

Dissertation Summary

Brian T. Miller

When does a perceptual experience as of my hands provide justification for me to believe that I have hands? One initially plausible *positive* requirement is that I must have reasons to believe something else: that my experience is veridical. Also plausible is the *negative* requirement I must not have reasons to believe that my experience is in this case non-veridical. Dogmatists about perception reject the positive requirement and embrace the negative one, holding that perceptual justification is *immediate* — it doesn't rest upon any other justification that I possess — but it is also defeasible, and in particular *underminable*.

Whereas Dogmatism is a theory in the epistemology of perception, Bayesianism is a theory of coherence for partial beliefs and of how coherence it to be maintained in the face of new evidence. Just as it is incoherent to believe both *I have hands* and $\neg(I have hands)$, it is incoherent to be highly confident in both of those propositions. Bayesianism offers a formal account of this latter type of coherence. Importantly, though Bayesianism is a *theory of coherence*, it is not a *coherence theory* as defended by Davidson and Bonjour: it is not an attempt to explain all facts about justification in terms of facts about coherence. Hence the Bayesian's claim that partial beliefs are subject to norms of coherence is at least *prima facie* consistent with the Dogmatist's claim that some beliefs are immediately justified by experience.

I develop and defend Bayesian Dogmatism.

Due to an influential argument that the view is incoherent (by Cohen, Hawthorne, Schiffer, and White), Bayesian Dogmatism has remained mostly undeveloped in the literature. According to that argument, Bayesians are committed to the plausible positive requirement on the provision of justification by perceptual experience that I mentioned above: it requires antecedent justification to believe that the experience is likely to be veridical. That's because given the Bayesian formalism, it is claimed, in order for an experience to rationally license high levels of confidence that I have hands, I must have started out confident that I'm not a handless brain in a vat having hand-like experiences.

This line of argument is mistaken. Bayesian norms govern relations between credences, and I violate those norms when my credences are probabilistically incoherent or when I revise my credence in some proposition without conditionalizing appropriately. Experiences are not credences, and so they cannot enter into the type of relationship governed by Bayesian norms. It's not that Bayesians must deny that experiences affect credences, or deny that these effects are rationally evaluable. Nor is it to deny that Bayesians have a story about how I should react once an experience has affected one of my credences: I should conditionalize. It's just that the Bayesian formalism is silent on the questions of when an experience rationalizes a credence revision, or of the magnitude of that revision, or of precisely which credences ought to be revised in light of a particular experience.

My opponent's argument against Dogmatism requires a specific account of the immediate effect of experience: upon having an experience as of my hands, I should conditionalize only upon the proposition *I've had an experience as of my hands* rather than upon *I have hands* itself. That account isn't obviously false, but it is optional. I argue that by updating instead (or also) on the latter — the content of the experience rather than a fact about the experience — Dogmatists are able to more faithfully represent their account of perceptual justification while avoiding their opponents' objections.

My response to the previous objection requires conditionalization on the contents of experience rather than on propositions about experience. This is a very natural move for Dogmatists to make, as my belief in the content of my experience is precisely what they take to be immediately justified by that experience. Dogmatists also hold that this justification is underminable, and so my proposal requires that my credence in the proposition conditionalized upon must be underminable. However, recent work has shown that underminable inputs to the Bayesian framework are problematic (Weisberg). That's due to a property of Bayesian conditionalization (including of Jeffrey's version) called 'rigidity', which ensures that updating on P cannot change my credence in: Q conditional on P . This is inconsistent with common intuitions about how undermining defeaters work. For example, before I have a visual experience of the new car, *the car is green* and *my color perception is reliable* might be independent of one another, in the sense that changing my credence in either shouldn't affect my credence in the other. But if I experience the apparent greenness of the car and on that basis come to believe that it's green then those propositions should no longer be independent: if I were to learn that my color perception is unreliable then I should decrease my confidence that the car is green. This is what the rigidity of conditionalization appears to rule out: *the car is green* can't lose its independence of *my color perception is reliable* through updating on the former (or the latter).

The rigidity puzzle is a problem not just for Bayesian Dogmatists, but for any Bayesian who thinks that the inputs to the model are underminable. In 'Updating, Undermining and Perceptual Learning' I argue that since the rigidity of conditionalization prevents the introduction of probabilistic dependence within the model, the best response is to introduce it outside of the Bayesian model, i.e. before conditionalization. This is easily accomplished using formal tools readily available to the Bayesian at the point of selecting what to update upon (i.e. selecting the input partition), as becomes apparent with a close look at the ways in which credence revision is and is not constrained by the rigidity of conditionalization.

In the final chapter of my dissertation I consider an alternate response to the rigidity paradox. Jeffrey conditionalization is sometimes criticized for being insufficiently sensitive to the ways that background beliefs affect the rational significance of an experience (Weisberg, Christensen). An experience as of a rhinoceros has one effect on a person who believes that they're typing at their desk — it justifies *I must be hallucinating* — and quite another for someone who believes that they're on safari. The idea is that our update rule ought to model these effects, and Bayesian conditionalization doesn't do that.

One possible response is to scrap conditionalization — and hence Bayesianism — in favor of a more holistic updating rule. I consider and reject one such rule on which the epistemic significance of an experience is calculated using a three-step process: (i) the agent partitions their probability space by background theories; (ii) relative to each background theory the agent determines what credences are immediately affected by the experience and conditionalizes appropriately; and (iii) the new credence in P is the sum of the new values for: P conditional on each background theory, weighted according to how probable the agent finds each theory (Gallow). I argue that, far from being an alternative to Bayesian conditionalization, this updating rule *just is* Bayesian conditionalization together with some informal advice on choosing what to conditionalize upon. For that reason, I argue, 'holistic conditionalization' is able to resolve the rigidity puzzle only because — as I show in 'Updating, Undermining and Perceptual Learning' — Bayesian Conditionalization is able to resolve the rigidity puzzle as well. Hence the drastic response of rejecting Bayesianism is unmotivated.