*, the agent should exogenously revise their credence in *I've had an experience as if 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 agents 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 modelling proposal as both inaccurate and prejudicial.

In short, the Bayesian Argument together with the Inferentialist's approach to modelling begs the question against the Dogmatist, and the Bayesian Argument without that approach to modelling 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.²¹

Disclosure statement

No potential conflict of interest was reported by the authors.

References

- Christensen, D. 1992. Confirmational Holism and Bayesian Epistemology, *Philosophy of Science* 59/4: 540–57.
- Hájek, A. 2003. What Conditional Probability Could Not Be, *Synthese* 137/3: 273–323.
- Howson, C. and P. Urbach 1993. *Scientific Reasoning: The Bayesian Approach*, Chicago: Open Court.
- Jeffrey, R. 1983. *The Logic of Decision*, Chicago: The University of Chicago Press.
- Kung, P. 2010. On Having No Reason: Dogmatism and Bayesian Confirmation, *Synthese* 177/1: 1–17.
- Moretti, L. 2015. In Defence of Dogmatism, *Philosophical Studies* 172/1: 261–82.
- Pollock, J.L. and J. Cruz 1999. *Contemporary Theories of Knowledge*, Lanham, MD: Rowman & Littlefield Publishers.
- Pryor, J. 2000. The Skeptic and the Dogmatist, *Noûs* 34/4: 517–49.
- Pryor, J. 2005. There is Immediate Justification, in *Contemporary Debates in Epistemology*, ed. M. Steup and E. Sosa, Oxford: Blackwell: 181–202.
- Pryor, J. 2013. Problems for Credulism, in *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, ed. C. Tucker, New York: Oxford University Press: 89–132.
- Silins, N. 2007. Basic Justification and the Moorean Response to the Skeptic, *Oxford Studies in Epistemology*, Vol. 2, ed. T.S. Gendler and J. Hawthorne, Oxford: Oxford University Press: 108–42.
- Weatherston, B. 2007. The Bayesian and the Dogmatist, *Proceedings of the Aristotelian Society* 107: 169–85.
- Weisberg, J. 2009. Commutativity or Holism? A Dilemma for Conditionalizers, *The British Journal for the Philosophy of Science* 60/4: 793–812.
- White, R. 2006. Problems for Dogmatism, *Philosophical Studies* 131/3: 525–57.
- Williamson, T. 2000. *Knowledge and its Limits*, Oxford: Clarendon Press.

²¹Thanks to David Sosa, Sinan Dogramaci, Josh Dever, Miriam Schoenfeld, audiences at the University of Texas, Austin and the Arché Philosophical Research Institute, and two anonymous referees for comments and objections.

Zardini, E. 2014. Confirming the Less Likely, Discovering the Unknown: Dogmatism—Surd and Doubly Surd, Natural, Flat and Doubly Flat, in *Scepticism and Perceptual Justification*, ed. D. Dodd and E. Zardini, Oxford: Oxford University Press: 33–70.