Published Online First: October 23, 2017

Troubles for Bayesian Formal Epistemology? A Response to Horgan

Jonah N. Schupbach

Abstract: This paper responds to Terry Horgan's recent critique of Bayesian formal epistemology. I argue that each of Horgan's criticisms misses its mark when Bayesianism is viewed as putting forward an inductive logic of confidences. Along the way, I explore the nature, scope, and limits of a defensible brand of Bayesianism.

Terry Horgan's "Troubles for Bayesian Formal Epistemology" (2017) targets a particular brand of Bayesianism, committed to the following five tenets:

(1) A central and fundamental concept for theorizing about epistemological issues is the notion of credence—quantitative, 0-to-1, ratio-scale, (perhaps *rational*) *degree of partial belief*.
(2) The state of mind ordinarily called "belief" is identical to degree-1 "partial belief."
(3) Credence is a type of probability that figures most prominently in matters of rationality.
(4) Normative principles governing epistemic rationality should be, first and foremost, principles governing the generation and updating of a cognitive agent's credences.
(5) Credences should always be *coherent*; they should always conform to the probability axioms.

Against this five-point Bayesianism, Horgan argues (1) that there are no such psychological states as credences, (2) that epistemic probabilities are not rightly understood as credences, and (3) that epistemic probabilities, rightly understood, are scarce. The upshot is that Bayesian formal epistemology as characterized above is a subject without an object, that any credence-based formal epistemology will be posing norms for the regulation of mythological non-existents, and that *any* probabilistic approach to epistemology is going to be exceedingly limited in scope—and so epistemologically anemic.

Horgan's five-point Bayesianism stands in stark contrast to a more orthodox characterization of Bayesianism, involving commitment only to the synchronic principle of Probabilism (roughly Horgan's point 5) and to the diachronic rule for updating that is Strict Conditionalization (one of

Res Philosophica, *Vol. 95, No. 1, January 2018*
https://doi.org/10.11612/resphil.1652

the principles Horgan no doubt has in mind in his point 4).[1] Yet, while he duly acknowledges that there are other senses in which one can be a Bayesian, Horgan suggests that the above tenets describe something like what most contemporary Bayesians are committed to if they are consistent and honest with themselves. For example, while many Bayesians these days might claim to reject point 2, their talk of "degrees of belief" and "partial belief" along with their general enthusiasm for Strict Conditionalization as an even remotely useful model of human learning/inference suggests otherwise. Even so, Horgan claims that his most daunting criticisms hold against less committed versions of Bayesianism, insofar as these essentially involve a quantitative notion of credence.

I've been called a Bayesian, but I don't accept the above five tenets. Instead of playing devil's advocate on behalf of Horgan's five-point Bayesian,[2] I will focus here on another probabilistic approach one might take to formal epistemology—one that I happen to prefer—and argue that it is resilient to Horgan's attack. Many refer to the position I have in mind as "Bayesian"; indeed, I would guess that it is much more commonly endorsed by Bayesian epistemologists than Horgan's five-point version. The availability of such a view reveals that Bayesian formal epistemology is not a lost cause as Horgan suggests, at least not for the reasons he puts forward.

1 Bayesianism as Inductive Logic

When viewed as an approach to inductive logic, Bayesianism is supposed to give us a formal means of clarifying norms of epistemic rationality under conditions of uncertainty. But what features of our epistemic economy should we think of Bayesianism as directly explicating and regulating in the form of probabilistically coherent "credences"? I share Horgan's resistance to an interpretation of credences as "partial beliefs" according to which

[1] In fact, philosophers have grown increasingly liberal as to whether Bayesianism even requires this much. Specifically, many self-acclaimed Bayesians today (e.g., Titelbaum 2013) drop Conditionalization, suggesting that Probabilism alone (or perhaps Probabilism plus a suitable replacement for Bayes's Rule) suffices as a defining tenet of contemporary Bayesianism.

[2] The devil's advocate might, for example, argue against the conceptual confusion objection that Belief, not Credence, is the source of confusion. The ordinary usage of "belief" is ambiguous, having all-or-nothing and gradational meanings attached to it (and that's true even if, by "belief", we follow Horgan in intending a "certain kind of psychological commitment"). Formal precisifications of "belief" that disentangle these senses have the potential to do us a great service. Of course, Credence would be a great distortion of all-or-nothing Belief, but this is not a problem if it's only meant to explicate a gradational sense of Belief. The devil's advocate might furthermore point out that, even if the above were not correct—even if no ordinary concept of Belief was gradational—it still would be too quick to conclude that the idea of partial belief is inherently confusing matters. It could, after all, be a newly engineered concept that allows us to think more carefully about (philosophical and other) matters relating to, but admittedly distinct from, Belief. To borrow an example from Carnap (1950, 12-15), the fact that Warm is inherently all-or-nothing does not mean that a scientifically-engineered concept of Temperature distorts the ordinary concept!

maximal probability is identified with belief *simpliciter*. If instead, as is often expressed, maximal probability explicates unbending certainty, then credences are more naturally thought of as degrees of confidence.[3]

So construed, Bayesianism most straightforwardly provides a logic of confidences. As a normative theory, it examines how an agent's confidences *ought* to look. If theorized credences, or degrees of confidence, are properly thought of as being located in the psyche of any epistemic agent, it is only in regard to a fictional, ideal Bayesian agent.

Horgan briefly discusses this sort of interpretation. He also considers and defends a Carnap-inspired interpretation of probability as degree of logical-evidential support, but he finds no place for credences in this idea. The view that I have in mind combines these alternatives, using probabilities to articulate normative, logical-evidential constraints *on our credences (confidences)*. Credence retains a place in this logical-evidential picture, because our actual credence judgments (be they qualitative or quantitative, relational or absolute) are regulated by these logical-evidential constraints—or put another way, because any particular probability function that satisfies the pertinent logical-evidential constraints can be thought of as the theoretical credences of one idealized, normative Bayesian agent (out of potentially many). Note that one need not advocate a full-on Carnapian interpretation here. The point is simply to consider any view that would take Bayesianism to be giving us an inductive logic, which regulates actual confidences by comparing them to probabilistically articulated logical-evidential standards.

It seems to me that, though still involving a quantitative notion of credence, any such view is less vulnerable than Horgan realizes to his attack. Let me take each of his main criticisms in turn and argue that they lose their edge when Bayesianism is thought of as offering an inductive logic of degrees of confidence.

2 A Subject with No Object?

Horgan (2017) puts forward a number of skeptical doubts about the very concept of credence. This concept "embodies (i) a conceptual confusion that (ii) is superimposed upon a psychological myth, all of which (iii) is exacerbated by an equivocating conflation" (235). He acknowledges, however, that the move to degrees of confidence (rather than degrees of partial belief) naturally allows one to sidestep complaints (i) and (iii).

[3] I'm not sure how well my talk of "degrees of confidence" aligns with Horgan's, however. For one thing, Horgan takes it for granted that maximal or "complete confidence in p" is identical to "outright belief in p." But it seems clear to me that one could fully believe a proposition p with less than full confidence (where "less than full confidence" need not amount, as Horgan suggests, to just saying "less than decisive evidence"). Also, the concept that aptly corresponds to maximal degree of confidence, according to the idea of confidence I'm working with, is unbending certainty—as opposed to any variety of belief.

In arguing for (ii), Horgan maintains that the idea that there are psychologically real, 0-to-1, ratio-scale, degrees of confidence is a myth. The theoretical entities constituting the credence structure of an idealized epistemic agent have no correspondence to reality. So-called norms for degrees of confidence are thus purely theoretical, apt only for hypothetical super-agents with epistemic states more precise and structured than those of humans. But if these norms are worthless for actual agents, then "the discipline that nowadays calls itself 'Bayesian formal epistemology' is relevantly similar to past disciplines like alchemy and phlogiston theory; it is not about any real phenomena, and thus it also is not about any genuine norms that govern real phenomena" (241).

Notably, Horgan argues for this with non-logical-evidential interpretations of probability in mind. Because he neglects any positive connection between degrees of logical-evidential support and credences, he has no motivation to consider such interpretations in making his point about credences. However, on the proposed Bayesian view, there is such a positive connection. Credences are the rational degrees of confidence of an ideally Bayesian agent, those which satisfy logical-evidential constraints. Are there epistemic states of actual agents that would be appropriate to compare against these credences?

Horgan actually gives us reason later in his article to think so. He suggests that there are situations in which "an epistemic agent might contemplate and understand . . . a specific, quantitative, degree of evidential support that [evidence] confers upon [a proposition]." And he specifically clarifies that such degrees of evidential support are available to actual agents in cases where either objective chances are available or the agent is in a context of suitable "evidential indifference."

Horgan offers a variation on the Monty Hall game as an example in which an agent is in such a position. Given the evidential indifference that a contestant has between the possibilities that a grand prize is hidden between either one of three doors, Horgan allows that the contestant may appropriately assess each of these possibilities as having equal degree of evidential support. While it is not implied by our Bayesian approach that the contestant is rationally required to have a corresponding confidence judgment, this is exactly the sort of case where it is not at all psychologically unrealistic to think that an agent could—and that agents in such a position often do. If prompted, the contestant would be expected to express equal confidence in the three possibilities. But this real possibility suffices to demonstrate the point that quantitatively precise credences are not the fictional entities that Horgan takes them to be. After all, this confidence judgment (along with the contestant's background knowledge that there is only one such prize hidden behind exactly one such door) can be expressed accurately with quantitative precision. Choose whatever conventions of formal explication you like; on the standard Bayesian, 0-to-1 scale of real values, the exact associated quantitative credences are {1/3, 1/3, 1/3}.

The same point holds for Horgan's route from known objective chances to quantitatively precise credences. Varying his example, Horgan imagines that the contestant is informed "that the location of the prize was based on the following, biased, randomizing procedure: after a fair die was tossed, the prize was placed behind door 1 if the die's up-face was either 1 or 2, behind door 2 if the up-face was 3, and behind door 3 if the up-face was 4, 5, or 6." Again, our Bayesian approach need not be taken to imply that the contestant ought to have corresponding confidence judgments in this case; however, it's worth noting that this too is a case in which it is not unrealistic to think that the agent might actually have them. If prompted, given that the contestant knows these objective chances, he or she may affirm being twice as confident that the prize is behind door 1 than door 2—and three times as confident that it's behind door 3. Again given the contestant's background knowledge, these realistic confidence judgments are quantitatively precise. On the standard Bayesian, 0-to-1 scale of real values, the exact associated quantitative credences are {1/3, 1/6, 1/2}.

The upshot is that the sorts of circumstances which, as Horgan argues, allow for agents to assess quantitative degrees of evidential support also seem to be circumstances in which it is not psychologically far-fetched to think that agents may have correspondingly precise confidences. One important note to avoid confusion here: Saying that an agent has quantitatively precise credences in these cases is crucially different than saying that the agent has quantified his or her own precise credences using the Bayesian conventions, or has done so at all.

3 A Subject with Scant Object

Credences are thus not "superimposed upon a psychological myth," but instead correspond to actually existing quantifiable degrees of confidence. This response quickly gives rise to Horgan's next criticism however. As he argues, quantitatively precise degrees of evidential support are only assessable by actual epistemic agents in very special circumstances. Even if quantitatively precise degrees of confidence may be available to actual agents in such cases then, such credences are arguably scarce. While Bayesianism might not be a subject without an object, it's object is ill-suited to be of central epistemic concern. Credences exist, but the vast majority of epistemic states to be evaluated by epistemic norms do not involve them. And credence-based formal epistemology thus does not have the scope it is often claimed to have, the scope required if it is to give us an interesting, general epistemology.

This criticism is only worrying to the extent that the Bayesian standard only holds precise credences up as norms. However, the idealized structures of rational Bayesian agents contain more than information about such rational quantitative degrees of confidence. Logical-evidential constraints

on *qualitative* assessments of confidence can be read off the highly idealized structure of Bayesian agents.[4] Thus, Bayesian formal epistemology also speaks to cases in which actual agents can make such qualitative assessments. Horgan finds it plausible "that epistemic agents often have certain *qualitative* degrees of confidence in the various propositions that they can entertain and understand" (236). He clarifies that actual agents regularly express such qualitative confidence relations in absolute judgments (e.g., "a given proposition is *highly likely* given the evidence, or *somewhat more likely than not*, or *equally as likely as not*, and so on") and in relative judgments (e.g., "proposition p [is] *much more likely* than proposition q, or *somewhat more likely* than q, or *equally as likely as* proposition q, etc."). These common qualitative confidence judgments have equal standing as epistemic objects regulated by this form of Bayesianism.

What other epistemic judgments or attitudes might a Bayesian epistemology regulate? The obvious question is whether Bayesian epistemology can provide a logic of *beliefs*. The answer depends on whether beliefs can be read off of the credence structure of our ideal Bayesian agents. Horgan is skeptical, referring to "the awkwardness of trying to say something both plausible and normatively general about when, and why, a given less-than-1 credence in p would render it epistemically appropriate to *believe* that p" in light of the lottery paradox. He may be right; this is a matter of current debate. But even if the answer is no, and only judgments of confidence (quantitative and qualitative, absolute and comparative) can be regulated by the Bayesian machinery, this suffices to make formal epistemology very much worth our while. Horgan himself says as much when he asserts that such qualitative judgments alone "surely are both psychologically real and epistemologically important. Likewise, theoretical inquiry about epistemic norms regarding such judgments, and about epistemic norms regarding the interconnections between such judgments and the justification-status of beliefs, is surely within the appropriate purview of epistemology" (244).

4 Credences as Idealizations?

When Bayesianism is viewed as an approach to inductive logic, probabilities are interpreted as the rational credences of an ideal Bayesian agent. These credences themselves are idealizations contained in a theoretical, normative model of reasoning under uncertainty. Against the view that credences should be thought of as useful idealizations, Horgan (244) argues (1) that

[4] If one allows there to be multiple, rational, idealized Bayesian agents, then these would plausibly be the qualitative relations that hold robustly across all such agents. For example, such theoretical agents might diverge on the exact theoretical credence they assign to Heads versus Tails, but agree robustly in having a greater theoretical credence of Heads. In that case, the logical-evidential constraints encoded in our Bayesian inductive logic would not be so demanding as to lock agents who have quantitatively precise degrees of confidence down to one particular value, although they would still legislate actual, *qualitative, comparative* judgments of confidence.

useful idealizations typically "introduce only a relatively minor discrepancy between the actual workings of the explanatorily pertinent factors and the model's characterization of those workings" and (2) that theoretical credences are "no mere minor approximation to anything that is in fact psychologically real," but substantially distort the reality. In light of these considerations, Horgan concludes that the onus is on the Bayesian "to explain to us skeptics why your framework *does not* badly distort the very psychological and epistemic-normative phenomena you seek to illuminate" (247).

While I agree that Bayesians bear a burden of showing that their idealizations are not pernicious to the aims of their theory, I do not agree with Horgan on the nature of that burden. Theoretical models (in science, philosophy, and elsewhere) necessarily distort their target phenomena to some extent, and good models often distort the reality very "badly" in some respects. It seems to me that the formal epistemologist must convince us, as any modeler must, that the ways in which they have distorted the target are not confounding the desired results—as opposed to showing us that they have not much distorted the target. Since (as it seems to me) Bayesian models manifestly do substantially distort certain features of actual reasoners, this is a potentially easier burden to bear.

The start of such a response should clarify the nature of such distortions. Most obviously and at the very least, the ideal Bayesian agent is unlike actual agents on account of having probabilistically coherent credences over an infinitely vast set of propositions. While the first assumption—that these credences be probabilistically coherent—may indeed be a substantial distortion, it is in no way pernicious. By virtue of the fact that this feature of the model is meant to be prescriptive, not descriptive, the model is allowed to part from the psychological reality in this regard. If the Bayesian is right that probabilistic coherence is necessary for rational degrees of confidence, then the question of just how substantial this distortion is (relative to any particular agent's confidences) is a key question to ask when logically evaluating that agent.

It is more likely that the assumption that the ideal Bayesian agent has credences *over an infinitely vast set of propositions* confounds the results of our model. This would mean that our logical conclusions and evaluations grounded in this model are off-base because of this assumption, that this assumption leads us to call rational things irrational. But it seems to me that if we remain modest enough in our ambitions for just how much logical input Bayesianism is supposed to give us, then there is reason to think that this assumption does no such thing.

A normative standard need not be (and I would surmise that no normative standard should be) taken as prescriptive with respect to all of its features. Just because a normative theory portrays the ideal agent as having certain features, these features *per se* need not be prescribed to actual agents. Take a (perhaps) less controversial example. Standard deductive logic may

be used as an eminently useful model of good monotonic reasoning under conditions of certainty. The ideal deductive agent can usefully be thought of as one who's beliefs map onto a maximal, consistent (deductively closed) set of propositions.[5] This is not to say that, when we evaluate actual epistemic agents by the lights of deductive logic, we condemn them for not having deductively closed, infinite sets of beliefs, or for not considering all the infinite number of atomic propositions corresponding to the atoms of our formal system. These features of the ideal deductive agent do not serve as suitable standards for actual epistemic agents. Instead, we evaluate those beliefs had by any actual agent by exploring the extent to which they might fit into an extended, idealized deductive system. The salient evaluative question is, "Could these beliefs all be believed by an ideal deductive agent?" and not "Does this actual agent have all the same beliefs as an ideal deductive agent?"

Back to Bayesianism; as with deductive standards of rationality, so with inductive. The ideal Bayesian reasoner has coherent (probabilistic), 0-to-1, ratio-scale, degrees of confidence (i.e., credences) over an infinitely large algebra of propositions. But this is not to say that, when we evaluate actual epistemic agents by the lights of Bayesian logic, we should condemn them for not having 0-to-1, ratio-scale, degrees of confidence over some (or even the vast majority of) considered propositions, or for not considering in the first place all the infinite number of propositions represented in our algebra. These features of the ideal Bayesian agent do not serve as suitable standards for actual epistemic agents. Instead, we evaluate those credences (if any) that are had by any actual agent by exploring the extent to which they might be fit into an extended, idealized Bayesian system. The salient evaluative question is, "Could these credences all be had by an ideal Bayesian agent?" and not "Does this actual agent have all of the same credences as an ideal Bayesian agent?"

So modesty in our logical ambitions alleviates the worry that the ideal Bayesian agent sets an impossible standard for actual reasoners. That the ideal agent has 0-to-1, ratio-scale, degrees of confidence over a potentially infinite number of propositions need not be seen as a prescriptive (or descriptive) feature of the Bayesian model. But is there positive reason to think that the Bayesian standard is not confounded by this assumption and hence that it properly regulates the confidences (quantitative, qualitative, absolute, comparative) of non-ideal agents? To evaluate an actual agent's quantitative degrees of confidence, for example, why should it be pertinent to look to an agent that has the full gamut of credences? The start of an answer to such questions can defer to various standard defenses of Bayesianism. If an actual agent's quantitative credences, few as they may

[5] The ideal deductive agent could also be characterized using available external, material considerations. For example, knowing that certain propositions are true, the ideally deductive agent could be thought of as one who's beliefs map onto a maximal, consistent (deductively closed) set of propositions containing *inter alia* all these true propositions.

be, cannot fit into the epistemic state of any ideal Bayesian agent, then that is because on their own (without having to concoct non-actual credences for the agent or come up with a story about how the full gamut of credences is actually there if hidden) these credences already run the agent into trouble. The nature of this trouble will depend on what requirements must be satisfied by the ideal Bayesian agent. One can at least say that credences that cannot be held by any ideal agent *on pain of incoherence* will on their own already lead us to deem sure-loss betting scenarios as fair, be dominated in terms of expected inaccuracy by alternative sets of credences, etc. Of course, one might at this point join a longstanding debate by challenging the standard defenses of Bayesianism, but this would be to argue against Bayesianism for different reasons than those conveyed by Horgan.

Jonah N. Schupbach
E-mail : jonah.n.schupbach@utah.edu

References:

Carnap, Rudolf. 1950. *Logical Foundations of Probability*. Chicago: University of Chicago Press.

Horgan, Terry. 2017. "Troubles for Bayesian Formal Epistemology." *Res Philosophica* 94 (2): 233–255.

Titelbaum, Michael G. 2013. *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford: Oxford University Press.