

Penultimate draft. In *Res Philosophica* 94, 2 (2017): 233-255.

## Troubles for Bayesian Formal Epistemology

Terry Horgan, University of Arizona

Bayesian formal epistemology (henceforth, BFE) is nowadays very trendy. But is it a sound approach to epistemological issues, with serious prospects for yielding genuine illumination of those issues? I strongly doubt it. My principal goal in this paper is to articulate and motivate my doubts. My realistic-seeming hope is that fans of BFE will recognize that these doubts pose a serious challenge, and will feel significant pressure to address this challenge. My unrealistic-seeming hope is that I will persuade some fans of BFE that this research program is flawed and unpromising, and that therefore it probably is not worth pursuing. (How often does *that* form of persuasion occur in philosophy?) My cautiously optimistic hope is that I will instill in the minds of younger readers a healthy skepticism about BFE that will dispose them to resist its trendiness rather than being swept up by it. (Philosophers in general, and epistemologists in particular, like to think that their views on philosophical issues arise via pure reason, with sociological factors like trendiness playing no significant role. Baloney!)<sup>1</sup>

BFE, as I understand it, appears to rest (at least implicitly, and despite what advocates of BFE sometimes explicitly avow) upon the following foundational assumptions. First, a central and fundamental concept that should be employed in theorizing about epistemological issues is the notion of *credence*—with this notion being understood as quantitative, 0-to-1, ratio-scale, *degree of partial belief* (or, in some contexts, *rational degree of partial belief*). (The expression ‘partial belief’ often is used in full generality, with degree 0 and degree 1 counting as limit cases of partial belief.) Second, the state of mind ordinarily just called “belief” is identical to the top limit case: degree-1 “partial belief.” Third, rational credence—i.e., rational partial degree of belief—is one type of probability; indeed, it is the kind of probability that figures most directly in matters of epistemic rationality. (Another type of probability, nowadays commonly called *chance*, instead is supposed to be something like in-the-world objective propensity—e.g., the propensity of a fair coin to come up heads when flipped.) Fourth, normative principles governing epistemic rationality should be, first and foremost, principles governing the generation and updating of a cognitive agent’s credences. Fifth, one such normative principle is that

credences should always be *coherent* in the probabilistic sense—i.e., they should always conform to the axioms of probability theory.

There are three principal reasons why I contend that advocates of BFE, despite what they might say explicitly about credence, belief, and their inter-relations, are implicitly committed in their theoretical practice to all five of these foundational assumptions. First, they typically embrace the familiar claim that that the rationally appropriate way to update one's credences is by conditionalization. Conditionalization works this way: if one's prior probability for  $A$  given  $B$  is  $n$  and one subsequently acquires the evidence  $B$ , then one sets one's updated absolute probability for  $A$  equal to  $n$ ; i.e.,  $P(A) = P(A|B) = n$ . This effectively presupposes that “acquiring evidence  $B$ ” amounts to proposition  $B$  receiving credence 1 (and proposition  $\sim B$  thereby receiving credence 0)—which, in turn, effectively amounts to treating credence 1 for  $B$  as coinciding with full-fledged belief that  $B$  rather than coinciding with some stronger psychological state (e.g., absolute certainty that  $B$ ). Second, in practice the BFE advocates typically write as though the normative principles of epistemic rationality they seek to articulate are directly applicable to actual epistemic agents, insofar as those agents are rational. This presupposes that credences are psychologically real mental states of actual human beings. Third, in practice the BFE advocates typically just eschew the term ‘belief’ in favor of the term ‘credence’. This effectively amounts to replacing the traditional notion of belief with the graded notion of 0-to1, ratio-scale, partial belief (with credence 1 being treated as constituting full belief rather than something like absolute certainty, because of BFE's embrace of conditionalization).

I have extensive doubts about this foundational framework. I maintain that the concept of credence is deeply problematic, in multiple respects. I therefore believe (i) that there are no such psychological states as credences, (ii) that the probabilities that figure most directly in matters of epistemic rationality—I call them *epistemic* probabilities—are not rightly understood as credences, and (iii) that epistemic probabilities are not well suited to play a central role in epistemological theorizing.<sup>2</sup>

I do realize that some fans of BFE will explicitly repudiate some of the lately mentioned foundational assumptions—regardless whether or not this comports with their actual theoretical practice. I will consider several alternative potential versions of BFE as the paper proceeds. All the versions face the same—most serious—trouble; also, they each face significant additional troubles, although not exactly the same ones.

## 1. Skeptical Doubts about the Concept of Credence<sup>3</sup>

In my view, the concept of credence—i.e., the concept of 0-to-1 ratio-scale, “partial belief”—is triply problematic: it embodies (i) a conceptual confusion that (ii) is superimposed upon a psychological myth, all of which (iii) is exacerbated by an equivocating conflation. Let me explain.

First, the conceptual confusion. The ordinary concept of belief, I contend, is an all-or-nothing notion—which means that the very idea of “partial” belief is inherently confused, a distortion of the ordinary concept. Belief that  $p$  is a certain kind of psychological *commitment* state: viz., a commitment to it’s being the case that  $p$ . One either believes  $p$ , or one disbelieves  $p$  (i.e., one believes *not-p*), or one neither believes  $p$  nor disbelieves  $p$ . I.e., one either is is-wise committed with respect to  $p$ , or one is is-wise committed with respect to *not-p*, or one is not is-wise committed either way. Although it can be *vague* whether or not one believes  $p$ , the ordinary notion of belief simply does not allow for so-called “partial” beliefs.

Second, the psychological myth. Those who embrace the degree-of-belief conception of epistemic probability typically assume that for virtually any proposition  $p$  that an epistemic agent can entertain and understand, the agent has some specific, 0-to-1 ratio-scale, degree of *confidence* in  $p$ . Now, quite apart from the question whether such a degree of confidence should be considered a “partial belief” in  $p$  (and I claim it should *not* be), in any case it is extremely implausible that epistemic agents normally have such quantitative degrees of confidence for all—or for most—of the propositions that they can entertain and understand. The usual reason for supposing that they do appeals to their potential betting behavior with respect to a potential infinity of successively more fine-grained betting decisions—as systematized, for example, by Savage’s axioms that supposedly govern a (rational) agent’s subjective “utilities” and subjective “partial beliefs.” But there are at least two serious problems with this. First, if one were actually faced with *forced choices* regarding some such betting decisions, then sooner or later one would find oneself making choices that one experiences as *arbitrary* rather than as reflecting one’s actual bet-preferences; and the strong defeasible presumption is that fine-grained forced choices that one experiences as arbitrary are indeed psychologically arbitrary. (Prediction: If one experimentally requires human subjects to make such progressively more fine-grained betting decisions, and one does so more than once in separate sessions in such a way that in later sessions the subjects no longer remember their

earlier forced-choice decisions, one will find no discernible cross-time consistency in the forced-choice decisions that the subjects themselves describe as being experientially arbitrary.) A further problem with appealing to hypothetical betting behavior, in order to try to underwrite the putative psychological reality of 0-to-1 ratio-scale degrees of confidence, is that this is guilty of the equivocating conflation to which I will shortly turn.

Much more plausible than this psychological myth about putative quantitative degrees of confidence is the contention that epistemic agents often have certain *qualitative* degrees of confidence in various of the propositions that they can entertain and understand but which they neither believe nor disbelieve. These involve qualitative non-comparative categories, expressible in judgments of the form “*p* is highly likely,” “*p* is moderately likely,” “*p* is highly unlikely,” and such; they also involve qualitative comparative categories, expressible in judgments of the form “*p* is much more likely than *q*,” “*p* is somewhat more likely than *q*,” and “*p* and *q* are equi-likely.”<sup>4</sup> These qualitative categories are not cashable in terms of putative *quantitative* degrees of confidence on a 0-to-1 ratio scale, because often enough an epistemic agent will have certain qualitative degrees of confidence concerning certain propositions without having any corresponding quantitative degrees of confidence concerning them.<sup>5</sup> For instance, often enough an epistemic agent will regard *p* and *q* as equally likely without having any specific quantitative degree of confidence in *p* and *q* at all. Limit cases on this qualitative-confidence scale are outright belief in *p*, which constitutes complete confidence in *p*, and outright disbelief in *p* (i.e., outright belief in *not-p*), which constitutes complete lack of confidence in *p*.<sup>6</sup>

Third, the equivocating conflation. Among one’s beliefs, all of which fall on the complete-confidence end of the qualitative-confidence scale just mentioned, often some are held *more strongly* than others—typically because some of them are regarded as being more strongly supported by one’s available evidence than others. This is so even though one also regards them all as being sufficiently well evidentially supported to be belief-worthy, which is why one actually does believe them all; and one’s believing them all is a psychological fact that locates them all at the full-confidence end of the lately described qualitative-confidence scale. I myself believe, for example, both that Washington D. C. is the national capitol of the United States, and that Aristotle was a teacher of Alexander the Great. But I believe the former much more strongly than I believe the latter, because I regard my evidence for the former as much stronger than my evidence for the latter. (I don’t now even remember having read or having been

taught that Aristotle was a teacher of Alexander the Great, but it now seems to me that I must have learned it somehow.) Since I believe them both, I count as being psychologically *fully confident* in them both.

The conflation comes when one does two things together: (1) construing full-fledged belief as degree-1 confidence on a 0-to-1 ratio scale, and (2) construing different strengths of belief as different positions on that very same 0-to-1 ratio scale. (If there really were such a thing as 0-to-1 ratio-scale partial belief, then all full-fledged beliefs would have degree 1 on *that* scale, regardless of their strength.)<sup>7</sup> This mistake is deeply ingrained in much contemporary “received wisdom” about epistemic probability, and of course it is only abetted by the idea that a person’s actual and potential betting behavior underwrites the putative psychological reality of 0-to-1 ratio-scale degrees of belief for each proposition that the person can entertain and understand. Often enough, one’s betting behavior would be better construed as reflecting the comparative strengths of certain *full-fledged* beliefs one has, rather than reflecting one’s putative, quantitative, degrees of *partial* belief. (I myself would bet a lot more money on the proposition that Washington D.C. is the national capitol of the United States than I would on the proposition that Aristotle was a teacher of Alexander the Great—even though I believe both propositions.)

Let me make two further points. First, I have been using the terms ‘confidence’ and ‘strength’ somewhat stipulatively, in order to mark an important conceptual distinction. In ordinary parlance, one could certainly express a comparative-strength judgment by saying “I am *more confident* that Washington D.C. is the national capitol of the United States than I am that Aristotle was a teacher of Alexander the Great, even though I believe both propositions.” Second, strength of belief is usually a qualitative matter, not a quantitative one. (The rationale for this claim is the same as for the corresponding claim about confidence.) For example, I myself believe *very strongly* that Washington D.C. is the national capitol of the United States, whereas I believe only *modestly strongly* that Aristotle was a teacher of Alexander the Great; and I believe the former proposition (*much*) *more strongly* than I do the latter. Thus, each of the two key degree-notions I have distinguished here is qualitative, rather than quantitative: on one hand, degrees of *confidence* (both comparative and non-comparative), which accrue in non-limit cases to propositions that one neither believes nor disbelieves, and on the other hand, degrees of *strength* (both comparative and non-comparative), which accrue to propositions that one believes.

## 2. Skeptical Doubts about Differently-Construed Credence

There is a somewhat different way of construing the concept of credence, often officially endorsed in the philosophical literature, which evades at least some of my critique in the preceding section. On this alternative picture, degree-1 credence does not constitute the psychological state of full belief, but instead coincides with the psychological state of *absolute certainty*.<sup>8</sup> Belief, ordinarily so-called, is a psychological commitment-state vis-à-vis a proposition  $p$ ; it is distinct from any credence-state; and normally the belief that  $p$  will coincide with some credence for  $p$  that is high while yet less than 1. (Perhaps the level of credence that is required for belief varies with conversational context.)

Those who officially advocate this alternative construal are apt to paint themselves into an awkward theoretical corner, if they then go on to embrace the familiar claim—a mainstay of so-called “Bayesianism”—that the rationally appropriate way to update one’s credences is by conditionalization. The problem, as I said above, is that this effectively presupposes that “acquiring evidence  $B$ ” amounts to proposition  $B$ ’s acquiring a credence of 1—which, according to the construal of credence now under consideration, virtually never happens because credence 1 constitutes *absolute certainty*. Thus, what advocates of this construal should say, in order to maintain a theoretically consistent position, is that updating of one’s credences normally should happen some other way than by ordinary conditionalization—say, by “Jeffrey conditionalization.” Yet those who call themselves “Bayesians” typically fail to do this; instead they advocate updating by ordinary conditionalization. In effect, I think, they thereby commit themselves to the partial-belief construal of the notion of credence that I criticized in section 1—even if they officially embrace the alternative construal now being discussed.

Be that as it may, consider the view that (a) does embrace this alternative construal of the notion of credence, (b) asserts that updating of credences normally should happen by Jeffrey conditionalization (and not by ordinary conditionalization), (c) eschews the notion of partial belief altogether, (d) asserts that the belief that  $p$  normally co-exists with some credence for  $p$  that is high but yet less than 1, and (e) asserts that full belief that  $p$  is a psychological state that is distinct from any credence-state vis-à-vis  $p$ .<sup>9</sup>

This view evades my first objection from section 1: it eschews the notion of “partial belief,” and it does not identify beliefs with credences of any kind (either degree-1 credences or sufficiently high quantitative-degree credences). It also evades my third objection: rather than conflating degree of confidence with degree of strength, in effect it posits a single, 0-to-1 ratio-scale, range of credences that

includes both quantitative degrees of confidence for not-believed propositions and quantitative degrees of strength for believed propositions. (As one might put it, the psychological states posited by this approach are precise, ratio-scale, degrees of *credibility* that an epistemic agent supposedly experiences vis-à-vis propositions that the agent can contemplate and understand. Degree-one credibility coincides not with mere belief, but rather with absolute certainty.)

Nonetheless, my second objection is just as applicable to this alternative construal of the notion of credence as it is to the construal discussed in section 1. The idea that there are psychologically real, zero-to-one, ratio-scale, degrees of credibility is just as mythological, and for exactly the same reasons, as is the idea that there are such psychologically real degrees of *confidence*.

### 3. **Skeptical Doubts about BFE**

If my critique of the concept of credence is correct, what is the import for BFE? Prima facie, the import is quite seriously negative, as I will now explain. (In subsequent sections I will consider several ways of trying to evade this pessimistic assessment, and I will argue that none are promising.)

I have argued that typically, for any given epistemic agent and any proposition  $p$  that the agent might contemplate and understand, there is no such psychological state as that agent's credence in  $p$ . This is so regardless whether (1) credence is construed as quantitative degree on a zero-to-one confidence-scale for which degree-1 credence is full belief and lesser degrees are merely-partial belief, or (2) credence instead is construed as quantitative degree of credibility on a zero-to-one credibility-scale for which degree-1 credence is absolute certainty. Either way, the question, "Does the agent's credence in  $p$  conform to epistemic norms governing the rational formation and rational updating of credences?" is a bad question. It suffers presupposition failure, because there is no such state as *the agent's credence in p*. Indeed, because of this the question also thereby suffers an additional presupposition failure: there are no such things as *the norms governing the rational formation and rational updating of credences*. This being so, 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. Not a happy outcome!

Even apart from this principal reason for skeptical doubts—a reason that is equally applicable relative to either conception of credence—the two respective conceptions each suffer additional problems

as well. A version of BFE deploying the first conception encounters yet another kind of presupposition failure: it presupposes, in effect, that there are no such psychological phenomena as varying degrees of strength among full-fledged beliefs. But there *are* such phenomena (although degree of strength is really only a qualitative matter, not something quantitatively precise); and prima facie, they themselves are subject to epistemic norms of rationality. Yet BFE deploying the first conception of credence renders this phenomenon virtually invisible, by virtue of the way it conflates degrees of confidence with degrees of strength.<sup>10</sup> (Even if the putative kinds of credence were psychologically real, epistemically apt accommodation of new evidence would not necessarily involve adjusting one's credence; instead, sometimes it would involve adjusting the strength of one's full-fledged belief (one's degree-one credence), while leaving the credence itself at degree one.)

A version of BFE employing the second conception of credence would encounter additional problems too, over and above the principal problem of psychological mythology. One difficulty is the awkwardness of trying to say something both plausible and normatively general about when, and why, a given less-than-1 degree of credence in  $p$  would render it epistemically appropriate to *believe* that  $p$ .<sup>11</sup> The lottery paradox looms large here, for example: outright belief that one's lottery ticket will lose does not seem epistemically justified, no matter how high are the odds against winning.

In addition, the appeal to Jeffrey conditionalization is problematic, relative to the second conception of credence. Christopher Meacham (2016) explains Jeffrey conditionalization as follows:

One common complaint about Conditionalization is that it requires us to adopt a credence of 1 in our evidence.... Like Conditionalization, Jeffrey Conditionalization is a constraint on how a subject's credences should change over time in light of evidence. But on this picture, one's evidence isn't a proposition  $E$ ; instead it's a weighted partition of propositions  $S = \{(E_1, x_1), (E_2, x_2), \dots\}$  (where  $x_1-x_n$  are real numbers in the  $[0,1]$  interval that sum to 1). Given such evidence, Jeffrey Conditionalization tells us to update as follows:

**Jeffrey Conditionalization:** If a subject with credences  $cr$  gets evidence partition  $S = \{(E_1, x_1), (E_2, x_2), \dots\}$ , she should adopt new credences  $cr_S$  such that:

$$cr_S(\cdot) = \sum_i x_i \cdot cr(\cdot \mid E_i), \text{ if defined.}$$

Intuitively, Jeffrey Conditionalization tells us that upon receiving evidence partition  $S$ , we should assign each  $E_i$  a credence of  $x_i$ , and then renormalize; that is, move our credence in the  $E_i$ s to the

indicated amount, and then distribute the credence assigned to each  $E_i$  among the propositions entailing  $E_i$  in a way that keeps the ratios between them the same. (Meacham 2016, p. 769)

I myself would offer the following alternative gloss, worded similarly to Meacham's. Intuitively, Jeffrey conditionalization tells us that upon receiving evidence partition  $S$ , one should first update one's credence for each  $E_i$  this way:

$$(1) \quad cr_S(E_i) = x_i$$

One should then update one's credence for each proposition ( $H \& E_i$ ) this way:

$$(2) \quad cr_S(H \& E_i) = cr_S(E_i) \cdot cr(H | E_i), \text{ if defined}$$

(This is the natural generalization of ordinary conditionalization.) Since the respective  $E_i$ s constitute a partition, one's updated credence for proposition  $H$  therefore is this:

$$(3) \quad cr_S(H) = \sum_i cr_S(H \& E_i) = \sum_i x_i \cdot cr(H | E_i), \text{ if defined.}$$

In the present context, Meacham's intuitive gloss on Jeffrey conditionalization—and my own gloss, which cleaves closely to Meacham's terminology—raise some troubling questions. What are those  $x_i$ s, considered as constituents of an evidence partition  $S$  and before they are adopted as new credences for the respective  $E_i$ s in  $S$ ? Does Jeffrey conditionalization effectively smuggle in some notion of evidence-involving probability, distinct from the notion of credence itself? If one adopts the new credence  $x_i$  for the proposition  $E_i$ , is this itself a form of credence-updating that is distinct from either ordinary conditionalization or Jeffrey conditionalization? Finally, isn't the claim that evidence normally has the form required for applicability of Jeffrey conditionalization just wildly mythological anyway?

Perhaps the best response to these troubling questions, on behalf of a fan of BFE who seeks to construe credences in the manner described in Section 2 above, is to say that the reception of new evidence really amounts to an experiential "jolt" whose essence is to *directly re-set* one's credences for the respective elements of a partition of propositions  $S^* = \{E_1, E_2, \dots\}$ , with the respective new credences being those cited in the corresponding weighted partition  $S = \{(E_1, x_1), (E_2, x_2), \dots\}$ . But this response raises troubling questions of its own. Wouldn't some ways of being thus "jolted" by an evidentially pertinent new experience be more epistemically appropriate than others? Wouldn't it be an important task for epistemology to sort out the difference between an epistemically appropriate jolt and epistemically inappropriate ones?

The additional problems I have been raising about the two respective versions of BFE are subsidiary worries—although serious ones in their own right. The really big problem is that BFE, conceived as an enterprise that treats credence as its central notion, is a discipline which—like phlogiston theory—evidently has no genuine subject matter. There are no such psychological states as credences, and

hence there are no such epistemic norms of rationality as norms governing the rational generation and rational updating of credences.

#### 4. Credences as Idealizations?

One response I often hear, when voicing to fans of BFE my doubts about the notion of credence, goes as follows. Even if it is true that talk of credences does not track any real psychological states of human epistemic agents, nevertheless the notion of credence is a useful and informative theoretical posit for purposes of modeling human epistemic rationality. Theoretical models, it is urged, frequently deploy idealizing assumptions that are actually false; nevertheless, these can be good models, which provide genuine explanatory understanding and/or genuine normative insight.

Remarks of this kind, I take it, are intended to suggest pertinent analogies to scientific models of natural phenomena. For example, Isaac Newton put forward highly accurate, explanatorily powerful, mathematical models of the trajectories of celestial objects through the night sky; and in doing so he relied on an idealizing assumption that he knew perfectly well was false—viz., the assumption that these objects all are *point-masses*.

If one posits credences in one's epistemological theorizing, even though one concedes that the idea that humans really instantiate such psychological states is a myth, then is one engaging in a kind of modeling that is pertinently similar to what Newton was doing when he made the idealizing assumption that celestial objects are point-masses?

I very much doubt it. One feature that is common to many idealizing assumptions deployed in scientific modeling is this: the idealizing assumption(s) introduce only a relatively minor discrepancy between the actual workings of the explanatorily pertinent factors and the model's characterization of those workings—and one has very good reason to think so. In such cases, one greatly simplifies one's representation of the operative factors, and often one also greatly simplifies the mathematical application of theoretical laws to one's specification of the specific circumstances.

Is anything of *that* kind going on when one posits these fictional psychological entities—credences—in an effort to illuminate the norms governing epistemic rationality? That is, is one basically just simplifying the pertinent realities of human psychology in an idealizing way, and in a way that one knows is retaining all but minor and negligible psychological factors that are pertinent to epistemic

rationality? Prima facie, it hardly seems so. On the contrary, it seems that one is *substantially distorting* the pertinent realities of human psychology, by positing idealized psychological states that not only are unreal themselves, but also are no mere minor approximation to anything that is in fact psychologically real. Such reality-distortion, when defended on the grounds that it supposedly can facilitate useful modeling of real epistemic normativity, should be regarded as “guilty until shown innocent.”

Well, what genuine psychological phenomena might there be which one might think are aptly modeled by the positing of credences? The best candidates, I take it, are psychologically real states whose content involves matters of *qualitative* likelihood, relative to a body of available evidence. For example, an epistemic agent might regard a given proposition as *highly likely* (given the evidence), or *somewhat more likely than not*, or *equally as likely as not*, and so on. Similarly, an epistemic agent might regard proposition *p* as *much more likely* than proposition *q*, or as *somewhat more likely* than *q*, or as *equally as likely as* proposition *q*, and so on. Qualitative-degree judgments about likelihood, both absolute and comparative, 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.

So I suggest that the question of the viability of credence-positing as modeling really comes down to this: Does one non-distortingly and illuminatingly model psychological realities concerning matters of *qualitative* likelihood (both absolute and comparative) by theorizing about idealized agents whose psychology includes quantitatively precise *credences* for virtually all the propositions that those idealized agents can entertain and understand?

Speaking for myself, I very much doubt it. Theorizing this way is an instance of something more generic: viz., trying to model the vague and imprecise via modeling-apparatus that is non-vague and precise. By my lights, modeling of that kind is bound to be more distorting than illuminating. To illustrate, consider the most popular way among philosophers of modeling the semantical workings of vagueness: viz., supervaluationism. The phenomenon to be modeled can be usefully described as “No sharp status-transitions among successive items in a sorites sequence.” For example, no sharp boundaries between heaps and non-heaps, or between bald guys and non-bald guys, etc. Supervaluationist semantics treats a vague predication as true simpliciter if it is true under all “permissible precisifications” of our

vague discourse, as false simpliciter if it is false under all such permissible precisifications, and otherwise as neither true nor false.

So, suppose that one sets forth supervaluationist semantics for vague discourse, and that one does so within a metalanguage governed by classical logic and classical two-valued semantics. Well then, one's logically precise semantical account of vagueness ends up committed, in spite of itself, to sharp boundaries among successive items in a sorites sequence: one sharp boundary between the bald guys and the neither-bald-nor-not-bald guys, and another one between the latter guys and the not-bald guys; one sharp boundary between the piles of sand that are heaps and those that are neither-heap-nor-not-heap, and another one between the latter piles and those that are not-heaps; etc. The account thereby distorts the fundamental essence of vagueness, rather than adequately modeling it.<sup>12,13</sup>

Likewise *mutatis mutandis*, I say, for the attempt to model psychological facts about qualitative-likelihood judgments via the theoretical machinery of credences. Doing so is much more likely to distort the phenomena of interest, rather than illuminating them.

Fans of BFE as modeling will perhaps respond by first stressing the *normative* character of BFE, and then claiming that the "ideal Bayesian reasoner" serves as a normative ideal for actual *human* reasoners. The guiding thought is this: even if humans do not have precise, quantitative, credences for most propositions they can entertain and understand, and indeed even if humans are psychologically incapable of having such credences for most such propositions, nevertheless one should strive to approximate, as best one can, the credences that would be possessed by an ideal Bayesian reasoner who possessed just the same evidence as oneself.

But the problem now is that, according to orthodox versions of BFE themselves, various ideal Bayesian reasoners who all possess just the same evidence as oneself (both synchronically and also diachronically) will have different credences from one another—because their respective histories of credence-updating by conditionalization will have commenced from different prior credences.<sup>14</sup> So, since there is no good reason to suppose that some unique one of these ideal Bayesian reasoners is more similar psychologically to oneself than all the others, Bayesian BFE provides no determinate normatively ideal "target" for one to aspire to emulate.<sup>15</sup>

At this point in the dialectic, yet another move can be expected from the fans of BFE as modeling: they are apt to urge a move away from *point-valued* credences, and toward something like

*interval-valued* credences. The thought will be that this move is a way for the technical modeling-machinery to come to terms with the fact that psychologically real likelihood-judgments (both absolute and comparative) normally are only qualitative rather than quantitative, and are inherently so.

To which I respond: Look again at the distorting effect of supervaluationist models of vagueness, insofar as those models are propounded within a metalanguage that is governed by classical logic and classical two-valued semantics. One gets sharp boundaries in spite of oneself, in the course of trying to model “no sharp boundaries.” Likewise, *mutatis mutandis*, for the appeal to interval-valued credences. Those beasties are sharply bounded, whereas categories like “highly likely, relative to the available evidence,” are not sharply bounded. Thus there is no reason to suppose that one’s own rationally appropriate *qualitative* degrees of belief in various propositions are more similar to some single specific assignment of sharp interval-valued probabilities to those propositions than to any of various other incompatible such assignments. So the same problem re-appears that was stressed two paragraphs ago: no determinate normatively ideal “target” is provided for one to aspire to emulate. When you model the vague and non-sharp via the precise and sharp, you distort the phenomena you are trying to model, rather than illuminating those phenomena.<sup>16</sup>

The upshot is that there is a heavy burden of proof on the fans of BFE as a putative modeling framework—viz., to explain to us skeptics why your framework *does not* badly distort the very psychological and epistemic-normative phenomena you seek to illuminate. Unless that burden can be successfully discharged, the possibility is all too real that the discipline of BFE is relevantly similar to previous disciplines like phlogiston theory and alchemy. Brief and sketchy remarks about credence as a putatively useful construct for modeling purposes will not suffice.<sup>17,18</sup>

## 5. Carnapian Formal Epistemology?

A potential response to what I say above would be claim that there is an epistemically relevant kind of probability that is inherent to evidence-relations themselves, that adheres to virtually any proposition that an epistemic agent can entertain and understand (relative to that epistemic agent’s body of available evidence), and that is not to be understood as some kind of psychological state (either real or ideal). One such view, advocated (among others) by Carnap, construes epistemic probability as a generalization of logical entailment. The leading idea is this: for any proposition  $p$ , an epistemic agent’s

corpus of available evidence bears some specific, quantitative, 0-to-1 ratio-scale, degree of *partial logical entailment* to  $p$ —this being  $p$ 's epistemic probability, for the agent. (Ordinary logical entailment is the limit case, with  $p$ 's epistemic probability being 1.)

If this Carnapian construal of epistemic probability as “logical” probability were plausible and viable, then epistemic-probability invoking formal epistemology would be back in business—still tied to the idea that normative epistemological theorizing should primarily involve normative principles governing epistemic probability, but no longer tied to the idea that epistemic probabilities are so-called “credences.” But nowadays it is widely acknowledged that the notion of partial logical entailment has never been worked out satisfactorily, despite Carnap’s heroic efforts to do so. Those who call themselves “Bayesians” typically think so; indeed, this is major reason why they fall back upon a psychologistic construal of epistemic probability. For the Bayesians, and for me too, it seems very likely (a *qualitative* judgment!) that the notion of logical entailment simply is not subject to the sought-for kind of generalization—one that would confer precise, quantitative, degrees of partial logical entailment on each proposition, relative to a given evidential corpus. (Typically, I believe, the degree of confirmation possessed by a proposition  $p$ , relative to a given corpus of evidence, will be at best a *qualitative* matter—e.g., highly likely, somewhat likely, equally likely as not.)

Nevertheless, I think there is an important kernel of truth in the idea that epistemic probability is partial logical entailment—viz., the more generic idea that epistemic probability is quantitative, 0-to-1 scale, degree of *evidential support*, relative to a given body of evidence. This idea can be—and I think should be—implemented in a different way, rather than being implemented by trying to generalize the notion of logical entailment. This is the heart of my own conception of epistemic probability, to which I turn presently.

Before doing so, however, let me note another important difference between Carnap’s approach to epistemic probability and mine. Carnap was very ambitious with respect to the presumed *scope* of epistemic probability: he assumed that virtually any proposition would have a determinate, quantitative, degree of partial logical entailment, relative to a given body of evidence. (He sought a fully general, systematic, quantitatively precise, confirmation theory, based on partial logical entailment.) My approach, on the other hand, is very modest regarding the scope of epistemic probability: although I believe that

*qualitative* degrees of evidential support arise frequently, I also believe that *quantitative*, ratio-scale, degrees of evidential support only arise under quite special circumstances.

If I am right about this, then it will be bad news for the prospects of a non-psychologistic, epistemic-probability invoking, formal epistemology. Epistemic probabilities will not be eligible to play a principal and central role in epistemological theorizing, for the simple reason that there just won't exist enough epistemic probabilities to serve that purpose.

## 6. Epistemic Probability as Quantitative Degree of Evidential Support

As I said, I contend that epistemic probability is rightly construed as quantitative, 0-to-1 scale, degree of evidential support, relative to a given body of evidence. The basic idea is this: under certain circumstances, the *qualitative* relation of evidential equi-likelihood among two or more propositions can give rise to *quantitative*, 0-to-1 ratio-scale, degrees of evidential support for those propositions. The pertinent circumstances are as follows: (1) one knows (or justifiably believes) of  $n$  distinct propositions, for known (or justifiably believed)  $n$ , that these propositions form a partition of the pertinent space of possibilities, i.e., they are mutually exclusive and jointly exhaustive, and (2) these propositions are all equally likely, relative to one's available evidence. Under such circumstances, each of the propositions has epistemic probability  $1/n$ . This quantity is the quantitative, 0-to-1 ratio-scale, degree of evidential support possessed by each of the propositions in the partition. (I call the resulting assignment of probabilities to the respective propositions a *simple probability-partition*—for short, an SPP.)

This basic idea requires supplementation, in two ways. First, sometimes one will know (or justifiably believe), of one or several of the possibilities in an SPP, that it/they sub-divide into a partition of distinct sub-possibilities—where one knows (or justifiably believes) that these sub-possibilities form a partition themselves, one knows (or has a justified belief about) how many sub-possibilities there are, and each of the sub-possibilities is equally likely relative to one's available evidence. One will then be able to assign further, equi-likelihood based, epistemic probabilities to the sub-possibilities. (I call the resulting assignment of probabilities a *hierarchical probability-partition*—for short an HPP. An SPP is thus a special case of an HPP.) If the main possibility has epistemic probability  $1/n$ , and there are  $m$  sub-possibilities in the sub-partition, then the possibilities in the sub-partition each have epistemic probability  $(1/n \times 1/m)$ .

Second, sometimes some of the pertinent possibilities under consideration will be such that one knows their *objective chances* (or one justifiably believes of them that they possess those objective chances). If so, then *ceteris paribus*, each of these possibilities has a quantitative degree of evidential support—i.e., an epistemic probability—that is equal to its known (or justifiably believed) objective chance.<sup>19</sup> Some of the probability-assignments in a HPP can be based on such known (or justifiably believed) objective chances.

To illustrate this approach to epistemic probability in application, consider the setup in the famous Monty Hall problem. You are a contestant in a game show in which you know that there is a valuable prize behind one of three doors, and nothing behind either of the other two doors. Letting ‘ $P_i$ ’ symbolize ‘The prize is behind door  $i$ ’, the following SPP gives the epistemic probabilities of the prize location, based on evidential indifference concerning this matter:

*Table 1*

P1, 1/3

P2, 1/3

P3, 1/3

Suppose that you choose door 3. Monty informs you that he knows the location of the prize, and that he will now open a door that (i) is not your chosen door, and (ii) he knows has no prize behind it. Letting ‘ $O_j$ ’ symbolize ‘Monty opens door  $j$ ’, the following HPP gives a more elaborated assignment of pertinent epistemic probabilities, based on two kinds of evidential indifference—the first concerning the prize location, and the second concerning which door Monty opens if he has an option:

*Table 2*

P1, 1/3

O2, 1/3

P2, 1/3

O1, 1/3

P3, 1/3

O1, 1/6

O2, 1/6

The subcases in an HPP are conjunctive possibilities—in this example, possibilities of the form ( $P_i$  &  $O_j$ ).

Suppose that Monty now opens door 2. This excludes the fine-grained possibilities (P2 & O1) and (P3 & O1), thereby rendering each of the remaining fine-grained possibilities twice as probable as before. (Renormalizing in this way the probabilities of the non-excluded fine-grained possibilities—i.e., by preserving the ratios of their prior probabilities—implements updating-by-conditionalization. I call this *intuitive Bayesian updating*.) Thus, the correct updated epistemic probabilities for P1 and P3 are  $2/3$  and  $1/3$  respectively. So you should switch doors if Monty gives you the opportunity to do so.

For an application of my approach where known chance figures in, consider a variation of the Monty Hall problem in which Monty also informs you 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 either 4 or 5 or 6. Suppose, again, that you have chosen door 3. Here the pertinent epistemic probabilities are given in the following HPP—where the probabilities of the main cases are equal to their known chances<sup>20</sup> and the probabilities of the two sub-cases under P3 then arise by evidential indifference.<sup>21</sup>

*Table 3*

P1,  $1/3$

O2,  $1/3 (= 4/12)$

P2,  $1/6$

O1,  $1/6 (= 2/12)$

P3,  $1/2$

O1,  $1/4 (= 3/12)$

O2,  $1/4 (= 3/12)$

If Monty now opens door 2, thereby eliminating the fine-grained possibilities (P2 & O1) and (P3 & O1), then the updated probabilities for the remaining two fine-grained possibilities (P1 & O2) and (P3 & O2) are  $4/7$  and  $3/7$  respectively. (This results by intuitive Bayesian updating, i.e. renormalizing in the way that retains the 4-to-3 ratio of the prior probabilities of those two fine-grained possibilities—which implements updating-by-conditionalization.) Thus, the updated probabilities for P1 and P3 are  $4/7$  and  $3/7$ . So you should switch doors if offered the opportunity.

I contend that epistemic probability is quantitative, 0-to-1 ratio-scale, degree of evidential support, relative to a body of available evidence. But what counts as “available evidence”? In my view this is somewhat context-dependent matter, depending on one’s specific purposes in doing probabilistic reasoning. Typically, I think, the evidence must be *psychologically available* to the epistemic agent for whom the epistemic probabilities arise. Also, typically the evidence must be fairly readily *retrievable* by the agent, in order to count as psychologically available. But in some contexts it might be appropriate to relax such psychological assumptions—e.g., contexts in which all logical and mathematical truths, no matter how abstruse, count as having epistemic probability 1. (In most contexts, I think, highly abstruse logical and/or mathematical truths either have no epistemic probability at all for a given agent, or else have epistemic probability  $\frac{1}{2}$  by virtue of evidential indifference as between truth and falsity.)

What epistemic status must a given proposition have for a given agent, in order for that proposition to be a component of the agent’s “available evidence”? This too I regard as a somewhat context-dependent matter. Typically, such a proposition must be *epistemically justified* for the agent, and indeed must be justified in such a way that it counts as knowledge provided that it is true. In some contexts, perhaps, only propositions that *are* known by the agent—and hence are true—will count as available evidence. And perhaps there are contexts where the evidence need not meet either condition, as long as it is present in the agent’s environment and is readily enough available.

## 7. Evidentialist Formal Epistemology?

Suppose that one accepts my critique of the kind of formal epistemology that gives center stage to the notion of credence, and that one also holds that there is no hope for a Carnap-style confirmation theory based on the notion of logical probability. One might yet seek to develop an alternative kind of epistemic-probability based formal epistemology—one that embraces my contention that epistemic probability is 0-to-1, ratio-scale, degree of evidential support (relative to a given evidential corpus). The idea would be to articulate epistemological norms that determine, for virtually any proposition  $p$  that an epistemic agent might contemplate and understand (and relative to the agent’s evidential corpus  $E$ ), a specific, quantitative, degree of evidential support that  $E$  confers upon  $p$ .

Yes, one might seek to develop that kind of formal epistemology. But one shouldn’t, for the following reason. If indeed epistemic probabilities are quantitative degrees of evidential support that arise

in the ways I have described, then such probabilities arise only relatively rarely. On one hand, only relatively rarely does one have knowledge (or justified belief) about objective chances of the possibilities one is considering.<sup>22</sup> And on the other hand, only relatively rarely does the pertinent space of possibilities constitute a partition (perhaps hierarchical) exhibiting the strong kind of symmetry, vis-à-vis one's evidence, that generates evidential-indifference based epistemic probabilities for all possibilities in the partition other than those (if any) whose epistemic probabilities are their known (or justifiedly believed) chances.<sup>23</sup> Most propositions that an epistemic agent can contemplate and understand will not have any epistemic probability at all, relative to the agent's corpus of available evidence. Instead, normally a given proposition will have at best some *qualitative* degree of evidential support, relative to that corpus.

## 8. Conclusion

Theorizing using numbers is tempting. After all, there is lots that we know how to do with numbers (adding, multiplying, and much more); also, with mathematics comes rigor. Fans of BFE have become gripped by the pull of such number-temptation. But BFE badly distorts the phenomena of epistemic rationality that it seeks to illuminate. A better approach to epistemology will be one that takes to heart the famous remark of Aristotle about ethics: "Our account...will be adequate if it achieves such clarity as the subject-matter allows; for the same degree of precision is not to be expected in all discussions, any more than in all products of handicraft." (*Nicomachean Ethics*, Book 1, Chapter 3)<sup>24</sup>

## References

- Elga, Adam. (2000). Self-Locating Belief and the Sleeping Beauty Problem. *Analysis* 60 (2): 143-147.
- Elga, Adam. (2010). Subjective Probabilities Should Be Sharp. *Philosophers' Imprint* 10 (5): 1-11.
- Horgan, Terence. (1994). Robust Vagueness and the Forced-March Sorites Paradox. *Philosophical Perspectives* 8 (Logic and Language): 159-188. Reprinted in Horgan (forthcoming b).
- \_\_\_\_\_. (1995). Transvaluationism: A Dionysian Approach to Vagueness. *Southern Journal of Philosophy* 33 (Supplement 1): 97-125. Reprinted in Horgan (forthcoming b).
- \_\_\_\_\_. (1997). Modeling the Noncomputational Mind: Reply to Litch. *Philosophical Psychology* 10 (3): 365-371.

- \_\_\_\_\_. (2004). Sleeping Beauty Awakened: New Odds at the Dawn of the New Day. *Analysis* 64 (1): 10-20. Reprinted in Horgan (forthcoming b).
- \_\_\_\_\_. (2008). Synchronic Bayesian Updating and the Sleeping Beauty Problem: Reply to Pust. *Synthese* 160 (2): 155-159. Reprinted in Horgan (forthcoming b).
- \_\_\_\_\_. (2010). Transvaluationism about Vagueness: A Progress Report. *Southern Journal of Philosophy* 48 (1): 67-94. Reprinted in Horgan (forthcoming b).
- \_\_\_\_\_. (2015) Generalized Conditionalization and the Sleeping Beauty Problem, II. *Erkenntnis* 80 (4): 811-839. Reprinted in Horgan (forthcoming b).
- \_\_\_\_\_. (Forthcoming a). Epistemic Probability. Chapter 14 in Horgan (forthcoming b).
- \_\_\_\_\_. (Forthcoming b). *Essays on Paradox*. Oxford and New York: Oxford University Press.
- Horgan, Terence and Tienson, John. (1990). Soft Laws. *Midwest Studies in Philosophy* 15 (1): 256-279.
- \_\_\_\_\_. (1996). *Connectionism and the Philosophy of Psychology*. Cambridge, MA: MIT Press.
- \_\_\_\_\_. (1997). *Precis of Connectionism and the Philosophy of Psychology*. *Philosophical Psychology* 10 (3): 337-356.
- Horgan, Terence. and Timmons, Mark. (2009). What Does the Frame Problem Tell Us about Moral Normativity? *Ethical Theory and Moral Practice* 12 (1): 25-51.
- Horgan, Terence. and Mahtani, Anna. (2013). Generalized Conditionalization and the Sleeping Beauty Problem. *Erkenntnis* 78 (2): 333-351.
- Horgan, Terence. and Potrč, Matjaž. (2008). *Austere Realism: Contextual Semantics Meets Minimal Ontology*. Cambridge, MA: MIT Press.
- \_\_\_\_\_. (Forthcoming). Ontological Vagueness: Why It's Impossible, and Why the Metaphysical and Semantical Consequences are Severe. In *The Sorites*, edited by Ali Abasnezhad and Ottavio Bueno. Berlin: *Synthese* Library, Springer Verlag.
- Mahtani, Anna. (2016). Imprecise Probabilities and Unstable Betting Behavior. *Nous* 50 (3), preprint.
- Meacham, Christopher. (2013). Understanding Conditionalization. *Canadian Journal of Philosophy* 45 (5-6): 767-797.

---

<sup>1</sup> The label ‘formal epistemology’ can be applied to enterprises other than BFE—for instance, the enterprise of systematizing epistemic normativity via an axiomatic theory consisting of exceptionless general normative principles—while perhaps also seeking to express these principles in a formal-language format that could constitute a computer program. Thus generically described, this enterprise is orthogonal to BFE, because one might seek such normative principles that do not mention credence at all, and also do not mention probability construed some other way. I also have serious doubts about the generic enterprise of fully systematizing epistemic normativity (although these doubts are not my present topic), because I believe that epistemic normativity (and likewise moral normativity) is too subtle and too complex to be fully systematizable by exceptionless general principles; cf. Horgan and Tienson (1996, 1997), Horgan (1997), Horgan and Timmons (2009). A more viable enterprise for epistemology (and likewise for ethics), I think, is to seek general normative principles that have ineliminable *ceteris paribus* clauses, thereby systematizing epistemic normativity (/moral normativity) partially but not completely. Such principles were dubbed “soft laws” in Horgan and Tienson (1990).

<sup>2</sup> The label ‘epistemic probability’ sometimes has been appropriated for a specific conception of the kind of probability that figures most directly in matters of epistemic rationality—e.g., the idea that the probability conferred on a proposition by a body of evidence is the degree of “partial logical entailment” between the evidence and the proposition—rather than being employed in my more generic way. By my lights, this is an instance of what I call *the tyranny of terminology*: it tends to make invisible any potential alternative conceptions of rationality-involving probability that might be equally worthy of the label—including my own conception, to be described below.

<sup>3</sup> This section and section 5 below draw on material from Horgan (forthcoming b).

<sup>4</sup> My use of the word ‘likely’ in such characterizations of qualitative degrees of confidence is different, of course, from the stipulative, technical, use of ‘likelihood’ in statistics under which the “likelihood” of  $A$  on  $B$  is the conditional probability of  $B$  given  $A$ .

<sup>5</sup> This entails that terminology expressing qualitative degrees of confidence is not rightly understood as involving implicit quantification over numerical point-values (or over numerical interval-values) of confidence. Thus, “ $p$  is highly likely” does not say that  $p$  has some numerical confidence-value close to 1, because  $p$  need not have any numerical confidence-value at all for an epistemic agent who regards  $p$  as highly likely. Indeed, I contend that it is almost always the case that people’s qualitative-confidence judgments occur without any underlying psychological states that constitute quantitative confidence-levels.

<sup>6</sup> A referee remarks that my skeptical doubts about the psychological reality of quantitative degrees of confidence also seem applicable, *mutatis mutandis*, concerning the putative psychological reality of quantitative degrees of desire/desireability—i.e., concerning the putative psychological states called “subjective utilities” in decision theory.

---

Indeed so. And in fact I am doubly skeptical about decision theory as a general normative model of practical rationality, because I believe (1) that normally the various possible outcomes of an envisioned act will only have *qualitative* degrees of comparative and/or non-comparative desirability for an agent, and (2) that normally the various possible states of nature will only have only qualitative degrees of comparative or non-comparative likelihood for the agent. To say that decision theory cannot serve as a *general* normative model of rationality does not mean that it has no applications at all, however. Sometimes an agent can and does have quantizable preferences concerning the available potential outcomes (e.g., preferences over monetary outcomes that are linear with the quantity of money obtained), and also assigns principled epistemic probabilities to the alternative possible states of nature; in such situations, normative decision theory becomes applicable. In Section 6 I explain my own construal of epistemic probabilities, and I discuss the conditions under which I think they arise.

<sup>7</sup> A referee points out that this observation can be accommodated by construing strength of belief not as a position on a partial-belief scale, but in some other way—say, as a measure of how robust a subject’s credence is in light of various pieces of possible future evidence (so-called “resilience”), and/or as a measure of how strongly the subject believes that the evidence supports the belief. Fair enough. But although the conflation is potentially avoidable in such ways, it often seems to occur anyway. Consider, for instance, the contemporary epistemological debate about rational disagreement between epistemic peers. Those who deploy the conceptual machinery of BFE in addressing this debate always construe “conciliation” with an epistemic peer about a proposition *p* as a matter of altering one’s credence for *p* in the direction of the peer’s credence for *p*; likewise, they always construe “steadfastness” as a matter of retaining one’s original credence for *p*. It never occurs to them that conciliation might take the form of a change in strength that accompanies persistence of one’s original credence.

<sup>8</sup> Those who espouse this view typically assume uncritically that certainty differs from ordinary full-fledged belief. I think this is too quick and too flat-footed. Although the notion of certainty *can* be used in a way that is more demanding than the notion of belief (as when one speaks of *absolute* certainty), I maintain that it is not always used that way. Rather, ‘certain’ is governed by contextually implicit parameters, and in many contexts those parameters work this way: any belief that qualifies as knowledge—or that would so qualify if it were true—thereby qualifies, in context, as certain. (Witness the extreme unacceptability of assertions or judgments of the form “I know that *p*, but I am not certain whether *p*.”) And ordinary full-fledged beliefs often *do* qualify as knowledge, or would so qualify if true. (Admittedly, there also are contextually appropriate uses of ‘believe’ under which one employs this term (especially when stressed) to express significant doubt—as in “I *believe* this road leads to Timbuktu, but I’m not sure”—but these uses are not pertinent here.) Thus, there need not be any conflict between equating degree-one

---

credence with full-fledged belief and also equating degree-one credence with certainty, because in many contexts the notions of belief and certainty coincide with one another.

<sup>9</sup> A radically revisionary variant of this alternative construal of credence would instead advocate simply dropping the notion of belief altogether, rather than maintaining that it coincides with some close-to-one level of credence. My remarks in the present section and the next will apply to this revisionary version too, *mutatis mutandis*—apart from the one exception cited in footnote 11.

<sup>10</sup> Strictly speaking, it need not do so, because it might acknowledge differences in strength of belief and might invoke some construal of them other than different degrees of sufficiently-high credence; cf. note 7 above. However, doing this would significantly undermine the epistemological centrality that this version of BFE assigns to the notion of credence. For example, it would longer be unproblematic to construe “steadfastness” in a peer-disagreement situation as retaining one’s credence for proposition  $p$  despite one’s peer’s differing credence, or to construe “conciliation” as altering one’s credence for  $p$  in the direction of the peer’s credence.

<sup>11</sup> This problem is avoided by the radically revisionist construal of credence mentioned in the footnote 9.

<sup>12</sup> In my view, such distortion is very substantial. Genuine vagueness, I maintain, involves the following kind of inherent logical incoherence: the vague statuses of the respective items in a sorites sequence—statuses like *baldness* or *heaphood*—are governed by status-principles that are mutually unsatisfiable. This kind of inherent incoherence does not preclude vagueness from arising in language and thought, because the mutually unsatisfiable status-principles can be honored in one’s affirmatory practice. (Never assign incompatible statuses to adjacent items in a sorites sequence; never assign statuses to all items in such a sequence.) But it renders ontological vagueness impossible—which has enormous consequences for metaphysics. Texts in which these claims are elaborated and defended include Horgan (1994, 1995, 2010), Horgan and Potrč (2008, forthcoming).

<sup>13</sup> One way to avoid commitment to sharp boundaries is via metalinguistic iteration of one’s semantic machinery: treat the first-order metalanguage as vague, and hence as subject to supervaluationist semantics expressible in the second-order metalanguage; likewise for the second-order metalanguage; and likewise *ad infinitum*, all the way up the metalinguistic hierarchy. To this I have no objection, although I maintain that it does not by any means eliminate the specific kind of logical incoherence that I claim is inherent to vagueness.

<sup>14</sup> If such ideal Bayesian reasoners update their credences by Jeffrey conditionalization rather than by standard conditionalization, then there is yet another reason why these ideal reasoners can differ from one another in their credences—viz., that an evidentially pertinent experience can “jolt” the pre-updating credences of one such reasoner differently than this same experience jolts the corresponding pre-updating credences of another such reasoner—even

---

if the affected pre-jolt credences are the same for the two reasoners. (See the penultimate paragraph of Section 3.) This in turn would lead to different credence-updatings by those respective reasoners.

<sup>15</sup> This paragraph was prompted by Jonah Schupbach’s commentary on an earlier version of the present paper, at the 2016 *Res Philosophica* conference “Bridging Formal and Traditional Epistemology.”

<sup>16</sup> Anna Mahtani (2016) argues, convincingly in my view, that the so-called “two-bets” scenario, deployed by Adam Elga (2010) in support of the claim that rational agents cannot have credences that are interval-valued but not point-valued, actually supports not only Elga’s claim but also the claim that rational agents cannot have point-valued credences either. Mahtani’s overall argument, I think, thereby bolsters my contention that credence-involving models—regardless whether they posit point-valued credences or instead posit interval-valued credences—distort epistemic rationality rather than illuminating it.

<sup>17</sup> Rejecting BFE as a general normative framework for modeling epistemic rationality does not entail that issues discussed by fans of BFE are of no genuine interest. Rather, many of those issues can be re-framed: not as issues about epistemic rationality per se, but rather as issues about normatively appropriate reasoning about epistemic probabilities. Among these are questions about normatively appropriate updating of epistemic probabilities in light of newly acquired information. Also, new issues arise too about apt probabilistic reasoning, once one eschews the notion of credence and instead construes epistemic probability in the way I describe in Section 6 below. Among these new issues are questions about when, and why, principled epistemic probabilities arise in the first place, and questions about when, and why, certain generalized versions of standard conditionalization are normatively appropriate in updating one’s epistemic probabilities. I discuss these matters in Horgan (forthcoming a) and also in several other chapters of Horgan (forthcoming b).

<sup>18</sup> Remarks analogous to those in the preceding footnote also apply to decision theory. As I said in footnote 6 above, decision theory is sometimes applicable to real-life decision problems, even though I think it cannot provide a general normative model of practical rationality. Issues debated in decision theory—e.g., the debate between advocates of causal decision theory and advocates of evidential decision theory—can be framed as pertaining specifically to decision problems in which principled epistemic probabilities and quantitative utilities exist. Some of those issues have analogues that arise for a wider class of decision problems—for example, the original version of Newcomb’s problem, in which it is specified that the predictor is *very likely* to have correctly predicted your upcoming choice (but no determinate probability of correctness has been stipulated). Several chapters in Horgan (forthcoming) address both versions of Newcomb’s problem, including its import for decision theory. Also, some chapters in Horgan (forthcoming b), discussing the two-envelope paradox, argue that because epistemic-probability contexts are weakly hyper-intensional, familiar definitions of expected utility in decision theory sometimes generate

---

several kinds of “non-standard” expected utility for a given decision problem that rank the available acts differently and hence cannot all be maximized; this raises the question of when, and why, a given kind of non-standard expected utility *should* be maximized, a question addressed in those essays.

<sup>19</sup> *Cetera* can fail to be *paria* in certain epistemic circumstances. One such circumstance, I maintain, is illustrated by the infamous Sleeping Beauty problem. On Monday when Sleeping Beauty has been awakened in the lab by the sleep experimenters, she does not know whether today is Monday or Tuesday, and hence does not know whether or not the Monday night coin-flip has yet occurred. I claim—and so do many others—that in this situation her epistemic probability for the proposition *The outcome of the coin-flip is heads* is  $1/3$ , even though the known chance of that proposition is  $1/2$ . Cf. Elga (2000), Horgan (2004, 2008), Horgan and Mahtani (2013), Horgan (2015).

<sup>20</sup> On some views of objective chance, once an event occurs or fails to occur its objective chance becomes either 1 (if it has occurred) or 0 (if it has failed to occur). But matters of “known chance” can remain pertinent to *epistemic* probability even after the fact of occurrence or non-occurrence, provided that one does not know the actual event-outcome but one does know the pre-outcome chances. So it is with this version of the Monty Hall problem, with respect to the possible outcomes of the die-toss. (Similarly, in the Sleeping Beauty problem the proposition *The outcome of the coin-flip is heads* has a known chance of  $1/2$  for Beauty on any day in which she is awakened by the experimenters in the lab, whether or not the coin has yet been flipped by that day.)

<sup>21</sup> In saying that the probabilities of the two sub-cases under P3 arise by evidential indifference, I am assuming (in effect by implicit stipulation, which I hereby make explicit) a version of the Monty Hall scenario in which your total evidence gives you no reason to regard (P3 & O1) as either more likely or less likely than (P3 & O2), despite the known (and unequal) chances of P1, P2, and P3. In the context of an actual game-show environment, however, one might well believe (say) that (P3 & O1) is less likely than (P3 & O2) because of one’s background beliefs about how the differing chances of P1 and P2 and P3 might well affect Monty’s motivation about which door to open if P3 obtains. E.g., one might now believe that (P3 & O1) is (qualitatively!) less likely than (P3 & O2). (This underscores an important fact about epistemic probability as I construe it: only under *very* special evidential circumstances do quantitative probabilities arise from qualitative relations of evidential indifference among the epistemically pertinent possibilities.) If you like, finesse this complication by varying the case this way: Monty also tells you that if the prize is behind your chosen door, then which of the other doors he opens depends on the outcome a fair-coin toss: if Heads then he’ll open the left-most unchosen door, and if Tails then he’ll open the right-most unchosen door.

<sup>22</sup> This is so even allowing that statistical information sometimes yields knowledge of objective chances regarding certain possibilities one is considering—e.g., the objective chance that one has AIDS after having learned that the result of one’s AIDS test was positive. Epistemology certainly can legitimately concern itself with normative issues

---

regarding statistical inference, including issues about statistically based probability-judgments. But that is a far less global enterprise than the kind of formal-epistemology project I am currently discussing.

<sup>23</sup> In Horgan (2015) I discuss in some detail the question of which kinds of symmetry in a hierarchical partition of possibilities ground indifference-based assignments of epistemic probabilities to possibilities in the partition, and which kinds do not. A moral of that discussion is that in general the requisite type of symmetry must be quite strong.

<sup>24</sup> My thanks to Jonah Schupbach for his comments on an earlier version of this paper at the 2016 *Res Philosophica* conference, to audiences at that conference and at the University of Konstanz, and to Miloud Belkoniene, Don Fallis, Kay Mathiesen, and an anonymous referee.