Sleeping Beauty awakened: new odds at the dawn of the new day

TERRY HORGAN

1. The story of Sleeping Beauty is set forth as follows by Dorr (2002):

> Sleeping Beauty is a paradigm of rationality. On Sunday she learns for certain that she is to be the subject of an experiment. The experimenters will wake her up on Monday morning, and tell her some time later that it is Monday. When she goes back to sleep, they will toss a fair coin. If the outcome of the toss is Heads, they will do nothing. If the outcome is Tails, they will administer a drug whose effect is to destroy all memories from the previous day, so that when she wakes up on Tuesday, she will be unable to tell that it is not Monday. (2002: 292)[1]

Let HEADS be the hypothesis that the coin lands heads, and let TAILS be the hypothesis that it lands tails. The Sleeping Beauty Problem is this. When Sleeping Beauty finds herself awakened by the experimenters, with no memory of a prior awakening and with no ability to tell whether or not it is Monday, what probabilities should she assign to HEADS and TAILS respectively?

Elga (2000) maintains that when she is awakened, P(HEADS) = 1/3 and P(TAILS) = 2/3. He offers the following intuitively plausible argument

[1] Sleeping Beauty also knows the following about the experiment, although Dorr neglects to say so: if the outcome of the coin toss is Tails, then she will be awakened *by the experimenters* again on Tuesday. This feature of the story is needed in order to guarantee that if the outcome is Tails, then her Tuesday awakening will leave her unable to tell that it is not Monday.

ANALYSIS 64.1, January 2004, pp. 10–21.

(2000: 143–44). If the experiment were performed many times, then over the long run about 1/3 of the awakenings would happen on trials in which the coin lands heads, and about 2/3 on trials in which it lands tails. So in the present circumstance in which the experiment is performed just once, P(HEADS) = 1/3 and P(TAILS) = 2/3.[2]

Lewis (2001) maintains, contrary to Elga, that when Sleeping Beauty is awakened, P(HEADS) = P(TAILS) = 1/2. He offers the following intuitively plausible argument (2001: 173–74). Initially on Sunday when Sleeping Beauty was still awake, surely the probabilities of HEADS and TAILS were 1/2 each. But only new relevant evidence should produce a change in probability, and when Sleeping Beauty is awakened she receives no new evidence that is relevant to HEADS v. TAILS. (She already knew on Sunday that she would find herself being awakened at least once by the experimenters with no memory of a prior awakening.) Hence, when Sleeping Beauty is awakened it is still the case that HEADS and TAILS each have probability 1/2.

I side with Elga in this dispute. But if indeed Lewis's argument is mistaken, then there should be a way of explaining why and how it goes wrong. The challenge is to make clear why Sleeping Beauty's evidence upon being awakened is relevant to HEADS v. TAILS – more specifically, why she has evidence that makes it the case that P(HEADS) = 1/3 and P(TAILS) = 2/3. If such an account cannot be given, then it would seem that Lewis is right after all. This would mean that in the Sleeping Beauty problem, what would happen over the long run if the experiment were repeated many times does not reflect the single-case probabilities of HEADS v. TAILS.[3,4]

2 Another plausible argument for this position is as follows. When Sleeping Beauty is awakened by the experimenters, P(HEADS *given that* today is Monday) = P(TAILS *given that* today is Monday) = 1/2, because if today is Monday then the coin-toss has not yet occurred. Also, P(TAILS and today is Monday) = P(TAILS and today is Tuesday), because her evidence is indifferent between those two possibilities. Under the laws of probability theory, and given Sleeping Beauty's background knowledge about her situation, these probability assignments entail that P(HEADS) = 1/3 and P(TAILS) = 2/3. For an elaborated version of this argument, see Elga 2000: 144–45.

3 For a plausible argument that the long-run frequencies do not reflect the single-case probabilities in the Sleeping Beauty problem, see §4 of Arntzenius 2002. Arntzenius does not endorse Lewis's position, however; instead he maintains 'that one should distinguish degrees of belief from acceptable betting odds, and that some of the time Sleeping Beauty should not have definite degrees of belief in certain propositions' (53–54).

4 Adopting Lewis's position would also mean that Elga's argument involving conditional probabilities (described in n. 2) should be rejected – perhaps by biting the bullet and siding with Lewis in his highly counterintuitive claims (1) that when Sleeping Beauty is awakened by the experimenters, P(HEADS *given that* today is Monday) =

Dorr (2002) also sides with Elga. Dorr employs a 'soritical argument by analogy': he appeals to the alleged parallelism between a case he constructs, in which it seems intuitively clear that the probabilities of HEADS and TAILS respectively are 1/3 and 2/3, and the original Sleeping Beauty case – where the putative analogy is bolstered by a sorites sequence of intermediate cases. Although I think that Dorr's reasoning is suggestive and on the right track, soritical arguments by analogy do run a serious risk of being slippery slope fallacies. I believe that what I say in the present paper captures the underlying spirit of Dorr's approach, but without any appeal to analogies or to soritical reasoning. I address the connection between my own argument and Dorr's below.

2. Upon being awakened by the experimenters and finding herself (as expected) with no memory of a prior awakening, Sleeping Beauty no longer knows what day it is; today might be Monday or it might be Tuesday. She also does not know whether the coin toss comes up heads or tails. She contemplates the following partition of statements, all pertaining to the current day:

H_1: HEADS and today is Monday
H_2: HEADS and today is Tuesday
T_1: TAILS and today is Monday
T_2: TAILS and today is Tuesday.

First she asks herself about the *prior* probabilities of these four statements, as determined by the evidence she possessed prior to being put to sleep. She rightly realizes that these prior probabilities are features these statements have *presently* – viz. probabilities the statements have *now* relative to the evidence available to her *then*.[5] Exercising her impeccable rationality, she correctly judges that each of the four statements has a prior probability of 1/4.[6] (The evidence available to her from Sunday is consistent with H_2. Although the Sunday evidence guarantees that she is awakened by the

2/3, and (2) that later on Monday when she is told that today is Monday, the unconditional probability of HEADS gets updated (via Bayesian conditionalization) from 1/2 to 2/3.

[5] The point is entirely general: the prior probability at time *t* of a statement S is a probability had by S *at t* – albeit a probability that is relative to a body of evidence that is known (at *t*) to be the evidence that was available *prior* to the most recent change in evidence. (Thanks to Sarah Wright for emphasizing this to me.) Typically – but not invariably – the prior probability of S at *t* is identical to the current-probability of S *at that earlier time* – i.e. the probability *then* possessed by S relative to the total evidence *then* available. More on this below.

[6] Likewise, she also correctly judges that the prior *conditional* probabilities of (H_1 given $-H_2$), of (T_1 given $-H_2$), and of (T_2 given $-H_2$) are each 1/3.

experimenters at least once, it does not guarantee that she is awakened by the experimenters *today*. Today might be Tuesday, after all.)

Next she asks about the *current* probabilities of H_1, H_2, T_1, and T_2, as determined by her total current evidence. Again exercising her impeccable rationality, she updates the prior probabilities in the light of the fact that she has just been awakened by the experimenters: the probability of H_2 goes to zero, and since the remaining three statements are exclusive and (given her current evidence) exhaustive, their current probabilities are ascertained by multiplying their respective prior probabilities by the common normalization-factor of 4/3.[7] So, after updating, $P(H_1) = P(T_1) = P(T_2) = 1/3$. But P(HEADS) = $P(H_1)$, and P(TAILS) = $P(T_1 \text{ v } T_2)$. Hence, P(HEADS) = 1/3 and P(TAILS) = 2/3.

So Sleeping Beauty's evidence is indeed relevant to HEADS v. TAILS, and it does indeed drive their probabilities to 1/3 and 2/3 respectively even though they each had probability 1/2 on the preceding Sunday. This is because the current evidence excludes the possibility expressed by the statement H_2 – a statement which, when Sleeping Beauty is awakened, has a prior probability of 1/4 as determined by her Sunday evidence.

3. I will offer an abstract characterization of the key factors at work in the Sleeping Beauty problem. Before doing so, several preliminary points need to be made. First, the kind of probability under discussion is *epistemic*, in the sense that it is essentially tied to available evidence. Epistemic probability perhaps should be construed as *degree of evidential support*, or instead perhaps should be construed as 'credence': *degree of belief* or *rational degree of belief*.

Second, epistemic-probability contexts are *intensional*: when a statement S* is obtained from a statement S by substitution of co-referring singular terms, the probability of the resulting statement S* can differ from the probability of the original statement S.[8] (As one might say, epistemic-probability contexts lack the feature of substitutivity *salva probabilitate*.[9]) The Sleeping Beauty problem illustrates this phenomenon. At any time during the experiment, the following claim is true, for Sleeping Beauty:

P(HEADS and Tuesday I am awakened by the experimenters) = 0.

[7] This is the standard, Bayesian, form of updating: conditionalization. See n. 6.

[8] See Horgan (2000) for further discussion of this insufficiently appreciated feature of epistemic probability, with specific application to the two-envelope paradox. Objective probability, on the other hand – i.e. *chance* – need not be intensional, even though epistemic probability is.

[9] Thanks to Sarah Wright for suggesting this phrase to me.

Suppose, however, that she has just been awakened by the experimenters and that (unbeknown to her),

Tuesday = today.

The following claim, obtained by substitution of co-referring terms, is false:

P(HEADS and today I am awakened by the experimenters) = 0.

Third, the syntax and semantics of epistemic-probability ascriptions needs to be construed in some way that accommodates this intensionality. For present purposes, it will be convenient to continue the practice already adopted in §2 above: I will take *statements* to be the items to which probabilities are ascribed.[10]

With these three points as background, let me now describe abstractly the various interconnected elements crucially at work in the Sleeping Beauty problem. First, something occurs that results in a *loss of self-location information* for the epistemic agent – in the case of Sleeping Beauty, a loss of information about what day it is.[11] Second, a *potential cognitive mishap* is involved in this loss of self-location information, a mishap that might have occurred (for all the agent can tell) but need not have – in the case of Sleeping Beauty, being injected with a drug that erases all memories of the preceding 24 hours.

Third, the loss of self-location information generates, for the agent, a way of conceiving and describing one's self-location – in the case of Sleeping Beauty, 'today' – that is *essentially indexical*, in the sense that the agent does not know what location it is that is thus conceived (and thus expressed) indexically. Fourth, the loss of self-location information thereby

[10] I leave it open exactly what sorts of entities statements are, qua bearers of epistemic possibility. (I do assume that there are *indexical* statements.) One might try accommodating the intensionality of probability in other ways than by taking the bearers of epistemic probability to be statements. For instance, perhaps epistemic probability attaches to very finely-individuated *possibilities* (including indexical ones) – so that the possibility *being awakened on Tuesday* is distinct from the possibility *being awakened today*, even if today = Tuesday. (Possibilities like *being awakened today* might get represented as classes of so-called *centered worlds*: possible worlds with designated individuals-at-times within them.) Or perhaps epistemic probability is an attribute of possibilities under *descriptions* and/or under *modes of presentation* (including indexical ones). Advocates of some such views would need to find ways to reformulate, within their preferred idiom, the claims I make in the text about probabilities of indexical statements.

[11] Loss of self-location information is also emphasized in relation to the Sleeping Beauty problem by Monton (2002) and by Arntzenius (forthcoming), both of whom endorse the claim that after Sleeping Beauty is awakened, P(HEADS) = 1/3. For Arntzenius this is a change from the position taken in Arntzenius 2002; see my n. 3 above.

generates a partition of essentially indexical *statements*, each of which is consistent with the evidence possessed by the agent prior to the loss of self-location information – in the case of Sleeping Beauty, a partition comprising the statements H_1, H_2, T_1 and T_2.

Fifth, this partition is *irreducibly indexical*, in this sense: the agent lacks a non-indexical way of conceiving (or describing) her own self-location that can be substituted *salva probabilitate* for the relevant, essentially indexical, thought-constituent (term) within each of the agent's probability judgments (ascriptions) concerning the statements in the partition. (Thus, the irreducible indexicality of the partition reflects the intensionality of epistemic probability.) In the case of Sleeping Beauty, for instance, the irreducible indexicality of the statement-partition $\{H_1, H_2, T_1, T_2\}$ is illustrated by the following substitutivity-failure involving statement T_1: even if today happens to be Monday (unbeknown to her), and even though $P(T_1)$ = P(TAILS and today is Monday) = 1/3, it is not the case that P(TAILS and Monday is Monday) = 1/3. Rather, P(TAILS and Monday is Monday) = P(TAILS) = 2/3.

Sixth, the irreducibly indexical partition of essentially indexical statements only arises once the loss of self-location information occurs, because this information-loss is what generates such a partition in the first place. This in turn means, seventh, that the prior probabilities of these essentially indexical statements only arise once the information-loss occurs. What makes the relevant probabilities count as *prior* probabilities – here and in general (as emphasized in §2) – is that these are the probabilities *now* possessed, relative to the evidence *then* available, by the statements in the partition. An atypical feature of the situation, however, is that these present prior-probabilities are not probabilities that *obtained* prior to the information-loss; for, the irreducibly indexical statement-partition had not yet arisen.

Eighth, the very episode that generates the information-loss also furnishes the agent with conclusive evidence, concerning a specific one of the essentially indexical statements in the partition (or perhaps a specific disjunction of them), that it does not obtain. Thus, ninth, the information-loss and the acquisition of this new evidence occur *simultaneously*, with the new evidence pertaining specifically to the irreducibly indexical statement-partition generated by the information-loss. So tenth, the current probability of each of the remaining statements in the partition, for the agent, is the prior *conditional* probability of that statement given this evidence – where this prior conditional probability arises, via the information-loss, simultaneously with the current probability itself.

The soritical argument by analogy in Dorr (2002), mentioned in §1 above, is relevant to the features described in the preceding two paragraphs. In my terms, Dorr's argument can be understood as follows. He first describes a case in which the information-gain occurs some time later than the information-loss that generates the essentially indexical statement-partition; then he describes a sorites sequence of similar cases in which the time-interval becomes progressively shorter; the limit case in this sequence is the original Sleeping Beauty scenario. In effect, the argument is (1) that for each successive case prior to the limit case, information is gained (after the information-loss) that lowers the probability of HEADS to 1/3 and raises the probability of TAILS to 2/3, and (2) that the limit case itself is relevantly similar in this respect to the other cases, even though in the limit case there is no time-gap between the information-loss itself and the information-gain that occurs against the background of this information-loss. This line of reasoning strikes me as correct, rather than as a slippery-slope fallacy; the point is that there is no evident reason why the limit case should be treated differently from the others.[12]

The ten features lately described constitute an abstract recipe for constructing Sleeping-Beauty-type cases. One such case is a version of the Sleeping Beauty story in which the coin is tossed on Sunday evening, before Sleeping Beauty is put to sleep. Another such case is the following variant story in which sleeping plays no role (and in which the relevant temporal indexical does not necessarily refer to a different day from the day on which the agent still knows what day it is):

> On Wednesday you learn for certain that you are to be the subject of an experiment. At 9:00 p.m. on Wednesday the experimenters will inject you with a drug that will cause 24 hours of insomnia. On Thursday evening they will toss a fair coin. If the outcome is Tails, then at 9:00 p.m. on Thursday they will inject you with a drug that obliterates all memories of the previous 24 hours and all effects of the insomnia – so that the moment after the injection, you will be unable to tell that it is not a moment after 9:00 p.m. on Wednesday.

Initially the probability of HEADS is 1/2. But shortly after you find yourself having just been injected, with no memory of any prior injection or any intervening period of insomnia, what is the probability of HEADS? The

[12] Bradley (2003) argues that Dorr's initial case is crucially disanalogous to the original Sleeping Beauty scenario. (He does not discuss Dorr's sorites sequence.) He writes, 'In the variant case, a certain possibility has been eliminated. It *could* have turned out that it was Heads and Tuesday. … In the original case, there is no such possibility' (267–68). On the contrary: in the original case too, the possibility *HEADS and today is Tuesday* is consistent with Sleeping Beauty's Sunday evidence, and thus now has prior probability 1/4 as determined by her Sunday evidence. Here too this possibility has been eliminated by her current total evidence – specifically, by her newly acquired, essentially indexical, knowledge that she has been awakened *today* by the experimenters (with no memory of any prior awakening).

discussion in §2 applies here, *mutatis mutandis*, because this case too exemplifies the ten features described in the present section. So now the probability of HEADS is 1/3.

This scenario nicely illustrates an important fact: sometimes a statement's present prior-probability is not identical to its preceding current-probability. Suppose that (unbeknown to you) today is still Wednesday, as you find yourself having just been injected and with no memory of a prior injection or of 24 hours of insomnia. You consider the four statements of the form 'X and today is Y' obtained by substituting 'HEADS' or 'TAILS' for 'X' and substituting 'Wednesday' or 'Thursday' for 'Y'. The present prior-probability of the statement *HEADS and today is Wednesday*, as determined by your pre-injection evidence, is 1/4. But the preceding current-probability possessed by this statement before the injection, as determined by the then-current evidence, was 1/2.[13] This difference between the preceding current-probability and the present prior-probability reflects the intervening loss of self-location information.

The discussion in §2 also applies, *mutatis mutandis*, to a variant of the original Sleeping Beauty story in which Sleeping Beauty is informed on Sunday that if the coin-toss comes up heads on Monday, then she will be killed immediately. When Sleeping Beauty is awakened by the experimenters and she contemplates the statement *HEADS and today is Tuesday and I am now dead*, she rightly judges that this statement has prior probability 1/4, as determined by her Sunday evidence. Given her *current* evidence, of course, she updates the probability to zero.[14]

Our abstract recipe probably can be further generalized. For instance, the second feature – the presence of a potential cognitive mishap – presumably is replaceable by various other ways of losing self-location

[13] There is a parallel difference between presently-prior and previously-current *conditional* probabilities. Both conform to the standard definition of conditional probability: P(A given B) = P(A&B)/P(B). The presently-prior conditional probability of (A given B) is determined, in accordance with the definition, by the present prior-probabilities of A and B, whereas the previously-current conditional probability of (A given B) is determined by the preceding current-probabilities of A and B. For instance, after you find yourself having just been injected and no longer knowing whether today is Wednesday or Thursday, the presently-prior conditional probability of [(HEADS and today is Wednesday) *given that* it's not the case that (HEADS and today is Thursday)] is 1/3. But the previously-current conditional probability of this statement, back before the experiment began and when you still knew that it was Wednesday, was 1/2. In situations where previously-current conditional probabilities differ from presently-prior conditional probabilities, Bayesian updating should employ the latter rather than the former.

[14] Justin Fisher impressed upon me the need to provide a treatment of the Sleeping Beauty problem that would generalize to cases like those described in recent paragraphs.

information. Also, a statement-partition that is irreducibly indexical presumably can arise through the loss of information not only about self-*location* but also about self-*identity*. So my recommended treatment of the Sleeping Beauty problem probably is applicable to cases involving loss of self-identity information, and/or to cases in which the loss of self-location or self-identity information results from some factor other than potential cognitive mishap.[15]

4. Before being put to sleep on Sunday, Sleeping Beauty rightly judged that the probability of HEADS was 1/2. Upon being awakened on Monday, she rightly judges that P(HEADS) = 1/3. Elga claims that this belief change does not result from acquiring new information. He says:

> This belief change is unusual. It is not the result of your receiving new information – you were already certain that you would be awakened on Monday. … Neither is this belief change the result of your suffering any cognitive mishaps in the intervening time – recall that the forgetting drug isn't administered until well after you are first awakened. (145)

He goes on to urge the following moral:

> Thus the Sleeping Beauty Problem provides a new variety of counterexample to Bas van Fraassen's 'Reflection Principle' (1984: 244, 1995: 19), even an extremely qualified version of which entails the following: Any agent who is certain that she will tomorrow have credence x in proposition R (though she will neither receive new information nor suffer any cognitive mishaps in the intervening time) ought *now* to have credence x in R. (146)

But in the sense of 'new information' that is contextually most appropriate with respect to questions of how available evidence affects probabilities, Sleeping Beauty *does* acquire new information upon being awakened on Monday. This awakening-event produces, simultaneously, both an information loss and also an information gain that is predicated upon that information loss. When she is awakened by the experimenters on Monday, she thereby loses self-location information: the awakening-event generates the irreducibly indexical partition of statements H_1, H_2, T_1, and T_2 – each of which now expresses an epistemic possibility relative to her Sunday evidence, and each of which now has a prior probability of 1/4 as determined by her Sunday evidence. Against the backdrop of this loss of self-location information, the awakening-event simultaneously constitutes evidence that conclusively excludes the epistemic possibility expressed by H_2; although

[15] For discussion of a range of cases that exhibit these features and appear otherwise similar to the Sleeping Beauty problem, see Arntzenius, forthcoming and Elga, manuscript.

H_2 is consistent with her Sunday evidence, it is ruled out by her total *current* evidence. Exclusion of epistemic possibilities counts as acquisition of new information, in the context of ascertaining probabilities on the basis of current evidence. So it appears that the Sleeping Beauty problem does not really constitute a counter-example to the core principle that Elga cites – the principle that is entailed by even an extremely qualified version of the Reflection Principle – provided that the phrase 'new information' in this core principle is appropriately construed.

Elga uses the phrase 'new information' in a more coarse-grained way. He says, 'To say that an agent receives new information (as I shall use that expression) is to say that the agent receives evidence that rules out possible worlds not already ruled out by her previous evidence' (2000: 145, n. 4). On this usage, evidence that only rules out a *within*-world possibility concerning one's present temporal location (e.g. Tuesday in a HEADS-world), but does not rule out any possible world altogether (e.g. a HEADS-world or a TAILS-world), does not count as new information. So much the worse for Elga's usage, in the present context.

It bears emphasis that mere changes in the referents of self-locating indexicals should not be construed as generating new information all by themselves, on pain of trivializing the Reflection Principle.[16] If several days go by, and one knows all the while what day one is currently located within, then the mere change in referent of the term 'today' does not constitute or generate 'new information' about one's self-location. (Perhaps one knows all along all relevant information concerning each successive day, expressible non-indexically. If so, then one also knows all along how to reformulate various aspects of that information indexically – where such reformulations do not alter probabilities.) By contrast, the statement-partition that Sleeping Beauty contemplates after being awakened is *irreducibly indexical*: because 'today' is essentially indexical for her, some statements in the partition have probabilities (both prior probabilities and current probabilities) different from the probabilities of the corresponding non-indexical statements that result from replacing essential indexicals by co-referential non-indexicals. Exclusion of an epistemic possibility expressed by a statement in an irreducibly indexical statement-partition *does* constitute genuinely new information.

One might reply, 'But Sleeping Beauty already knew on Sunday that she would be awakened on Monday with no memory of a prior awakening. So didn't she already possess, on Sunday, all the information that she would possess upon being awakened on Monday?' The answer is no. Although Sleeping Beauty did know on Sunday that she would be awakened on Monday, and although she also knew on Sunday that on Monday

[16] Justin Fisher impressed upon me the importance of this point.

she would possess information that would be expressible by saying or thinking

It is not the case that (today is Tuesday and the coin comes up heads),

the information she thus expresses on Monday is *essentially indexical*. She did not yet possess this information on Sunday. Although she already knew on Sunday that she would obtain essentially indexical new information on Monday, and although she even knew on Sunday how she we would indexically describe this new information on Monday, she did not yet have the new information itself.[17,18]

University of Arizona
Tucson, AZ 85721-0027, USA
thorgan@email.arizona.edu

References

Arntzenius, F. 2002. Reflections on Sleeping Beauty. *Analysis* 62: 53–62.

Arntzenius, F. Forthcoming. Some problems for conditionalization and reflection.

Bradley, D. 2003. Sleeping Beauty: a note on Dorr's argument for 1/3. *Analysis* 63: 266–68.

Dorr, C. 2002. Sleeping Beauty: in defence of Elga. *Analysis* 62: 292–96.

Elga, A. 2000. Self-locating belief and the Sleeping Beauty problem. *Analysis* 60: 143–47.

Elga, A. Manuscript. Defeating Dr. Evil with self-locating belief.

Horgan, T. 2000. The two-envelope paradox, nonstandard expected utility, and the intensionality of probability. *Noûs* 34: 578–602.

Lewis, D. 2001. Sleeping Beauty: reply to Elga. *Analysis* 61: 171–76.

[17] *Note added in proof*: Since this paper was accepted I have come to realize that for both the original scenario and its variants, the initially-assigned probabilities of 1/4 each for H_1, H_2, T_1, and T_2 obtain relative to a portion of the agent's total current information comprising not only her *prior* information about the experiment but also her *current*, disjunctive, information about what the current day is (viz. Monday or Tuesday in the original problem, Wednesday or Thursday in this one). So I now think that these *preliminary* probabilities (as I now call them) should not be called *prior* probabilities. I would now put my key claim this way: the current probabilities of H_1, H_2, T_1, and T_2 are their preliminary *conditional* probabilities with respect to the statement 'I am awakened today by the experimenters.' Such Bayesian updating is a generalization of standard Bayesian updating of prior probabilities, since ordinary prior probabilities are a special case of preliminary probabilities.

[18] I have pestered numerous people about the Sleeping Beauty problem. Thanks to Justin Fisher and Sarah Wright for especially valuable discussion and feedback, and to Robert Barnard, David Chalmers, Ned Hall, John Hawthorne, Dianne Horgan, Kelly Horgan, Jenann Ismael, Keith Lehrer, David Papineau, John Pollock, Eric Schwitzgebel, John Tienson, Mark Timmons, Michael Tye, Brian Weatherson and Ruth Weintraub.

Monton, B. 2002. Sleeping Beauty and the forgetful Bayesian. *Analysis* 62: 47–53.
van Fraassen, B. C. 1984. Belief and the will. *Journal of Philosophy* 81: 235–56.
van Fraassen, B. C. 1995. Belief and problem of Ulysses and the sirens. *Philosophical Studies* 77: 7–37.

ANALYSIS 64.1, January 2004, pp. 21–29.