

Synchronic Bayesian updating and the Sleeping Beauty problem: reply to Pust

Terry Horgan

Published online: 23 November 2006
© Springer Science+Business Media B.V. 2006

Abstract I maintain, in defending “thirdism,” that Sleeping Beauty should do Bayesian updating after assigning the “preliminary probability” $1/4$ to the statement S: “Today is Tuesday and the coin flip is heads.” (This preliminary probability obtains relative to a specific proper subset I of her available information.) Pust objects that her preliminary probability for S is really zero, because she could not be in an epistemic situation in which S is true. I reply that the impossibility of being in such an epistemic situation is irrelevant, because relative to I, statement S nonetheless has degree of evidential support $1/4$.

Keywords Sleeping Beauty · Probability · Epistemic probability · Preliminary probability · Credence · Bayesian updating · Synchronic Bayesian updating

In Horgan (2004, 2007) I defend the “thirder” position concerning the Sleeping Beauty problem, by appeal to a form of reasoning about epistemic probabilities that in Horgan (2007) I call “synchronic Bayesian updating.” Pust (forthcoming) argues that my recommended treatment fails. Here I will argue that his objection, despite being initially plausible, does not succeed. (Since Pust describes very clearly both the Sleeping Beauty problem itself and my own approach to it, I will not rehearse those here.)

Bayesian updating of epistemic probabilities, as I construe it, is a matter of conditionalizing on what I call “preliminary probabilities” — i.e., epistemic probabilities that obtain relative to a *proper subset* of an epistemic agent’s total available information. So-called “prior probabilities” are a special case of preliminary probabilities; they are preliminary probabilities relative to a body of information that *previously* constituted one’s *total* pertinent information, but *now* (because of recently acquired additional information) constitutes only a *portion* of one’s total pertinent information. Bayesian

T. Horgan (✉)
Department of Philosophy, University of Arizona,
Tucson, AZ 85721-0027, USA
e-mail: thorgan@email.arizona.edu

updating on prior probabilities is therefore really just a special case of what should properly be called Bayesian updating. Sometimes the pertinent preliminary probabilities are not prior probabilities, but instead involve a range of essentially indexical, self-locational, epistemic possibilities that only arise by virtue of the agent's having recently lost relevant self-locational information (e.g., information about whether today is Monday or is Tuesday). Such *synchronic* Bayesian updating, as I call it, is what I claim should be employed in the Sleeping Beauty problem.¹

Pust does not raise an objection either to this generalized way of conceiving Bayesian updating, or to the claim that synchronic Bayesian updating is sometimes an appropriate way to reason about epistemic probability. Instead he maintains that my *specific* proposal concerning the Sleeping Beauty problem is undermined by the following argument (p. 4):

- (1) An epistemic probability is the degree to which an agent in some logically possible epistemic situation ought (rationally) to believe some statement.
- (2) Any logically possible agent in any logically possible epistemic situation ought to be absolutely certain that the statement 'I am conscious now' is true.
- (3) Thus, (when she is awake on Monday) Beauty's preliminary probability for 'I am conscious now' is one. [1, 2]
- (4) Beauty's preliminary probability for 'I am conscious now only if I am awakened today by the experimenters' is one.
- (5) Thus, Beauty's preliminary probability for 'I am awakened today by the experimenters' is one. [3, 4]
- (6) Beauty's preliminary probability for H2 with respect to the statement 'I am awakened today by the experimenters' is zero.

So, Beauty's preliminary probability for H2 is zero. [5, 6]

My response will involve rejecting premise (1) of this argument, and also rejecting line (3). Let me begin by commenting on how I take it that Pust intends premise (1) to be understood. On one construal, the premise says that an epistemic probability of a statement S, for an agent A in a logically possible epistemic situation E at a time t, is the degree to which A in E ought rationally (at t) to believe S. This claim about the nature of epistemic probabilities applies not only to fully updated probabilities, but also to preliminary probabilities. But this construal just obliterates the distinction between preliminary probabilities and updated probabilities. Although one might not wish to contest the claim that a *fully updated* epistemic probability of a statement S, for an agent A in an epistemic possible epistemic situation E, is the degree to which A in E ought rationally to believe S, one certainly should deny that a *preliminary* probability of S, for A in E, is the degree to which A ought rationally to believe S *in E itself*. After all, in E itself the agent A has additional relevant evidence that bears on S—evidence that is deliberately ignored when A assigns a preliminary probability to S.

This point can be further appreciated by attending carefully to ordinary Bayesian updating on prior probabilities, in typical updating situations—situations in which there is no call for *synchronic* Bayesian updating (because there has been no peculiar

¹ Bradley Monton has pointed out to me that what I call preliminary probabilities are familiar in the literature of philosophy of science under the label 'ur-probabilities', and are sometimes invoked as a proposed solution to the problem of old evidence. See, for instance, Barnes (1999), Glymour (1980, pp. 87–91), Howson (1984, 1985, 1991), Jeffrey (1995), and Monton (forthcoming).

recent loss in pertinent, indexical, self-location information, or anything of the sort). Prior probabilities are a special case of preliminary probabilities. At the point in time where one commences one's updating, these preliminary probabilities do not coincide with degrees of belief that are *now* (i.e., at that moment) rationally appropriate within the agent's actual present epistemic situation. Rather, insofar as one thinks of epistemic probability in terms of rational degree of belief, the preliminary probability *presently* possessed by a given hypothesis is the degree of belief that was *formerly* rationally appropriate, within the agent's prior epistemic situation before the acquisition of the pertinent new evidence. It is in that sense that the (present) preliminary probability is a "prior" probability.

How then should premise (1) be understood, vis-à-vis preliminary probabilities? What Pust actually has in mind, I take it, is something like the following (a generalization of the penultimate sentence of the preceding paragraph). The preliminary probability of a statement S, for an agent A in an epistemic situation E at a time t and relative to a proper subset P of S's total pertinent information in E (at t), is the degree of belief in S that *would* be rational for A to possess if A *were* in an epistemic situation E* in which S's total pertinent information is P. I will hereafter construe premise (1) in this way.

So construed, I admit, the premise does have considerable prima facie plausibility – as does the remainder of Pust's argument. The intuitive thrust of his objection is nicely captured in the following passage late in his paper, a passage which also fits very well with the present construal of premise (1):

Horgan's suggestion is, in effect, that we should determine what credence Beauty ought to have in each of the four hypotheses [viz., H1, H2, T1, T2] by considering what credence she ought to have in each hypothesis were she to (1) have all the evidence she actually has on Monday *except for* the evidence inconsistent with H2 and then (2) conditionalize upon that evidence. The difficulty with this suggestion is that the knowledge Beauty has which is inconsistent with H2 is her knowledge that she is awakened today by the experimenters. However, given her stipulated *certainty* regarding the conditions of the experiment, she *could not* lack this knowledge unless she lacked the knowledge that she is presently conscious. So Horgan's proposal must appeal to epistemic probabilities relative to an epistemic situation in which Beauty [a] has positive epistemic probability for each of the relevant statements and [b] has certain knowledge of the experimental protocol but [c] lacks knowledge that she is presently conscious. However, such an epistemic situation is impossible. (p. 6)²

My response is the following. Beauty's preliminary probabilities for hypotheses H1, H2, T1, and T2 should *not* be equated with the degrees of belief in these respective hypotheses that she ought rationally to have if she were in an epistemic situation with features [a]–[c]. Such an epistemic situation is indeed impossible. Rather, her preliminary probability for each hypothesis, relative to the body of information comprising "all the evidence she actually has on Monday *except for* the evidence inconsistent with H2," is *the degree of evidential support for that hypothesis, relative to the specified information*. Although I need not necessarily contest the claim that a *fully*

² A referee has posed the following objection, similar in spirit to the passage just quoted. If, while you are deliberating about epistemic probabilities, you are stripped of any reason to endorse the claim that you are conscious now, then you seem to be stripped of adequate reason to *trust* your current deliberations.

updated epistemic probability is, as premise (1) asserts, “the degree to which an agent in some logically possible situation ought (rationally) to believe some statement,” I do contest this claim as applied to preliminary probabilities like those that figure crucially in the Sleeping Beauty problem.³ Beauty should assign preliminary probabilities not by contemplating how she would assign non-preliminary probabilities in a certain epistemic situation in which she is not conscious (since she would not assign probabilities at all in such a situation), but rather by contemplating herself from a detached, third-person, perspective. As she might say to herself:

Statements H1, H2, T1, and T2 are all consistent with the conjunction of (1) the information I had on Sunday with (2) the information that today is either Monday or Tuesday. In particular, statement H2 is consistent with that specific information—notwithstanding the fact that if H2 were true, then today I would not be conscious. So, *relative to the information just mentioned, and only that information*, the four statements each have the same degree of evidential support. Since they are also mutually exclusive and jointly exhaustive, they therefore have preliminary probabilities of 1/4 each (relative to that evidence). (Likewise, the statement ‘I am conscious today’ has preliminary probability 3/4, relative to that evidence.) These preliminary probabilities ignore the fact that I am currently conscious, of course, which means that in assigning them I am considering my present self from a third-person perspective. Assigning a non-zero preliminary probability to H2, a statement entailing (given what I know about the experiment) that I myself am not presently conscious, feels a bit strange, to be sure—since obviously I could not be doing this if I were not *in fact* presently conscious. But hey, I’m a paragon of rationality! I understand clearly that preliminary probability is degree of evidential support relative to a proper subset of one’s total evidence. And, relative to the proper subset I am now considering, H2 does indeed have degree 1/4 of evidential support (as do H1, T1, and T2).⁴

³ I myself would make the following claims, although my reply to Pust requires only a special case of the second one. First, even if fully updated epistemic probability *necessarily coincides* with the degree to which an agent in some epistemic situation ought rationally to believe some statement, such probability is not *identical* to rational degree of belief. Rather, second, epistemic probability (either preliminary or fully updated) is *quantitative degree of evidential support, relative to a body of available evidence possessed by some agent in some epistemic situation*. (My reply to Pust appeals to a special case of this claim, concerning preliminary probability only.) Third, *fully updated* epistemic probability is degree of evidential support relative to the *total* body of available evidence possessed by some agent in some epistemic situation. Fourth, epistemic probability (either preliminary or fully updated) is not identical to rationally-constrained-but-otherwise-*subjective* degree of belief. Rather, fifth, epistemic probability (either preliminary or updated) is an *objective* feature arising from available information together with (1) objectively appropriate indifference principles and/or (2) warranted beliefs about matters of objective chance. (Indifference principles typically generate epistemic probabilities this way: when n specific statements are (relative to a given body of information) mutually exclusive, jointly exhaustive, and equi-likely under a suitable indifference principle, then each of these statements has epistemic probability $1/n$ (relative to that body of information).) Sixth, statements acquire determinate, quantitative, epistemic probabilities only relatively rarely, because the epistemic preconditions for the existence of such probabilities arise only relatively rarely. (And, while I am registering various claims I would wish to make about epistemic probability, let me also register my deep suspicion of the notion *quantitative degree of belief*.)

⁴ Beauty should reply in a similar vein to the objection posed in footnote 2. She should say, “When I assign preliminary probabilities relative to a body of information that *does not include* the information that I am conscious now, I do not thereby deprive myself of adequate reason to *endorse* this claim; on the contrary, I continue to know full well that I am conscious now, even while assigning preliminary probabilities relative to a portion of my total information that does not include this fact.”

Being a paragon of rationality, Beauty should go on from here to do synchronic Bayesian updating using her preliminary probabilities, thereby taking account of the additional evidence provided by the fact that she is now conscious. By conditionalizing on these preliminary probabilities, she correctly concludes that the probability of HEADS is $1/3$.⁵

References

- Barnes, E. C. (1999). The quantitative problem of old evidence. *British Journal for the Philosophy of Science*, 50, 249–264.
- Glymour, C. (1980). *Theory and evidence*. Princeton: Princeton University Press.
- Horgan, T. (2004). Sleeping Beauty awakened: New odds at the dawn of the new day. *Analysis*, 64, 10–21.
- Horgan, T. (2007). “Synchronic Bayesian updating and the generalized Sleeping Beauty problem. *Analysis*, 67, 50–59.
- Howson, C. (1984). Bayesianism and support by novel facts. *British Journal for the Philosophy of Science*, 35, 245–251.
- Howson, C. (1985). Some recent objections to the Bayesian theory of support. *British Journal for the Philosophy of Science*, 36, 305–359.
- Howson, C. (1991). The ‘old evidence’ problem. *British Journal for the Philosophy of Science*, 42, 547–555.
- Jeffrey, R. (1995). Probability reparation: The problem of new explanation. *Philosophical Studies*, 77, 97–101.
- Monton, B. (forthcoming). God, fine-tuning, the the problem of old evidence.
- Pust, J. (forthcoming). Horgan on Sleeping Beauty.

⁵ Thanks to David Chalmers, David Henderson, Jenann Ismael, Bradley Monton, Joel Pust, Mark Timmons, an anonymous referee, and members of my spring 2005 epistemology seminar for helpful comments and discussion.