Erkenntnis, in press

Generalized Conditionalization and the Sleeping Beauty Problem

Terry Horgan

University of Arizona

Anna Mahtani

Oxford University

The Sleeping Beauty problem, much discussed in the recent philosophical literature, goes as follows. Sleeping Beauty is an experimental subject in a sleep-research laboratory. On Sunday she is reliably given the following information, all of which she correctly believes to be true. On Sunday evening the researchers will put her into dreamless sleep in the laboratory. They will flip a fair coin. On Monday they will awaken her in the laboratory, and she will spend the day there. If the coin came up heads, then she will continue sleeping dreamlessly until Wednesday morning, when she will wake up by herself. But if the coin came up tails, then they will erase all her Monday memories while she is sleeping on Monday night; then on Tuesday they will again a waken her in the laboratory, and she will spend all day there, and they will again put her into dreamless sleep on Tuesday evening, and she will sleep dreamlessly from Tuesday evening until Wednesday morning, when she will wake up by herself.

The next thing she knows, she finds herself being awakened by the experimenters in the laboratory, with no memory of a prior awakening—something that she knew would happen at least once and maybe twice. She knows that it is now Monday or Tuesday, but she does not know which. The question is put to her, "What is the probability that the coin-flip comes up heads?" This question, for the scenario as just described, is the Sleeping Beauty problem.[1]

Let P_(HEADS) be Beauty's epistemic probability on Sunday for the proposition that the coin-flip comes up heads, and let P(HEADS) be her epistemic probability for this proposition upon being awakened by the experimenters.[2] Some (e.g., Lewis 2001) claim that P(HEADS) = ½, while others (e.g., Elga 2000) claim that P(HEADS) = 1/3. In this paper we present a new argument for the claim that P(HEADS)=1/3. Our argument will depend on a principle that we will call 'generalized conditionalization'. We begin by explaining the respects in which generalized conditionalization is a generalization of the familiar, conventional kind of conditionalization, and we will set forth a rationale for the legitimacy of generalized conditionalization. We will then lay out the new argument for thirdism. Thereafter we will respond to an initially plausible-looking objection to the argument, and we will discuss one of our argument's advantages: it is immune to two attacks that Joel Pust has leveled at other arguments for thirdism (Pust 2008 and Pust forthcoming).

1. Generalized Conditionalization

In this section we describe and defend the principle of generalized conditionalization, which is required for our new argument for thirdism.

In conventional conditionalization, one's reasoning deploys a space of candidate-possibilities that we will call *diachronic*, because these candidate-possibilities persist from the moment that the epistemic agent first acquires pertinent information that bears upon the epistemic probabilities of these candidate-possibilities.[3] The candidate-possibilities are expressed by a set of propositions $A_1, \ldots, A_n$. One begins with certain *prior probabilities* $P\_(A_1), P\_(A_2), \ldots, P\_(A_n)$ for propositions $A_1, \ldots, A_n$ —probabilities that accrue to these propositions relative to the body of information that previously was one's total pertinent information. (The horizontal line in the notation '$P\_(\quad)$' indicates that these are prior probabilities.) One also invokes prior *conditional* probabilities of the form $P\_(A_1|B), P\_(A_2|B), \ldots, P\_(A_n|B)$, where B is a proposition that was not known to be true at that prior time, but which one subsequently has come to know is true (without one's having subsequently acquired any other pertinent new information). On the basis of the newly acquired information B, one now updates one's prior unconditional probabilities for the propositions $A_1, \ldots, A_n$, by setting their current unconditional probabilities equal to their prior probabilities conditional on B: i.e., for each i, $1 \leq i \leq n$, $P(A_i) = P\_(A_i|B)$. (The absence of the horizontal line in the notation '$P(\quad)$' indicates that these are updated probabilities.)

The method we are calling *generalized* conditionalization goes as follows. In a situation where one's total available information includes the information that some proposition B obtains, one begins by "bracketing" that information and then assigning *preliminary* probabilities relative to a specific proper portion of one's total, current, pertinent information—intuitively speaking, a proper portion consisting of all that information *minus* the information that B obtains.[4] One assigns preliminary probabilities $P\_(A_1)$,

$P\_(A_2)$, …, $P\_(A_n)$ to various propositions A_1, …, A_n — probabilities that accrue to these propositions relative to the information one is taking into consideration. (Now the horizontal line in the notation 'P_( )' indicates that these are preliminary probabilities.) One also invokes preliminary conditional probabilities of the form $P\_(A_1|B)$, $P\_(A_2|B)$, …, $P\_(A_n|B)$, where B is the proposition being bracketed. One now factors the information B back in via conditionalization—i.e., one updates one's preliminary probabilities for the propositions A_1, …, A_n, by setting their updated probabilities equal to their preliminary probabilities conditional on B: $P(A_i) = P\_(A_i|B)$. (Once again, the absence of the horizontal line in the notation 'P( )' indicates that these are updated probabilities.) So the principle of generalized conditionalization is this: one's epistemic probability $P(A_i)$ relative to all the information one possesses should be equal to one's preliminary epistemic conditional probability $P\_(A_i|B)$ relative to some proper portion of one's total information that excludes proposition B.

There are two respects in which generalized conditionalization constitutes a *generalization* of conventional conditionalization. First, the "bracketed" information expressed by the proposition B need not be *newly acquired* information; rather, it can be virtually any item of information in the agent's total corpus of current pertinent information.[5] (Cases where bracketed information was previously possessed we will call "conditionalization on old information.") Second, the items in the pertinent space of candidate-possibilities need not be *diachronic* candidate-possibilities, expressible non-indexically and present from the moment that the epistemic agent first acquires pertinent information that bears upon the epistemic probabilities of these candidate-possibilities. Rather, they can be *synchronic* candidate-possibilities, as we will call them—e.g. essentially temporally-indexical possibilities that did not yet exist (because of their temporal indexicality) at the time when one first acquired some of one's currently pertinent information. (Cases involving synchronic candidate-possibilities we will call "synchronic conditionalization.") Our new argument for thirdism applies generalized conditionalization in a way that exhibits both of the two features mentioned.

The following simple and powerful rationale can be given in support of the method of generalized conditionalization. Let *I* be the total corpus of information, available to a given epistemic agent at a given time, that is pertinent to the epistemic probabilities of the propositions A_1, …, A_n, and let *I*-B be a proper portion of *I* that consists of *I* itself "minus" the information expressed by the bracketed proposition B.[6] If A_1, …, A_n express diachronic candidate-possibilities and B is newly acquired information that excludes some of these candidate-possibilities (and is one's strongest pertinent newly-acquired information), then the obvious and natural way to update the prior probabilities of the propositions $A_1, \ldots, A_n$ is via conditionalization on the proposition B. (This is conventional conditionalization.) But, from a purely logical point of view, it makes no difference whether the proposition B constitutes new information that previously was not a part of one's total pertinent information, or instead was already known to obtain even before one commenced one's probabilistic reasoning concerning propositions $A_1, \ldots, A_n$. Likewise, it makes no difference whether the items in the space of candidate-possibilities—the items to which one is assigning preliminary and updated probabilities—are diachronic candidate-possibilities or synchronic candidate-possibilities. For, from a purely logical point of view, the preliminary absolute probabilities of $A_1, \ldots, A_n$ (relative to *I*-B), the preliminary conditional probabilities of $A_1, \ldots, A_n$ conditional on B (relative to *I*-B), and the updated absolute probabilities of $A_1, \ldots, A_n$ (relative to *I*) all depend only on the *content* of the information *I*-B and the information *I*—not on when various portions of the information were acquired, and not on whether the space of candidate-possibilities to which $A_1, \ldots, A_n$ belong happens to be diachronic or synchronic.

The upshot, so far, is this. Generalized conditionalization, when applicable, appears to be a quite legitimate way to ascertain epistemic probabilities.

2. The New Argument for Thirdism

We remarked in section 1 that for purposes of assigning preliminary epistemic probabilities to a set of propositions in order to do generalized conditionalization, one may bracket virtually any piece of information in one's total corpus of current pertinent information. This means, for instance, that the bracketed information need not be newly acquired information, as opposed to old information. And one can reason probabilistically about essentially indexical synchronic possibilities, bracketing old information that it is not essentially indexical: we call this *synchronic conditionalization on old information*.

This latter observation is the key to the way of reasoning about the sleeping Beauty problem that we will now propose. The reasoning goes as follows. Sleeping Beauty is awakened on Monday, and considers a body of information *I** that is a proper portion of her total corpus of current pertinent information. *I** includes the essentially indexical information that she herself was awakened *today* by the experimenters, that the coin is flipped on Sunday night, and the essentially indexical information that today is either Monday or Tuesday; *I** also includes most, but not all, of the pertinent diachronic information that she already possessed on Sunday. *I* excludes* the information embodied in the following rule governing the sleep experiment, all of which is diachronic information that she already possessed on Sunday:

(R_1) On Monday morning Beauty is awakened by the experimenters, and she spends the day awake in the lab. If the coin came up heads, then she sleeps dreamlessly from Monday evening until waking up by herself on Wednesday morning. If it came up tails, then her Monday memories are erased while she is asleep on Monday evening, she is awakened again by the experimenters on Tuesday, and she again spends the day awake in the lab.

But although *I** excludes the information constituting rule R_1, it *includes* a disjunctive fact that is a logical consequence of R_1 together with the other included information. This disjunctive fact is that the experimenters behave in accordance with one of four (mutually exclusive) rules—either R_1 or else one of the following three. (Note that the order in which the components of a given rule are stated does not fully correspond to the temporal sequence in which the components would get implemented; this is to facilitate maximum parallelism in the way the respective rules are formulated.)

(R_2) On Monday morning Beauty is awakened by the experimenters, and she spends the day awake in the lab. If the coin came up heads on Sunday evening, then her Monday memories are erased while she is asleep on Monday evening, she is awakened again by the experimenters on Tuesday, and she again spends the day awake in the lab. If it came up tails, then she sleeps dreamlessly from Monday evening until waking up by herself on Wednesday morning.

(R_3) On Tuesday morning Beauty is awakened in the lab the experimenters, and she spends the day awake in the lab. If the coin came up heads on Sunday evening, then she sleeps dreamlessly from Sunday evening until she is awakened by the experimenters on Tuesday. If it came up tails, then she is awakened by the experimenters on Monday morning, she spends the day awake in the lab, and her Monday memories are erased while she is asleep on Monday evening.

(R_4) On Tuesday morning Beauty is awakened by the experimenters, and she spends Tuesday awake in the lab. If the coin came up heads on Sunday evening, then she is awakened by the experimenters on Monday morning, she spends Monday awake in the lab, and her Monday memories will be erased while she is asleep on Monday evening. If it came up tails, then she sleeps dreamlessly from Sunday evening until she is awakened by the experimenters on Tuesday.

Here, the proposition R_1 constitutes the information that is part of her total current pertinent information *I* but is being bracketed, and thus is excluded from *I**. R_1 is being bracketed in a way that leaves behind, in *I**, a logically weaker "residue" of information that she actually possesses by virtue of actually possessing R_1 itself—viz., the disjunctive information (R_1 or R_2 or R_3 or R_4).[7]

Beauty, upon being awakened in the lab by the experimenters, now sets herself the task of assigning preliminary epistemic probabilities, relative to the body of information *I**, to various candidate-possibilities. These are the unconditional epistemic probabilities that she *would* assign if *I** were her total available relevant information. She considers the following table, which characterizes in a certain way the space of pertinent candidate-possibilities:

1. HEADS
   - 1.a. $\text{Today}_{\text{Mon}}$
     - 1.a.i. R_1
     - 1.a.ii. R_2
     - 1.a.iii. R_4
   - 1.b. $\text{Today}_{\text{Tues}}$
     - 1.b.i. R_2
     - 1.b.ii. R_3
     - 1.b.iii. R_4
2. TAILS
   - 2.a. $\text{Today}_{\text{Mon}}$
     - 2.a.i. R_1
     - 2.a.ii. R_2
     - 2.a.iii. R_3
   - 2.b. $\text{Today}_{\text{Tues}}$
     - 2.b.i. R_1
     - 2.b.ii. R_3
     - 2.b.iii. R_4

Table 1

The information expressible by Beauty as "I was awakened today by the experimenters" plays a crucial role in structuring the space of candidate-possibilities indicated by Table 1. This essentially indexical information excludes, for each of the cases 1.a, 1.b, 2.a, 2.b, one of the four rules from being operative in the given case—a different rule being excluded in each case. Table 1 also makes explicit the fact that the candidate-possibilities exhibit what we will call a *hierarchicial partition-structure*: they subdivide into a coarse-grained partition {HEADS, TAILS} concerning the outcome of the coin-flip, with each cell of this partition subdividing into a finer grained sub-partition {$\text{Today}_{\text{Mon}}$, $\text{Today}_{\text{Tues}}$} concerning which day it is today, with each of these latter cells further subdividing into yet-finer grained sub-sub-partition consisting of three candidate-possibilities concerning which rule is being followed. This partition structure is *strongly symmetrical,* in this sense: the two cells in the outer partition {HEADS, TAILS} subdivide symmetrically into a four-way partition of matching pairs of sub-cells {1.a, 1.b, 2.a, 2.b}, and these sub-cells then subdivide symmetrically into a 12-way partition comprising four structurally parallel sets of sub-sub-cells—with each of the four structurally parallel sets comprising three candidate-rules and excluding one rule (a different rule in each case).

In light of the fact that the space of candidate-possibilities possesses the strongly symmetrical hierarchical partition-structure exhibited in Table 1, Beauty concludes that the evidence provided by information *I** is indifferent with respect to the two outer cells HEADS and TAILS and the four intermediate cells 1.a, 1.b, 2.a, and 2.b. She therefore assigns preliminary probabilities of ½ each to HEADS and TAILS, and preliminary probabilities of ¼ each to 1.a, 1.b, 2.a, and 2.b.

What about the four structurally parallel sets of three sub-sub-cells, each set consisting of three possibilities concerning which rule is being used? Beauty rightly notes that each of these four sets exhibit a certain kind of asymmetry: two of the three rules dictate that the experimenters will awaken her on another day (in addition to doing so today), whereas the remaining rule dictates that they awaken her only today. (For example, in case 1.a—the case (HEADS & $\text{Today}_{\text{Mon}}$)—the rules R_2 and R_4 would each require the experimenters to awaken her again on Tuesday, whereas the rule R_1 would require them not to.) But this kind of asymmetry is not a reason to consider any one of the three rules more likely, or less likely, to be operative (within the given case) than the others. Since no such reason is present, she therefore concludes, regarding each of the four structurally parallel sets of three sub-sub-cells in Table 1, that the three candidate-possibilities within the set are equally likely, and hence each have preliminary epistemic probability 1/12.[8] Thus, the preliminary probabilities that accrue to the candidate-possibilities in Table 1 are as follows:

1.	HEADS	1/2	
1.a.	$Today_{Mon}$	1/4	
1.a.i.	R_1		1/12
1.a.ii.	R_2		1/12
1.a.iii.	R_4		1/12
1.b.	$Today_{Tues}$	1/4	
1.b.i.	R_2		1/12
1.b.ii.	R_3		1/12
1.b.iii.	R_4		1/12
2.	TAILS	1/2	
2.a.	$Today_{Mon}$	1/4	
2.a.i.	R_1		1/12
2.a.ii.	R_2		1/12
2.a.iii.	R_3		1/12
2.b.	$Today_{Tues}$	1/4	
2.b.i.	R_1		1/12
2.b.ii.	R_3		1/12
2.b.iii.	R_4		1/12

Table 2

From the preliminary probabilities in Table 2 it follows that Beauty also must assign the following preliminary conditional probabilities, again envisioning being in the epistemic situation lately described[9]:

$P\_[(Today_{H,Mon} \& R_1)|R_1] = 1/3$

$P\_[(Today_{H,Mon} \& R_2)|R_1] = 0$

$P\_[(Today_{H,Mon} \& R_4)|R_1] = 0$

$P\_[(Today_{H,Tues} \& R_2)|R_1] = 0$

$P\_[(Today_{H,Tues} \& R_3)|R_1] = 0$

$P\_[(Today_{H,Tues} \& R_4)|R_1] = 0$

$P\_[(Today_{T,Mon} \& R_1)|R_1] = 1/3$

$P\_[(Today_{T,Mon} \& R_2)|R_1] = 0$

$P\_[(Today_{T,Mon} \& R_3)|R_1] = 0$

$P\_[(Today_{T,Tues} \& R_1)|R_1] = 1/3$

$P\_[(Today_{T,Tues} \& R_3)|R_1] = 0$

$P\_[(Today_{T,Tues} \& R_4)|R_1] = 0$

Table 3

She now resorts to a version of generalized conditionalization that we above called synchronic conditionalization on old information. It is *synchronic* in two respects: first, because the information *I** includes the essentially indexical information that Sleeping Beauty can express by saying "I was awakened today by the experimenters" and "Today is either Monday or Tuesday," and second, because the conditionalization pertains to a space of synchronic, essentially indexical, possibilities—possibilities that have only arisen in the wake of her recent (partial) loss of self-locational information. The bracketed information deployed in this form of conditionalization is *old* information, because claim R_1 is temporally non-indexical and has been possessed all along since she was initially informed about the experimental setup back on Sunday. She applies this form of conditionalization to the preliminary conditional probabilities in Table 2 together with her knowledge of the bracketed proposition R_1 to obtain these updated absolute probabilities:

$P(Today_{H,Mon} \& R_1) = P(Today_{H,Mon}) = 1/3$

$P(Today_{T,Mon} \& R_1) = P(Today_{T,Mon}) = 1/3$

$P(Today_{T,Tues} \& R_1) = P(Today_{T,Tues}) = 1/3$[10]

And so she rightly reaches the thirder conclusion:

$P(H) = P(Today_{H,Mon}) = 1/3$

$P(T) = [P(Today_{T,Mon}) + P(Today_{T,Tues})] = 2/3$[11]

A further observation deserves emphasis. We stressed at the end of section 1 that it makes no difference, logically speaking, whether the bracketed information that one deploys in doing conditionalization is “old” information or instead is recently-acquired information. That fact is clearly pertinent here. If Beauty were in an alternative sleep-experiment situation in which her Sunday information included the disjunction (R_1 of R_2 or R_3 or R_4) but did not include proposition R_1, then upon being awakened by the experimenters it would clearly be appropriate for her to assign epistemic probabilities of 1/12 each to possibilities *HEADS and today is Monday and* R_1, *HEADS and today is Monday and* R_2, etc. And if she were to later acquire knowledge of proposition R_1, then it would clearly be appropriate for her to do synchronic conditionalization on this new non-indexical information, thereby correctly concluding—by reasoning that exactly parallels the reasoning set forth in the immediately preceding paragraph—that $P(H) = 1/3$ and $P(T) = 2/3$. So, since it makes no logical difference whether the information embodied in R_1 is old information or newly acquired information, this same reasoning correctly yields the conclusion that in the actual Sleeping Beauty problem, $P(H) = 1/3$ and $P(T) = 2/3$.

3. An objection and a reply

Halfers tend to think that suitable application of the Principal Principle (cf. note 8) and the Lewis-Elga indifference principle (cf. note 8) yields the conclusion $P(H) = P(T) = ½$. One such line of reasoning, structured so as to first generate preliminary probabilities relative to the information I* and then deploy synchronic conditionalization on old non-indexical information, would go as follows.[12] By the Principal Principle,

$P\_(H \& R_i) = P\_(T \& R_i)$ for each i, $1 \leq i \leq 4$.

This is because (the halfer would claim) the two uncentered possibilities (H & R_i) and (T & R_i) differ only in those respects determined by the outcome of a fair coin toss, with the information in *I** providing no further relevant evidence as to their comparative likelihood. So, any distribution of preliminary probabilities over the centered possibilities that also conforms to the Lewis-Elga indifference principle will lead, via synchronic conditionalization over old information, to the conclusion concerning non-preliminary probabilities that the halfers advocate: viz., $P(H) = P(T) = ½$. To illustrate this, suppose that Beauty assigns (on the basis of evidential symmetry, say) preliminary probabilities of ¼ each to the propositions R_1, R_2, R_3, and R_4. Applying the Principal Principal to each of these four equi-likely cases, she next obtains the following distribution of preliminary probabilities over the pertinent uncentered possibilities:

$P\_(H \& R_1) = 1/8$

$P\_(T \& R_1) = 1/8$

$P\_(H \& R_2) = 1/8$

$P\_(T \& R_2) = 1/8$

$P\_(H \& R_3) = 1/8$

$P\_(T \& R_3) = 1/8$

$P\_(H \& R_4) = 1/8$

$P\_(T \& R_4) = 1/8$

Table 4

Applying the Lewis-Elga indifference principle to the uncentered possibilities in Table 4, Beauty next obtains a distribution of preliminary probabilities over the pertinent *centered* possibilities. The following table—rather than Table 1 above—is the 12-way probability distribution over centered possibilities that she arrives at:

$P\_(\text{Today}_{H,Mon} \& R_1) = 1/8$

$P\_(\text{Today}_{H,Mon} \& R_2) = 1/16$

$P\_(\text{Today}_{H,Mon} \& R_4) = 1/16$

$P\_(\text{Today}_{H,Tues} \& R_2) = 1/16$

$P\_(\text{Today}_{H,Tues} \& R_3) = 1/8$

$P\_(\text{Today}_{H,Tues} \& R_4) = 1/16$

$P\_(\text{Today}_{T,Mon} \& R_1) = 1/16$

$P\_(\text{Today}_{T,Mon} \& R_2) = 1/8$

$P\_(\text{Today}_{T,Mon} \& R_3) = 1/16$

$P\_(\text{Today}_{T,Tues} \& R_1) = 1/16$

$P\_(\text{Today}_{T,Tues} \& R_3) = 1/16$

$P\_(\text{Today}_{T,Tues} \& R_4) = 1/8$

Table 5

And if Beauty engages in synchronic conditionalization on R_1, using the preliminary probabilities in Table 5, she arrives at the conclusion that $P(H) = P(T) = ½$.

So goes the objection. The heart of it is the claim that for each i, $1 \leq i \leq 4$, the two uncentered possibilities (H & R_i) and (T & R_i) differ only in those respects determined by the outcome of a fair coin toss, with the information in I^* providing no further relevant evidence as to their comparative likelihood. We reply, on the contrary, that each pair of uncentered possibilities differs in another respect as well—a respect that *does* provide further evidence as to their comparative likelihood. (Thus, the Principal Principal cannot be correctly invoked in support of Table 4, and the preliminary probabilities in Table 5 are incorrect.) The information I^* provides further relevant evidence because of the way Beauty's essentially indexical information expressible as "I was awakened today by the experimenters" interacts with the rest of the information in I^* to generate, for each of the four equi-likely cases $Today_{H,Mon}$, $Today_{H,Tues,Mon}$, $Today_{T,Mon}$, $Today_{T,Tues}$, an asymmetric set of evidentially indifferent subcases. The asymmetry, as reflected in Table 1, is this: for each of these four cases, there are three respective sub-cases involving three respective rules from the set $\{R_1, R_2, R_3, R_4\}$, one of these being a once-awakening rule and two being twice-awakening rules. The evidential indifference is this: within each case, the information I^* is evidentially indifferent with respect to the three subcases in the case's sub-partition. This evidential indifference within asymmetric sets of subcases, together with the strongly symmetrical hierarchical partition-structure revealed in Table 1 induces the preliminary probabilities in Table 2, including the preliminary-probability distribution of 1/1/2 each to the twelve respective members of the finest-grained partition {1.a.i, 1.a.ii., …, 2bii., 2.b.iii}. Given that distribution, the correct 8-way distribution over uncentered possibilities of the form (H & R_i) and (T & R_i) is not the one in Table 5, but rather this:

P_(H & R_1) = P_($Today_{H,Mon}$) = 1/12

P_(T & R_1) = P_($Today_{T,Mon}$) + P_($Today_{T,Tues}$) = 1/12 + 1/12 = 1/6

P_(H & R_2) = P_($Today_{H,Mon}$) + P_($Today_{H,Tues}$) = 1/12 + 1/12 = 1/6

P_(T & R_2) = P_($Today_{T,Mon}$) = 1/12

P_(H & R_3) = P_($Today_{H,Tues}$) = 1/12

P_(T & R_3) = P_($Today_{T,Mon}$) + P_($Today_{T,Tues}$) = 1/12 + 1/12 = 1/6

P_(H & R_4) = P_($Today_{H,Mon}$) + P_($Today_{H,Tues}$) = 1/12 + 1/12 = 1/6

P_(T & R_4) = P_($Today_{T,Tues}$) = 1/12

Table 6

In short: the objection is mistaken in claiming that the Principal Principle applies to pairs of uncentered possibilities of the form (*H* & R_i), (*T* & R_i). It doesn't apply, because the information *I** provides further relevant evidence that bears on the comparative likelihood of (*H* & R_i) and (*T* & R_i). Thus, the halfer's attempt to refute our argument is unsuccessful.

Let us also add three comments on the dialectical state of play, as regards the objection and our reply to it. First, halfers could always just insist that the Principal Principle really does apply to pairs of possibilities of the form (*H* & R_i), (*T* & R_i); but to do so without offering an independently plausible rationale for that claim, and an independently plausible critique of our own argument for thirdism, would be tantamount to simply begging the question at issue.

Second, in replying as we do to the objection, we ourselves do rely on an independently plausible argument—viz., our earlier argument for the assignment of preliminary probabilities described in Table 2. This argument yields the 8-way probability assignment in Table 6 as a subsidiary conclusion, and thereby provides an underlying rationale for our claim that the Principal Principle does not apply to pairs of uncentered candidate-possibilities of the form (*H* & R_i), (*T* & R_i).

Third, what halfers would need to do, in order to fend off our argument in a non-question-begging way, would be to find some independently plausible reason for denying that in each of the four equi-likely candidate-possibilities 1.a, 1.b, 2.a, and 2b in Table 1, the three constituent finest-grain candidate-possibilities each have an equal preliminary probability of 1/12. This would require providing a credible reason why it should matter to the comparative likelihood of those three finest-grain possibilities whether they involve a once-awakening rule rather than a twice-awakening rule. But the trouble is that this difference doesn't seem to *make* a difference, as regards the question of comparative likelihood. Absent some credible basis for treating such differences as evidentially relevant, the default presumption—viz., that the evidence is indifferent with respect to the three candidate-possibilities—is in force and remains undefeated.

4. How this argument is immune to two attacks from Joel Pust

Joel Pust raises two separate objections to existing arguments for thirdism. While we are not convinced by either objection, we grant that they exert a considerable pull. Here we show that our new argument for thirdism is immune to both objections.

4a. Pust's objection to Horgan's original argument for thirdism

Terry Horgan (2004, 2007) argues for the thirder position, and his argument relies on the principle of generalized conditionalization. The reasoning goes as follows (with some minor changes in terminology). When Beauty is awakened by the experimenters on Monday, she considers the following four candidate-possibilities:

$\text{Today}_{H,Mon}$: *H* and today is Monday.

$\text{Today}_{H,Tues}$: *H* and today is Tuesday.

$\text{Today}_{T,Mon}$: *T* and today is Monday.

$\text{Today}_{T,Tues}$: *T* and today is Tuesday.

She first assigns preliminary probabilities to these four candidate-possibilities, relative to a certain proper portion *I*** of her total current information. *I*** excludes her information that she has been awakened today by the experimenters and is now conscious; she is bracketing this information, in assigning preliminary probabilities. (But *I*** does include her information that today is either Monday or Tuesday.) These preliminary probabilities are:

$P\_(\text{Today}_{H,Mon}) = 1/4$

$P\_(\text{Today}_{H,Tues}) = 1/4$

$P\_(\text{Today}_{T,Mon}) = 1/4$

$P\_(\text{Today}_{T,Tues}) = 1/4$

She also infers the following preliminary *conditional* probabilities for these four candidate-possibilities, relative to the same proper subset of her total current information (and letting 'A_{Today}' symbolize 'I have been awakened today by the experimenters, and I am now conscious', which expresses the bracketed information):

$P\_(\text{Today}_{H,Mon}|A_{Today}) = 1/3$

$P\_(\text{Today}_{H,Tues}|A_{Today}) = 0$

$P\_(\text{Today}_{T,Mon}|A_{Today}) = 1/3$

$P\_(\text{Today}_{T,Tues}|A_{Today}) = 1/3$

She now updates her preliminary non-conditional probabilities by the method of generalized conditionalization on A_{Today}, and infers the following updated absolute probabilities:

$P(Today_{H,Mon}) = 1/3$

$P(Today_{T,Mon}) = 1/3$

$P(Today_{T,Tues}) = 1/3$

And so,

$P(H) = P(Today_{H,Mon}) = 1/3$

$P(T) = [P(Today_{T,Mon}) + P(Today_{T,Tues})] = 2/3$

This reasoning is an instance of synchronic conditionalization, because it pertains to a space of (essentially indexical) synchronic candidate-possibilities, rather than to a space of diachronic candidate-possibilities that have been available ever since the agent first acquired pertinent information. However, unlike the reasoning involved in our new argument, the reasoning here does not involve conditionalizing on old information: the information conditionalized on – viz. the essentially indexical information *I was awakened today by the experimenters, and I am now conscious* – is newly acquired information that Beauty did not have on Sunday.

Joel Pust (2008) objects to Horgan's treatment of the Sleeping Beauty problem. He does not complain about generalized conditionalization per se, or about synchronic conditionalization as a species of generalized conditionalization.[13] His objection is more specific, and directly targets Horgan's claims (i) that Beauty should bracket the indexical information expressible as *I was awakened today by the experimenters, and I am presently conscious*, and (ii) that, in light of this bracketing, she should assign preliminary probability ¼ to the proposition *HEADS and today is Tuesday*. The objection is that it allegedly makes no sense to assign non-zero epistemic probability—even non-zero *preliminary* epistemic probability—to a proposition that entails that one is unconscious at the very moment that one is assigning that probability. His leading idea is that the preliminary epistemic probability of a proposition *p*, relative to a body of information *I** that is a proper subset of one's total body *I* of information pertinent to *p*, is to be understood as the non-preliminary probability that would be appropriate to assign to *p* if *I** were one's total information pertinent to *p*. But if *p* is an essentially indexical proposition that entails that one is presently unconscious, then one could *never* be in a situation in which it would be appropriate to assign non-zero probability to *p*; for, assigning a probability to a proposition requires that one be conscious when one does so.

Horgan (2008) replied by saying that although Pust's objection does have considerable prima facie plausibility, nonetheless the recommended reasoning concerning the Sleeping Beauty problem is legitimate and correct under a proper construal of epistemic probability. Horgan's claim was that the preliminary epistemic probability of a proposition *p*, relative to a body of information *I*** that is a proper subset of one's total body *I* of information pertinent to *p*, should be construed as the *quantitative degree of evidential support* that *I*** lends to *p*. Thus, if *I*** is the proper subset of Sleeping Beauty's pertinent information that excludes her information that she was awakened today by the experimenters, and also excludes her information that she is now conscious, then Beauty can perfectly well appreciate presently that the degree of evidential support that *I*** provides for the proposition *HEADS and today is Tuesday* is ¼—even though she also realizes that she would presently be unconscious if that proposition were true.[14] We broadly agree with this response to Pust's objection.[15]

Our new argument for thirdism presented in this paper, however, simply bypasses Pust's objection altogether. In our new argument for thirdism we do not claim that Sleeping Beauty should assign any non-zero preliminary probability to the claim that she is currently unconscious. We apply the principle of Generalized Conditionalization in such a way that Beauty brackets some information about the set-up of the experimental situation. The claim that Beauty is awake and conscious is not bracketed, but rather is part of the proper portion of her information relative to which Beauty assigns her preliminary probabilities. She thereby reaches the "thirder" conclusion via a version of generalized conditionalization that fully conforms to the way of assigning preliminary epistemic probabilities that Pust regards as legitimate.[16] The objection that Pust raised against Horgan's earlier treatment of the Sleeping Beauty Problem, whatever its dialectical force might be, simply does not arise vis-à-vis the reasoning just presented. Our new argument for thirdism is dialectically quite important—inter alia, because it allows us to remain neutral about the nature of epistemic probability, and it also allows us to remain neutral about Pust's (admittedly plausible-looking) construal of preliminary epistemic probability.

4b. Pust's second Objection

The second objection from Pust that we consider is a general attack on thirder arguments to-date, and also an attack on arguments for the standard halfer position.

Pust's key claim is that it is impossible to conditionalize on a temporally indexical claim such as 'the meeting starts now' (uttered at noon), or 'it is Monday now' as contemplated by Sleeping Beauty when she awakes. His reason for this claim is as follows:

> According to a standard formulation of the principle of conditionalization, when one learns *e* (and nothing stronger) at t_2, $P_{t2}(h)$ should be equal to one's previous degree of belief in *h* conditional on *e*, i.e. $P_{t1}(h/e)$. The principle of conditionalization, then, requires that the proposition newly certain at t_2 *had some value in one's prior credence distribution at* t_1. Hence, if there is some reason why even an ideally rational agent must lack *any* prior credence (even zero) for some proposition of which she becomes certain, then it will be impossible to conditionalize on that proposition (Pust forthcoming: 14).

Pust discusses various alternative accounts of indexical belief, and concludes that under every account a temporally indexical claim cannot have a value in one credence function at two distinct times. For example, one of the accounts that Pust considers involves the claim that an indexical belief is a relation between a person and a proposition – where that proposition has only limited temporal accessibility. Under this account, Pust claims, "the temporally indexical proposition grasped at *t* when one considers 'It is *now* Monday' cannot be grasped at a distinct moment in time, *t*'" (Pust forthcoming: 18). Given Pust's claim that one can conditionalize at t_2 on some newly certain claim only if that claim had some value in one's credence function at an earlier time t_1, it seems to follow that it is impossible to conditionalize on a temporally indexical claim.

Pust shows how two well-known arguments for thirdism (from Elga and Cian Dorr), and the standard argument for halfism, all rely on the assumption that we *can* conditionalize on temporally indexical claims. If the assumption is incorrect, then all of these arguments are unsound. Below we briefly explain how each of the arguments that Pust discusses rely on the assumption in question:

Consider first Elga's argument for thirdism (2000). Elga claims that Beauty's credence in HEADS on first waking should not vary depending on whether the coin is tossed on Monday night or Sunday night, so we can consider a version of the scenario in which the coin is tossed on Monday night. Elga claims that if after waking on Monday, Sleeping Beauty were to learn that it was Monday – and so that the coin toss had not even happened yet – then she ought to assign a probability of ½ to HEADS. Thus, Elga claims, when Beauty first wakes up on Monday she ought to assign a probability of ½ to

(HEADS/It is now Monday). Note that this move assumes that Beauty can conditionalize on the temporally indexical claim ‘It is now Monday’ – an assumption that Pust argues is false. Elga relies on the claim that Beauty should assign equal probabilities to (HEADS/It is now Monday) and (TAILS/It is now Monday), together with the claim that Beauty should assign equal probabilities to (TAILS/It is now Monday) and (TAILS/It is now Tuesday) to derive the claim that Beauty should give equal weight to each of these claims: ‘HEADS & It is now Monday’, ‘TAILS & It is now Monday’, ‘TAILS & It is now Tuesday’, and so that Beauty should assign a credence of 1/3 to HEADS.

Now consider Cian Dorr’s argument for thirdism (2002).[17] Dorr’s argument involves a modified version of the sleeping beauty scenario in which Beauty is certain that she will wake on both Monday and Tuesday with no memories of the previous waking, but also certain that if it is Tuesday and the coin landed heads then her memories of the previous waking will come flooding back a minute after she awakens. Upon first waking on Monday it seems that Beauty should assign equal probabilities of ¼ each to these claims: ‘HEADS & It is now Monday’, ‘HEADS & It is now Tuesday’, ‘TAILS & It is now Monday’, ‘TAILS & It is now Tuesday’. After a minute, when no memories come flooding back, she learns ‘not (HEADS & It is now Tuesday)’. At this point, Dorr claims that Sleeping Beauty ought to conditionalize on this temporally indexical claim, giving her a probability of 0 for ‘HEADS & It is now Tuesday’, and a probability of 1/3 for each of the remaining options. Dorr then argues that Beauty’s information state at this point in the modified version of the scenario is, in all relevant respects, identical to Beauty’s information state when she first wakes up on Monday in the original scenario. This yields the thirder conclusion. Again, this argument relies on the assumption that Pust disputes – that it is possible to conditionalize on a temporally indexical claim.

Pust also considers the motivation for the standard halfer position. The halfer typically claims (and Pust agrees) that when Sleeping Beauty first wakes up, the probability of ‘HEADS and it is now Monday’ is ½, the probability of ‘TAILS & It is now Monday’ is ¼, and the probability of ‘TAILS & It is now Tuesday’ is ¼. Assuming that Beauty can conditionalize on the temporally indexical claim ‘It is now Monday’, it follows that Beauty’s conditional probability P(HEADS/It is now Monday) is 2/3, and so that in a scenario in which Beauty comes to learn that it is Monday – and thus that the coin has not even been tossed yet – she should assign a probability of 2/3 to HEADS. Pust claims that – because temporally indexical claims such as ‘it is now Monday’ cannot be conditionalized on – halfers can avoid this unwelcome result and can adopt a 'double-halfer' position, according to which Sleeping Beauty's epistemic probability for HEADS is ½ both when she first wakes up and when she later discovers that it is Monday.

Thus various well-known arguments for the thirder position and an argument for the standard halfer position depend on the assumption that it is possible to conditionalize on a temporally indexical claim. Pust claims that this assumption is false – and thus rejects all of these arguments. Our main aim here is to point out that our new argument for the thirder position does not depend on the assumption that it is possible to conditionalize on a temporally indexical claim, and so is immune to Pust's objection. Our new argument applies the principle of Generalized Conditionalization in such a way that the claim conditionalized on is a claim about the set-up of the experiment – specifically R1. This claim is not temporally indexical. Thus Pust's objection does not apply to our argument for thirdism.

If Pust has succeded in showing that the other arguments for thirdism, and the argument for the standard halfer position are unsound, then our new argument would one of the few arguments in the area left standing. However we have two reasons to doubt that Pust's objection is entirely successful:

1. Each of the arguments that Pust discusses does indeed depend on the assumption that Beauty can conditionalize on a temporally indexical claim – but the relevant indexical need not be referring to an *instant* of time, but rather to a whole day. For example, Elga's argument works perfectly well if the claim that Sleeping Beauty conditionalizes on is 'it is Monday *today*' rather than 'it is Monday *now*'. The claim 'it is Monday *today*' seems to be a claim that Sleeping Beauty can grasp throughout the whole day, rather than just at a given instant. Presumably then the temporally indexical claim 'it is Monday today' could have a value in Beauty's credence function at two distinct times on Monday, t_1 and t_2. This would meet Pust's requirement for conditionalization.

   Pust mentions this issue in a footnote, where he writes: 'There is a tendency in the literature to formulate the relevant hypotheses regarding Beauty's temporal location in terms of days (e.g. 'It is Monday') rather than moments. This obscures the fact that when Beauty considers whether it is Monday, she is really considering whether it is *now* Monday, i.e. whether the present moment occurs on Monday' (Pust forthcoming: footnote 23). Presumably then Pust would claim that Beauty can contemplate the day that she finds herself in only indirectly – as 'the day that *now* is part of'. If this were so, then Pust could reasonably claim that claims involving 'today' cannot have a value in Beauty's credence function at two different times – even at two different times within the same day. But Pust has given us no reason to think that Beauty cannot contemplate the day that she finds herself in directly – rather than indirectly via contemplating the instant that she finds herself in. If Beauty could directly contemplate the day that she finds herself in, then claims involving 'today' could presumably have a value in Beauty's credence function at two different times during that day, and so Beauty could conditionalize on these claims. Pust's objection to standard thirder and halfer arguments would then fail.

2. In section 1 we defended the principle of Generalized Conditionalization. According to that principle, an agent can conditionalize on virtually any item in his corpus of information. *When* the item was learnt does not matter. Once it is recognized that the principle of Conditionalization can be generalized in this way, there is no reason to think that we can conditionalize only on claims that can be grasped at two distinct points in time. Thus even if we agree with Pust that there is only one instant in which a claim such as 'it is now Monday' (uttered by Sleeping Beauty when she first awakes) can be grasped, we can still make sense of the idea of Sleeping Beauty conditionalizing on this claim: we simply assign preliminary probabilities relative to a proper portion of Sleeping Beauty's information at that instant (the portion with the claim 'it is now Monday' removed) and then, by conditionalization on the claim 'it is now Monday', we obtain SB's updated probabilities at the very same instant.

5. Conclusion

The method of generalized conditionalization, when applicable, appears to be an entirely appropriate way to ascertain epistemic probabilities, because there is no logically significant difference between conventional conditionalization and generalized conditionalization. Synchronic conditionalization, which typically becomes applicable when a space of pertinent, essentially indexical, possibilities arises in the wake of one's losing pertinent self-locational information, is a version of generalized conditionalization. So synchronic conditionalization, when applicable, also appears to be an entirely appropriate way to ascertain epistemic probabilities. One way to implement synchronic conditionalization vis-à-vis the Sleeping Beauty problem is to bracket certain *old* information—specifically, the non-indexical information embodied in the proposition R_1. That implementation yields the conclusion that P(*HEADS*) = 1/3. This argument for thirdism is immune to both the objections raised by Pust.[18]

References

Arntzenius, F. 2003. "Some Problems for Conditionalization and Reflection." *Journal of Philosophy* 100: 356-370

Barnes, E. C. 1999. "The quantitative problem of old evidence." *British Journal for the Philosophy of Science* 50: 249-64.

Dorr, C. 2002. "Sleeping Beauty: In Defence of Elga" *Analysis* 62: 292-295

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

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.

Horgan, T. 2008. Synchronic Bayesian updating and the Sleeping Beauty problem: Reply to Pust. *Synthese* 160: 155-159.

Howson, C. 1984. "Bayesianism and support by novel facts." *British Journal for the Philosophy of Science* 35: 245-51.

Howson, C. 1985."Some recent objections to the Bayesian theory of support." *British Journal for the Philosophy of Science* 36: 305-9.

Howson, C. 1991. "The 'old evidence' problem." *British Journal for the Philosophy of Science* 42: 547-55.

Jeffrey, R. 1995. "Probability reparation: the problem of new explanation." *Philosophical Studies* 77: 97-101.

Lewis, D. 2001. "Sleeping Beauty: Reply to Elga." *Analysis* 62: 47-53.

Monton, B. 2006. "God, fine-tuning, the the problem of old evidence." *British Journal for the Philosophy of Science* 57: 405-424.

Pust, J. 2008. "Horgan on Sleeping Beauty." *Synthese* 160: 97-101.

Pust, J. 2011. "Sleeping Beauty, evidential support and indexical knowledge: reply to Horgan." *Synthese* (online first)

Pust, J. forthcoming. "Conditionalization and Essentially Indexical Credence." *Journal of Philosophy*.

---

[1] In another commonly discussed version, the coin is flipped on Monday evening while Sleeping Beauty is asleep, rather than Sunday evening. It is widely agreed among those who have discussed the problem, and we concur, that there is no important difference between the two versions. We focus on the Sunday-flip version here (at the suggestion of a referee) because doing so slightly simplifies exposition. But our discussion below is readily adaptable to the Monday-flip variant too.

[2] Here and throughout, we will use the expression 'epistemic probability' rather than the term 'credence'. The latter evokes subjectivist conceptions according to which epistemic probability is degree of belief, or rational degree of belief, or degree of confidence, or something of this sort. We think that epistemic probability is better construed as *quantitative degree of evidential support*, relative to a specific body of information. (Quantitative degree of evidential support should be distinguished from *qualitative* degree of evidential support, as expressible by locutions like 'makes highly likely', 'makes highly unlikely', makes somewhat likely', etc.) But for the most part, except where we indicate otherwise, our discussion in this paper will be neutral about the nature of epistemic probability.

[3] The expression 'candidate-possibility' is intended to accommodate the thought that some of the items that fall under this rubric will count sometimes as epistemically possible and sometimes as epistemically impossible, depending on the specific body of background-information relative to which one attributes epistemic probabilities to various propositions.

[4] Typically there are a number of ways one could bracket the information B. One way would be to excise not only B itself, but also any information that one actually possesses only by virtue of actually possessing B. We will call this *strong* bracketing of B. But another way, which we will call *weak* bracketing of B, would be to retain certain information that is logically weaker than B but which one actually possesses only by virtue of actually possessing B—e.g., certain information that is logically entailed by B, which does not logically entail B, and which one actually possesses only because it is logically entailed by B.

[5] Not surprisingly, the idea of generalizing the method of conditionalization in this way has arisen before, and has been invoked approvingly in the literature on "the problem of old evidence" in philosophy of science. See, for instance, Barnes (1999), Glymour (1980, 87-91), Howson (1984, 1985, 1991), Jeffrey (1995), and Monton (2006).

[6] *I*-B might be the proper portion of *I* that results from strongly bracketing B, or it might be some proper portion of *I* that results from weakly bracketing B in one way or another. Cf. note 4.

[7] Thus, R_1 is being bracketed *weakly* rather than *strongly* (cf. note 4 above).

[8] In correspondence, Joel Pust has pointed out that if (as Horgan assumed in his original argument for thirdism) it were acceptable sometimes to assign a non-zero preliminary probability to a proposition entailing that one is not presently conscious, then Horgan's original assignment of equal preliminary probabilities of ¼ each to the four hypotheses $Today_{H,Mon}$,$Today_{H,Tues}$ $Today_{T,Mon}$ and $Today_{T,Tues}$ could be justified this way: "by the conjunction of (a) the Elga-Lewis limited indifference principle, according to which indistinguishable centered worlds associated with the same possible world should get equal credence, and (b) a suitable version of the principal principle according to which possible worlds which differ only in the outcome of a fair coin toss should, in the absence of relevant evidence, be assigned the same credence" (quoting Pust). These two principles alone are not enough to justify the assignment of preliminary probabilities in Table 1; rather, one also must appeal to evidential indifference as grounds for assigning equal probabilities to the three subcases within each of the four cases 1.a, 1.b, 2.a, 2.b. But because the claim that the evidence is indifferent vis-à-vis the three subcases is so intuitively obvious here and so unproblematic-looking, this claim strikes us as beyond serious question—even though the task of formulating suitable general principles of probabilistic indifference remains difficult and elusive.

[9] By the definition of conditional probability, $P\_[(Today_{H,Mon} \& R_1)|R_1] = P\_[(Today_{H,Mon} \& R_1)]/P\_(R_1)$. This is equivalent to $P\_[(Today_{H,Mon} \& R_1)]/(P\_[(Today_{H,Mon} \& R_1)] + P\_[(Today_{T,Mon} \& R_1)]+ P\_[(Today_{T,Tues} \& R_1)]) = (1/12)/((1/12)+(1/12)+(1/12)) = 1/3$. One calculates $P\_[(Today_{T,Mon} \& R_1)|R_1]$ and $P\_[(Today_{T,Tues} \& R_1)|R_1]$ in a similar way.

[10] The conditionalization that leads to these conclusions can be implemented in a non-calculation-intensive way, directly on the basis of Table 2 and without doing the calculations that generate Table 3. First eliminate, from the original twelve finest-grained candidate-possibilities in Table 2 (each of which has an unconditional preliminary probability of 1/12), the nine finest-grained candidate-possibilities that are incompatible with R_1; then normalize one's probability distribution over the remaining three finest-grained candidate-possibilities—viz., ($Today_{H,Mon}$ & R_1), ($Today_{T,Mon}$ & R_1), and ($Today_{T,Tues}$ & R_1)—in a way that preserves the pairwise ratios (all 1:1) of their preliminary probabilities. These three possibilities now have updated probabilities of 1/3 each, replacing their preliminary probabilities of 1/12 each.

[11] A referee has suggested, with reference to our presentation of this argument in earlier drafts, that the argument would work equally well with a two-way disjunction of potential rules—e.g., (R_1 or R_2)—rather than the four-way disjunction we employ. This is not so, although our previous drafts failed to make the matter sufficiently clear. The trouble is that a two-way disjunction of rules would not generate a *strongly symmetric* hierarchical partition-structure of candidate-possibilities. Without strong symmetry, it becomes much more tendentious what counts as the correct way to assign preliminary probabilities to the various cells in the relevant partition-structure—which gives the halfers various opportunities for defensive resistance that are precluded when the partition-structure is strongly symmetric.

[12] This objection was posed by a referee.

[13] In fairness to him, we acknowledge that Horgan's prior writings have not explicitly articulated either notion of conditionalization as clearly or explicitly as would be ideal.

[14] This reply to Pust invokes the specific conception of epistemic probability described in note 2 above.

[15] Pust has recently argued, in response to Horgan (2008), that "the most plausible account of quantitative degree of [evidential] support, when conjoined with any of the three major accounts of indexical thought in such a way as to rationally constrain rational credence, contradicts essential elements of Horgan's argument" (Pust 2011, p. XX). We are inclined to think that if this conclusion is correct, then it points up the need for a more adequate account of indexical thought.

[16] Let 'R^+_1' abbreviate 'R_1 and I know that R_1'. Presumably, R^+_1 is an item of information that is known by Sleeping Beauty in her actual epistemic situation. Does that mean that her synchronic conditionalization should be keyed to R^+_1, rather than to R_1? No. For, in the envisioned epistemic situation she considers when bracketing the information R_1, the proposition R^+_1 has epistemic probability zero—which means that conditional probabilities that are conditional on R^+_1 are undefined in that envisioned epistemic situation. Conditionalization should be geared to one's strongest bracketed *conditionalizable-upon* information that is relevant to epistemic probabilities. For Sleeping Beauty, that information is R_1. In order to apply conditionalization, one must know (at the time one does so) the conditionalized-upon proposition—which arguably means that if one is sufficiently rational, one thereby *knows* that one knows that proposition. But it is not required that the knowledge-claim is itself a component of the conditionalized-upon proposition. (This note was prompted by correspondence with Pust.)

[17] Frank Arntzenius gives a closely related argument in Arntzenius 2003.

[18] Thanks to Adam Elga, two anonymous referees, and especially Joel Pust for helpful comments on earlier drafts.