Sleeping Beauty Without Self-location and Without Chance? [i]

Jenann Ismael

jtismael@u.arizona.edu

Slacker Bob

Bob is a slacker. He is about to take a test. He hasn't studied at all. It's a true/false test and he is just as likely to get an answer right as to get it wrong. His teacher tells him that half of the correct answers on the pretest will be True and half of them will be False, and that Bob will be given a pretest and a post-test. The pre-test will be used to gauge what students in the class are not getting right and need to focus on. The post-test will be constructed from the pre-test by the following rule. All of the answers that Bob gets right will appear on the post-test once. All of the answers that he gets wrong will appear on the post-test twice. He has to take the post-test in on-line. The questions will appear one at a time on the screen, in no particular order. He doesn't see the next answer until he has answered the one he is viewing. he can't revise his answers. And it all happens to quickly to allow him to keep track of whether a question has already appeared on the test. [ii]

Bob is told all of this, but he isn't told which of his answers were right and which were wrong. He figures he has an even chance of getting things right as wrong in the beginning, so he answers T to all of the questions on the pre-test.

Consider this question which appeared on the pre-test

P: The 5400093rd digit of the expansion of p is less than 5.

During the pre-test, Bob's credence in P is 1/2.

What probability should he assign when P is re-encountered on the post-test? He might reason as follows. Although he hasn't learned anything more about the expansion of p, he knows that, because of the way that the post-test was constructed a question encountered on the post-test twice as likely to be false as true, because he knows that false questions outnumber true ones by a factor of two on the post test. So his credence in P, when he encounters it on the post-test ought to be 1/3. But if he knows now that tomorrow when he is administered the post-test, he will be presented with the question and be in the position of assigning a probability 1/3 that it is true, shouldn't he assign that probability now? And why do his probabilities change? Where has he acquired new information or lost any old information?

Sleeping Beauty

Compare Bob's predicament to the familiar case of Sleeping Beauty. Sleeping Beauty is an ideal reasoner who knows she will be administered a short term amnesiac. She knows that after she falls asleep a coin will be flipped. If it lands heads, she will be awakened on Monday and asked: "What is the probability that the coin landed heads?'. She will not be informed which day it is. If the coin lands tails, she will be awakened on both Monday and on Tuesday and asked the same question each time. The amnesiac guarantees that, if awakened on Tuesday she will not remember being woken on Monday. How should she answer? Some say1/2. After all, the coin is fair and nothing new has been learned. Others say 1/3.[iii] After all, there are three possibilities: the coin landed heads and so she is now being asked on Monday, the coin landed tails and she is now being asked on Monday, and the coin landed tails and she is now being asked on Tuesday. Sleeping beauty has no more evidence for one possibility rather than any other. So she should consider the three possibilities as equally likely.

The differences[iv]

In both cases there is an apparently good argument for assigning a credence of 1/3 to a proposition to which one has previously assigned a credence ½, even though no new information that seems pertinent to the content of the proposition has been gained. But there are some interesting differences between the two cases. First, in Sleeping Beauty the coin flip is treated as an objectively chancy event. [v] In Bob's case, chance plays no role. The propositions to which Bob is assigning a probability are mathematical propositions. Second, in the case of Sleeping Beauty, the source of the problem has long been thought to be her self-locating uncertainty. The fact that she doesn't know whether it is Monday or Tuesday interacts with her credence that the coin landed heads in a confusing manner. In the case of Slacker Bob, self-locating ignorance plays no evident role. He knows what day it is and what time when he takes the post-test. He knows where he is.

What is essential to generating the puzzle? Self-locating knowledge not ignorance

Ironically, there is a weak form of self-locating knowledge, not ignorance, that unites the two cases. Bob knows how many questions will appear on the post-test and the distribution of true and false answers. He knows that every question will appear at least once, that false ones will appear twice, and that there are exactly twice as many false answers as there are true ones.[vi] But if questions on the post-test appear in no particular order and he is too forgetful to use information about whether a question has already appeared on the post-test to guide his answer if it appears again, then from his perspective every question he is confronted with on the post-test is effectively a random pick from a string in which he knows only that 1/3 of the answers are true.[vii] The puzzle is generated by a divergence between credence in the content ⌈P⌉ and the probability that tokens of "P" encountered in the post-test setting express a truth.[viii] And so self-location does play a role in shaping Bob's credence. He knows that questions are twice as likely to be included in the post-test if they are false and so questions encountered on the post-test are twice as likely to be false as true. Probability assignments made in that setting can exploit what is known about the setting.

Something similar is happening in Beauty's case. For Sleeping Beauty, waking events are twice as likely to follow tails tosses as they are follow heads tosses, and so coin tosses encountered in retrospect are twice as likely to be tails tosses as heads. What self-location is doing in both cases is locating question tokens in a reference class that is more heavily populated by false ones. The epistemic relevance of self-locating information in this weak sense is a pervasive phenomenon. We use this sort of self-locating information virtually all the time in forming credences.[ix] If I am in France, I assign a higher probability that the stranger in front of me is a French-speaker than I do if I am in Germany. If I am in Australia I assign a higher probability to the belief that the hopping form in the distance is a kangaroo than I do if I am in Canada even though internally my epistemic situation may be indistinguishable. Knowing where I am and how things go generally around here can and should make a difference to my credences. It makes a difference to how I assess the very same evidence. It makes a difference to how I update.

Let us separate credence in the content of P from the belief that a particular token of P expresses a falsehood. The post-test setting is one in which these two come apart. Bob's credence in the content ⌈P⌉ is unchanged between the pre-test and post-test settings, but he knows that tokens of false questions are overrepresented on the post-test by a factor of two to one, so he assigns each token of "P" in that setting a probability 1/3 of being true. Just as in Sleeping Beauty, the post-toss setting is one in which credence in ⌈H⌉ and the belief that a particular waking event is preceded by a heads tossing come apart.

One of the things that is called into question by SB is Van Fraassen's Reflection Principle which holds that if P is now certain that tomorrow P will have degree of belief x in R, while suffering no 'cognitive mishaps' between now and tomorrow, then P ought to now have degree of belief x in R. Is this a violation of the Principle of Reflection? The short answer is 'yes'. I can know now that there are different settings in which I will assign different probabilities to the same proposition. Bob is indifferent about the truth or falsity of the propositions on the test. He assigns 1/3 probability of truth to every question during the test. If he hasn't kept track of the number of times a question appears on the test, his probabilities will revert to ½. Nothing mathematical was learned at any stage. The only thing he has updated was his self-locating information: first on the pre-test, next on the post-test, and afterward out in the wild.[x] So the real lesson here is that differences in setting can make a difference to probabilities assigned to the same events in internally indistinguishable epistemic conditions. It is not surprising that the same evidence can have a different significance in different settings. In Bayesian terms, these sorts of differences are captured in Bayesian epistemology by changes in one's priors as one moves from one setting to another.

The Everett Connection; probability without ignorance

There is a corollary of interest to those that have seen a connection between Sleeping Beauty and the problem of probability in Everettian interpretations of quantum mechanics.[xi] In an Everett universe, there is no ignorance and there is no chance.[xii] Some have created a prospective role for Born probabilities by introducing a bit of closed-eye guessing by post-measurement observers before they've seen the results of the measurement. [xiii] The idea here is that the Born probability of x is the probability that an observer in an Everett universe attaches to x once a result has been registered in his branch of the wave-function, but before he has looked at it. But is there a role for probability outside this narrow, and artificially created window of self-locating ignorance? We could give some retrospective significance to the Born probabilities as the degree of surprise one should feel upon seeing a given result upon measurement, that of all of the observers in an Everett universe, he should be one of those seeing what he in fact sees. Or we could see the Born Probabilities as a measure of the amount of self-locating information he gets from observing a given result (his degree of specialness among post-measurement cohorts).

These would be additional ways of assigning epistemic import to the Born Probabilities, but a good deal of recent work on probability in an Everett world has focused on their practical role. Indeed, there was a time when the practical role of probability was more secure then its epistemic role, a time when the claim that rational degrees of belief must obey the laws of probability was defended by Dutch Book Arguments, which establishes conformity to the laws of probability as a norm of prudential rationality by showing that expected utility maximizers whose partial beliefs violate these laws can be induced to behave in ways that are sure to leave them less well off than they could otherwise be. I think that the case of Bob is instructive on this score.

To avoid ambiguity, let us call the probabilities that get fed into one's decision procedure along with one's utility function to guide action 'decision probabilities'. If Bob is acting to maximize his overall score, there is little question about what he should do. He has credence in the content of the questions (outside the post-test setting, he assigns probability ½ that the contents are true), but knowing what he does about the distribution of true and false questions on the post-test, he knows that answering 'false' to all gives him an expectation of getting two thirds right, whereas answering 'true' to all gives him an expectation of getting one third right. The situation is no different from a practical standpoint than it would be if each question occurred once on the post-test, but his score was calculated from a grading scheme that penalized him 2 points for answering 'true' to a false question and only 1 point for answering 'false' to a true question. The 1/3 probability – i.e., the probability that a random token encountered in the post-test setting expresses a truth - is the one that he should use in his utility calculations. We could give this epistemic import, but the *even without this epistemic import*, probabilities have a practical role. The practical role is not parasitic on their epistemic role. Bob might not have any interest in forming credences at all. He might devise his strategy for answering, fill out his test sheet, and hand it in without so much as glancing at the questions. His epistemic probabilities for the mathematical propositions that appear as questions going into the test were ½, his epistemic probabilities coming out of the test were ½, he didn't take the opportunity to learn something mathematical by keeping track of which questions appeared once and which ones appeared twice. His sole purpose is to maximize his score. And if his sole purpose is to maximize his score, what matters to him is the distribution of true and false answers on the post-test.[xiv]

Let's see if the Born Probabilities might have a similar practical significance, riding roughshod over the subtleties of Everett interpretation, we'll assume non-relativistic quantum mechanics and picture the branching as a physical process in which the universe literally splits into an infinite number of branches in random selections of which measurements results are exhibited in proportions reflected by the Born Probabilities.[xv] In this setting, the Born probability of x is the probability that a random pick from post-measurement observers will be observing an x result. There is no ignorance, self-locating or otherwise for the pre-measurement Everettian agent, and we'll suppose that he has no interest in the self-locating thoughts or degree of surprise of his downstream descendents. Because of the structure of the Everett universe, he cannot form preferences for one of his successors over the others, and this means that he cannot prefer one result to another if the only difference between them is that there is a reallocation of descendents among branches.[xvi] He knows that every possible outcome will be realized in some branch of the post-measurement universe, and he is making choices to maximize utility across branches in a setting in which only the distribution of results can matter.[xvii]

The practical role of probabilities for the Everettian agent mirrors the role that the distribution of true and false answers on the post-test played for Bob. If the Everettian agent's purpose is to maximize expected utility, given the indifference constraint on his options, then what matters to him is the distribution of measurement results across downstream branches. And provided we can give the Born Probabilities quantitative significance as reflecting that distribution, then I suggest that there is a prudential norm that says that the numbers that he should feed into his decision procedure to maximize utility are the Born Probabilities. Or, to put it differently, there is a prudential norm that says that he should take the Born Probabilities as his decision probabilities.

Conclusion

I'm stepping into heavily contested waters, but I think that the case of Slacker Bob promises to shed light on Sleeping Beauty. If the analogy is successful, it shows that neither chance nor self-locating ignorance play an essential role in generating the puzzle. If self-location plays any role in Slacker Bob, it is self-locating *knowledge* rather than ignorance that plays a role, and only in a way that is perfectly pervasive. If the analogy fails, it will be instructive to see where. In either case, I think the connection between the case of Slacker Bob and Everett is clearer and more illuminating than the various attempts to connect Everett to Beauty.

---

[i] I'd like to thank Kenny Easwaran, Mike Titelbaum, and Terry Horgan for extremely useful discussion. And I'd like to thank Alan Hajek, John Cusbert, and especially Bob Spekkens for excellent comments.

[ii] Or if we like, perhaps we administer an amesia-inducing drug between questions to keep him from using information about the number of times a question appears on the test in gauging his answer.

[iii] Elga, Adam. 2000. "Self-Locating Belief and the Sleeping Beauty Problem." Analysis 60:143–7.

[iv] "Sleeping beauty and self-location: A hybrid model", Nick Bostrom. Synthese, Vol. 157, No. 1. (July 2007), pp. 59-78
The Relevance of Self-Locating Beliefs, Michael G. Titelbaum Philosophical Review 2008 Volume 117, Number 4: 555-606;
"Confirmation in a Branching World: The Everett Interpretation and Sleeping Beauty", The British Journal for the Philosophy of Science (2011) 62(2): 323-342

[v] See Elga, Ibid. and Lewis, D. (2001). Sleeping beauty: reply to Elga. Analysis, 61(271), 171–176. Chance plays a central role in both discussions.

[vi] This is because his teacher told him that questions on the pre-test were evenly divided between true and false. This is needed so that there is no uncertainty about the length or distribution on the post-test. Thanks to Kenny Easwaran for this point.

[vii] The order in which questions on the post test matters because if there is an algorithm which tells Bob the order in which questions on the post-test are generated from those on the pre-test, he can use that to work out whether a question has appeared once or twice. Suppose, for example that the algorithm is this. first the whole string of questions is repeated in the order in which they appeared on the original test, followed by the string of repeated questions, in no particular order. So if there were 50 questions on the original test, the post-test would have the same first 50 questions, followed by repeats of the questions that he got wrong on the original. In this case, there is no puzzle about what Bob should do. He should maintain his original credences of 1/2 until question 51. Once he is on question 51, he knows that every question he sees is a repeat, and so he should have no credence in any of them. Similar remarks will apply to any ordering of which he is apprised beforehand. If he knows what the ordering is and he knows where he is in the ordering, he should assign pr=1/2 to each question the first time it appears and pr=0 the second time.

[viii] This was ensured by having the questions appear in a random order and saying that it happens to quickly to allow him to keep track of whether a given question has appeared once or twice. We could have also done it by administering to him an amnesia-inducing drug after each question. Or we could have just focused on the credence he assigns the very first question. The forgetfulness effectively means that he treats every appearance of a question, as it was the first.

[ix] See http://opinionator.blogs.nytimes.com/2012/01/15/do-thrifty-brains-make-better-minds/ for a nice discussion of the growing evidence that the brain is more of a machine for making educated guesses, bringing everything it knows about the current situation to bear on reading sensory cues, than a camera that registers and records

[x] If he has kept track of the number of times, he has learned something, and his probabilities will be 1 or 0 depending on whether it appeared once or twice.

[xi] Lewis, P. J. 2007a. Quantum Sleeping Beauty. *Analysis* 67: 59–65. Vaidman, L. and S. Saunders. 2001. On Sleeping Beauty controversy.

[xii] At least not in the ordinary sense, but see Ismael, "A Modest Proposal About Chance", J.Phil for a different view.

[xiii] Vaidman, L. 1998. On schizophrenic experiences of the neutron or why we should believe in the many-worlds interpretation of quantum theory. *International Studies in the Philosophy of Science* 12: 245–61.

[xiv] Bob is choosing an answering strategy whose utility will be measured by number of correct answers. A strategy has the form of a complete assignment of T or F to questions on the post-test. In practice, he answers in stages, but since he is indifferent to the contents, and unable to keep track of the number of times a question appears, seeing the questions can't help him improve his strategy, so he might as well hand in his policy without ever taking the test.

[xv] It is more common nowadays to think in terms of divergent worlds, but the differences, and the innumerable subtleties needed to bring this into line with the more sophisticated understanding of the Everett Universe, don't matter here. See Everett, H. 1957. 'Relative state' formulation of quantum mechanics. *Reviews of Modern Physics* 29: 454–62.
Saunders and Wallace, "Branching and Uncertainty", *Br J Philos Sci* (2008) 59(3): 293-305, and Wallace, D. 2002. Worlds in the Everett interpretation. *Studies in History and Philosophy of Modern Physics* 33: 637–61. Greaves, H. 2004. Understanding Deutsch's probability in a deterministic multiverse. *Studies in History and Philosophy of Modern Physics* 35: 423–56.

[xvi] This is a matter of some subtlety. The pre-measurement Everettian observer cannot form de re thoughts about any of his successors because all of his physical relations to them. He can pick them out de dicto, as 'the occupant of a branch at which such and such a result is registered', but a moment of thought will show that these de dicto attitudes can't support pre-measurement preferences for policies that benefit one over another of his descendents. If he preferred one of his successors to the others, then it would be rational for him to act to maximize utility for that successor. As it is, he is just locked into an undiscriminating practical policy that allows only the distribution of results across branches to matter in utility calculations. In Bob's case, there was a similar constraint on his choice of answering strategies, but it stemmed from his epistemic indifference to the contents of the questions that he knew he would be presented with. In an Everett Universe, it is enforced by the ontological symmetry in the relations that a pre-measurement observer bears to her downstream descendents.

[xvii] One way of thinking of it is that the Everettian agent is not dividing belief among a number of possibilities among which he is epistemically indfferent, he is dividing utility among a number of descendents among whom he is conatively, or practically, indifferent.