Unformatted version, published in *Erkenntnis* 80 (2015): 811-839

Generalized Conditionalization and the Sleeping Beauty Problem, II

1. Introduction

The much discussed Sleeping Beauty problem involves a perfectly rational agent, Beauty, in an experiment in a sleep-research laboratory. She knows the following. She is put into dreamless sleep on Sunday night. She is awakened by the experimenters on Monday morning and then returned to dreamless sleep. If the toss of a fair coin lands heads, then she sleeps dreamlessly until Wednesday and awakens knowing that the experiment is over. If the toss lands tails, then her memories of the Monday awakening are erased and she is awakened by the experimenters again on Tuesday morning.

The problem is this. When Beauty is awakened during the experiment with no memories of a previous awakening—as she knew would happen at least once—what should be her epistemic probability for the statement "The coin toss lands Heads" (HEADS)? Some philosophers, "halfers," say that her epistemic probability for HEADS is ½. Other philosophers, "thirders," say that her epistemic probability for HEADS is 1/3.[1]

I am a thirder. I maintain that Beauty, upon being awakened in the lab with no memory of a previous awakening, now has new pertinent evidence, and that this new evidence combines with her other pertinent knowledge to yield an epistemic probability for HEADS of 1/3. The new evidence is temporally essentially indexical, and she can express it by saying "I was awakened today by the experimenters (and am now conscious and awake)." In several earlier papers (CITATIONS DELETED FOR BLIND REVIEW), I claimed that Beauty can reason soundly from this evidence to the conclusion that the epistemic probability of HEADS is 1/3 by deploying a form of conditionalization that I called "synchronic Bayesian updating." The reasoning proceeds in two stages. First she temporarily "brackets" her knowledge that she was awakened today by the experimenters (and is now conscious and awake), without bracketing her knowledge that today is either Monday or Tuesday; and, on the basis of her unbracketed pertinent evidence she assigns the following "preliminary probabilities" (as I call them) to the following statements:

HEADS and today is Monday, 1/4

HEADS and today is Tuesday, 1/4

TAILS and today is Monday, 1/4

TAILS and today is Tuesday, 1/4

Next, she updates these preliminary probabilities by conditionalizing on her knowledge that she was awakened today by the experimenters (and is now conscious and awake)—knowledge that precludes the second of the four preliminary possibilities just cited. This yields "synchronically updated" epistemic probabilities of 1/3 each for the remaining three preliminary possibilities—which entails that the epistemic probability of HEADS is 1/3.

Joel Pust is a halfer. He objects to the argument just described (Pust 2008). He maintains that Beauty could only have preliminary epistemic probabilities of ¼ each for the above four statements if there were a possible epistemic situation in which she lacks the "bracketed" evidence, possesses all and only the other relevant knowledge that she actually possesses (including the knowledge that today is either Monday or Tuesday), and *is conscious* while in that situation. But there is no such possible epistemic situation; for, if she lacked the information that she is now conscious while still possessing all and only the other pertinent information (including the information that today is either Monday or Tuesday), then she would "possess" that other information only dispositionally and non-occurrently, because the result of the coin toss would be heads, today would be Tuesday, and hence she would now be *dreamlessly asleep*.

I find this objection unpersuasive, for the following reason. Beauty knows that the above preliminary probabilities are the epistemic probabilities that would be attributed to the four propositions by a perfectly rational agent, *other than Beauty herself*, if (a) this agent did not know whether or not Beauty is awakened today by the experimenters, and (b) this agent's epistemic situation vis-à-vis Beauty and vis-à-vis the experiment were otherwise pertinently just like Beauty's actual epistemic situation. For instance, Beauty can imagine a hypothetical agent—call her Sleeping Betty—who knows everything about Beauty in the experiment that Beauty herself learned on Sunday, and who also knows the following. Betty herself is a subject in the experiment too, in a separate room from Beauty; Betty is definitely awakened from dreamless sleep by the experimenters twice, once on Monday and once on Tuesday; and Betty's Monday memories are definitely erased after her Monday awakening. Beauty can rightly say the following to herself:

> Suppose that this hypothetical rational agent, Sleeping Betty, were awakened in the lab by the experimenters without knowing whether or not they awaken *me* today. It would certainly be rational for Betty to assign epistemic probabilities of ¼ each to the above four statements. And if Betty were to then learn that the experimenters *do* awaken me today, then it would be rational for her to update by conditionalizing—excluding the statement "HEADS and today is Tuesday," and assigning updated epistemic probabilities of 1/3 each to the remaining three statements.
>
> Now, the fact that I myself am Beauty, and am not some other person like this hypothetical rational agent Betty, is not pertinent to the problem at hand. I can, and should, reason about this matter in a way that parallels how Betty would reason. First I should assign preliminary epistemic probabilities of ¼ each to the above four statements, because these are the all-in epistemic probabilities that Betty would assign to them upon being awakened in the lab by the experimenters and not knowing whether or not I, Beauty, am awakened today. Then I should conditionalize on these preliminary probabilities, factoring in my knowledge that I was awakened today by the experimenters—because that's what Betty would do, upon learning this fact about me.

In short, Beauty can rationally assign the above preliminary probabilities to the above four statements because these are the probabilities that would be assigned to those statements by some *other* rational agent whose total pertinent information about Beauty and the experiment were the same as Beauty's bracketed information about herself and the experiment. And Beauty can rationally do synchronic Bayesian updating using these preliminary probabilities, since such reasoning is parallel to the conventional, diachronic, conditionalization that she knows that Betty would use upon learning that Beauty is awakened today by the experimenters. Although it is true enough that Beauty herself couldn't be in Betty's epistemic situation, this doesn't matter.[2]

But although I myself find this response to Pust compelling, others may not. Furthermore, the rationale just given for Beauty's use of synchronic Bayesian updating, the rationale involving Sleeping Betty, falls prey to a *different* objection of Pust's that does not apply directly to the original argument itself. (Out of the frying pan, into the fire!) Pust raises this objection against several influential thirdist arguments (e.g., Elga 2000, Dorr 2002, Arntzenius 2003) that seek to apply conventional, diachronic, conditionalization to temporally indexical statements like "Today is Monday," in versions of the awakening scenario in which Beauty later in the day learns something pertinent about today itself (e.g., that today is Monday). Pust maintains, in response to such arguments—and this contention also applies to Beauty's lately-described soliloquy about Sleeping Betty—that "No temporally indexical claim can appear in temporally disjoint credence functions."[3]

I do not find this contention plausible, and others may agree—although we skeptics do need to confront the case he makes for it. But in any event, and given all this dialectical disputation, the case for thirdism certainly would be strengthened if a new and different argument could be provided for the 1/3 answer—say, an argument that somehow reveals the evidential relevance of "I was awakened today by the experimenters" without either (a) trying to bracket this very statement, or (b) trying to apply conventional, diachronic, conditionalization to some temporally indexical statement like "Today is Monday."

Horgan and Mahtani (2013) claimed to do exactly that. They offered a new argument for thirdism that relies on a form of conditionalization—they call it "generalized conditionalization"—that is even more of a generalization of conventional conditionalization than is synchronic Bayesian updating. Generalized conditionalization goes beyond conventional conditionalization in two respects: first, by deploying a space of synchronic, essentially temporal, candidate-possibilities that are not "prior" possibilities (as is also done by synchronic Bayesian updating); and second, by allowing for the use of preliminary probabilities that arise by first bracketing, and then conditionalizing upon, "old evidence"—evidence that one has had all along, rather than evidence that was acquired only recently and is now being conditionalized upon.

Pust (in press) replies to the Horgan/Mahtani argument, raising several objections. In my view his objections do not undermine the argument, but they do reveal a need to provide several further elaborations of it—elaborations that I think are independently plausible. In this paper I will address his objections, by providing the elaborations that I think they prompt. Along the way I will underscore some general morals that emerge concerning reasoning about epistemic probabilities, especially in cases involving losses and/or gains in essentially indexical, self-locational, information.

Before proceeding, let me say something about what is at stake. Since the present paper is a reply to a reply to a paper proposing one specific argument addressing one specific puzzle about probability, one might wonder whether the matters under dispute are too narrow to have much general philosophical interest. Not so. Here is a partial list of reasons why not. (1) The Sleeping Beauty problem itself has engendered an enormous amount of interest and discussion in philosophy since it was introduced into the philosophical literature in Elga (2000). (2) The problem connects to a wide variety of disputed philosophical questions; to quote Titelbaum (2013), "The problem raises unanswered questions concerning relative frequencies, objective chances, the relation between self-locating and non-self-locating information, the relation between self-location and updating, Dutch books, accuracy arguments, memory loss, indifference principles, the existence of multiple universes, and many-worlds interpretations of quantum mechanics" (p. 1003). (3) Although standard conditionalization cannot be applied to the Sleeping Beauty problem or to various related probability problems involving loss of self-locational or self-identity information—because, as Arntzenius (2003) says, "it is clear that conditionalization can only serve to "narrow down" one's degree of belief distribution (one really *learns* by conditionalization)" (p. 367)—there remains an important philosophical question whether or not some generalized variant of conditionalization can still apply to such problems. (4) The Horgan-Mahtani argument not only invokes generalized conditionalization, but also invokes a particular kind of symmetry as a basis for certain evidential-indifference claims; if such symmetry can legitimately be used this way—and the bulk of Pust's reply to Horgan and Mahtani seeks to show that it cannot be—then the form of reasoning deployed by Horgan and Mahtani not only constitutes a new way of defending thirdism about the Sleeping Beauty problem itself, but also is potentially applicable to a host of similar probability problems like those discussed in Arntzenius (2003).

2. The Horgan/Mahtani Argument

I will here re-state the Horgan/Mahtani argument, in a way that builds in more explicit dialectical structure than was present in their original formulation. The argument proceeds in five stages, each being a stage of reasoning that Beauty herself can employ.[4] The first stage consists in considering a hypothetical body of relevant information I* comprising some, but not all, of the pertinent information I that Beauty actually possesses when awakened by the experimenters. I* has the following two features. First, it excludes this information:

(R_1) I am awakened on Monday morning and then returned to sleep. If the toss of a fair coin lands heads, I sleep until Wednesday and awaken knowing that the experiment is over. If the toss of a fair coin lands tails, my memories of the Monday awakening are erased and I am awakened on Tuesday morning.   [Monday only if heads, both days if tails]

Second, although I* excludes the information R_1, nevertheless it includes a logically weaker fact consisting of R_1 disjoined with the disjunction of the following:

(R_2) I am awakened on Monday morning and then returned to sleep. If the toss of a fair coin lands tails, I sleep until Wednesday and awaken knowing that the experiment is over. If the toss of a fair coin lands heads, my memories of the Monday awakening are erased and I am awakened on Tuesday morning.   [Both days if heads, Monday only if tails]

(R_3) I am awakened on Tuesday morning and then returned to sleep. If the toss of a fair coin lands heads, I sleep until Wednesday and awaken knowing that the experiment is over. If the toss of a fair coin lands tails, my memories of the Monday awakening are erased and I am awakened on Tuesday morning.   [Tuesday only if heads, both days if tails]

(R_4) I am awakened on Tuesday morning and then returned to sleep. If the toss of a fair coin lands tails, I sleep until Wednesday and awaken knowing that the experiment is over. If the toss of a fair coin lands heads, my memories of the Monday awakening are erased and I am awakened on Tuesday morning.   [Both days if heads, Tuesday only if tails]

Beauty has here *bracketed* her information R_1, by excluding it from the portion of her total pertinent evidence that she is now considering. And she has *weakly* bracketed this information, in the following sense: rather than bracketing not only R_1 itself but also all information that she would not possess at all if she did not possess R_1 (this would be *strong* bracketing, vis-à-vis R_1), instead she includes in I* a logically weaker, disjunctive, item of "residue" information that she possesses only by virtue of also possessing the bracketed information R_1—viz., (R_1 v R_2 v R_3 v R_4).

At the second stage, Beauty organizes the space of pertinent possibilities, relative to the information I*, into the "hierarchical partition structure" given in Table 1. (Notation, in Table 1 and in all subsequent tables: HEADS = The coin toss lands heads; TAILS = The coin toss lands tails; MON = Today is Monday; TUES = Today is Tuesday.) She then makes an observation, about this partition structure, that Horgan and Mahtani put this way:

> This partition structure is *strongly symmetrical*, in this sense: the two cells in the outer partition {HEADS, TAILS} subdivide symmetrically into 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). (p. 338)

At the third stage, Beauty begins to assign preliminary probabilities to the various cells in Table 1—i.e., probabilities she regards as assignable relative to the information I* (and ignoring her bracketed information). She assigns the preliminary probabilities given in Table 2. Horgan and Mahtani describe her reasoning this way:

> In light of the fact that the space of candidate-possibilities exhibits the strongly symmetrical hierarchical partition-structure exhibited…, 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. (p. 335)[5]

Beauty further concludes that because of the strong symmetry, preliminary probabilities should be assigned in a parallel fashion for each set of three bottom-level cells within each of the four intermediate cells.

At the fourth stage, Beauty judges that for each set of three innermost cells in Tables 1 and *, the information I* is evidentially indifferent among the three cells. The idea is that the main difference between the three preliminary possibilities within each intermediate cell—viz., the fact that in two of them the experimenters awaken Beauty on another day, whereas in the third they do not—does not affect their comparative likelihoods. Horgan and Mahtani write:

> 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…. 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…, that the three candidate-possibilities within the set are equally likely, and hence each has preliminary epistemic probability 1/12. (p. 336)

The upshot of this reasoning at stage 4, in combination with the reasoning at stage 3, is the probability assignment given in Table 3.

At the fifth stage, Beauty invokes generalized conditionalization. On the basis of the preliminary probabilities in Table 3, she calculates the following preliminary *conditional* probabilities: P_(HEADS| R_1) = 1/3, P_(TAILS|R_1) = 2/3. She then updates her preliminary probabilities for HEADS and TAILS by conditionalizing on her information R_1, thereby concluding that P(HEADS) = 1/3 and P(TAILS) = 2/3.

Several features of this argument are noteworthy. First, it is indeed an application of *generalized* conditionalization, because (a) the conditionalized-upon information R_1 is "old" information that she already had back on Sunday, and (b) the main subcells and the sub-subcells in Tables 1-1** all represent preliminary epistemic possibilities that are essentially indexical (in their temporal self-location aspect).

Second, the argument is not susceptible to either of the two objections that Pust has raised against other arguments for thirdism: it does not assign non-zero preliminary probabilities to any statements entailing that Beauty is now unconscious, and it does not apply conditionalization to any temporally indexical statement like "Today is Monday."

Third, there is a strong and powerful rationale for generalized conditionalization—viz., that from a purely logical point of view, it makes no difference whether or not the conditionalized-upon information is newly acquired or previously possessed, and it also makes no difference whether or not the pertinent space of possibilities includes essentially indexical ones (sometimes called "centered possibilities").

Fourth, generalized conditionalization is actually much more common than is usually realized, because many uses of conditionalization that are typically regarded as instances of traditional, diachronic, conditionalization really are not that, strictly speaking. Suppose, for example, that one confronts a version of the Monty Hall problem that goes as follows. First the contestant chooses one of three doors. Next, Monty opens one of the other two doors, revealing that there is nothing behind it. *Thereafter*, Monty informs the contestant that he (Monty) knew where the prize is and deliberately opened an unchosen door with no prize behind it. Strictly speaking, the contestant was not able to assign *prior* probabilities to the various possibilities concerning Monty's door-opening, because the contestant didn't know the principle governing Monty's decision-making until after Monty had already opened one of the doors. Examples like this are legion. And although normally it is harmless enough to regard them as involving updating of probabilities by standard conditionalization—since it makes no difference to the pertinent reasoning whether or not one actually was able to ascertain the relevant conditional probabilities prior to acquiring the information upon which one then does conditionalization—nevertheless in such cases those conditional probabilities are not really *prior* probabilities at all.

Fifth, it seems beyond serious question that the strong symmetry exhibited in Table 1 constitutes a form of evidential indifference that yields the probabilities indicated in Table 2. Although less thoroughgoing forms of symmetry sometimes cannot be relied upon as yielding evidential indifference—as will emerge below—the fact remains that *this* kind of symmetry-based evidential indifference, in relation to *this* partitioning of preliminary possibilities vis-à-vis the information I*, should not be in dispute.

Sixth it is important to be clear about what exactly is being claimed about how strong symmetry bears upon the probabilities of the sub-sub-cells in Tables 1 and *. All that strong symmetry directly establishes, relative to information I*, is this: for any two of the four intermediate cells, probabilities should be assigned to the three sub-cells within the first intermediate cell *in a parallel fashion* to how probabilities are assigned to the three sub-cells within the second intermediate cell. Strong symmetry by itself does not justify assigning probability 1/12 to each of the sub-sub-cells, as is done in Table 3.

Seventh, stage 4 of the Horgan/Mahtani argument invokes an enormously plausible evidential-indifference claim, one that has only become applicable because I* excludes the information R_1 while yet including the symmetrical disjunctive information (R_1 v R_2 v R_3 v R_4). That claim can be put the following way. (I label it 'NDMD' as short for *no difference that makes a difference*.)

NDMD Relative to my information I*, which includes the fact that I was awakened today by the experimenters (with no memory of another awakening), I have no reason to regard the preliminary possibility that a once-awakening rule is in force as being either more likely or less likely than the preliminary possibility that a twice-awakening-cum-memory-obliteration rule is in force; and I also have no reason to regard any such twice-awakening rule as either more likely or less likely than any other.

To appreciate how plausible this principle is, just suppose that Beauty, upon having been awakened by the experimenters, now learns which of the four main sub-cells in Tables 1-1** describes the actual situation—say, cell 1.b, HEADS & TUES. Surely her total available evidence is now indifferent between the three possibilities that are consistent with HEADS & TUES, viz., R_2, R_3, and R_4 For, (i) because of the memory-erasure aspect of the two possibly operative rules R_2 and R_4 (the two that are twice-awakening rules vis-a-vis HEADS), her current experiential situation and her current total evidence would be exactly the same if the single-awakening rule R_3 were in force as it would be if either one of the twice-awakening rules R_2 or R_4 were in force; moreover (ii) the differences among these three possible scenarios would not involve any differences in essentially indexical temporal-location information, since she knows full well that today is Tuesday. Likewise, mutatis mutandis, for the sub-sub-cells in Table 2: one simply applies this same kind of indifference reasoning to each intermediate cell in the partition, on the supposition that one's actual situation corresponds to that intermediate cell.

Eighth, and in light of the previous two points, although the strong symmetry of the partition structure in Tables 1 and 1* does not itself justify assigning probability 1/12 to each of the sub-sub-cells, what it does do is *enable* the applicability of the overwhelmingly plausible indifference claim lately mentioned. And, given the first three stages of the Horgan/Mahtani argument, that additional indifference claim justifies the assignments of 1/12 each at stage 4.

Ninth, it is important to appreciate that the indifference reasoning deployed at stage 4 is much less tendentious than the following form of indifference reasoning concerning the Sleeping Beauty problem:

> Since my available evidence would be exactly the same in each of the three essentially-indexical possible situations HEADS&MON, TAILS&MON, TAILS&TUES, each of these essentially-indexical possibilities has an epistemic probability of 1/3.[6]

Halfers, of course, would object vigorously to such reasoning, claiming that if two non-indexical possibilities both have probability ½, and only one essentially indexical possibility is associated with the first non-indexical possibility whereas two distinct essentially indexical possibilities are associated with the other non-indexical possibility, then the right way to reason by indifference is this: leave the probabilities of the two non-indexical possibilities at ½, and apply indifference to the two essentially-indexical possibilities associated with the second non-indexical possibility. Doing indifference reasoning this way, one gets the following probabilities for the three essentially indexical possibilities at issue:

> HEADS&MON, 1/2
>
> TAILS&MON, 1/4

TAILS&TUES, 1/4

Thus, the question of how to correctly apply indifference reasoning to these three essentially-indexical possibilities is at the very heart of the disagreement between halfers and thirders. By contrast, however, the indifference principle NDMD is far less tendentious, and thereby is independently far more plausible. When one supposes, regarding one of the intermediate cells in Table 2, that this cell describes one's actual situation, one is making a supposition that *holds fixed* the temporal location that one supposes is now designated by 'today' (and also holds fixed the supposed outcome of the coin flip), and one considers each of three possibilities concerning what happens on the *other* day. Under the operative supposition, there is simply no independently plausible reason to think that a particular one of the two twice-awakening rules is either more likely or less likely to be in force than the single once-awakening rule—or that a particular one of the two twice-awakening rules is either more likely or less likely to be in force than the other twice-awakening rule. So, when one asks, for a given intermediate cell in Table 2, how to assign probabilities to the three finest-grained cells within that intermediate cells (each involving a different rule), the following is true. Since those three rules differ among themselves only about what happens on the other day and not about about what happens today, and since one's today-evidence would be exactly the same regardless of rich rule were in force, one's evidence is indifferent with respect to which rule is in force.

Finally, tenth, the Horgan/Mahtani argument provides a recipe, potentially applicable to a variety of tricky and tendentious probability puzzles, including those involving losses and/or gains in essentially indexical self-locational information. First, find some crucial bit of non-indexical "old information" that generates some tendentious form of asymmetry in the space of pertinent possibilities—an asymmetry that renders problematic the question of how correctly to apply indifference reasoning. Then, weakly bracket that information, leaving behind a symmetrical disjunction within which the bracketed information constitutes one disjunct. Then seek out a hierarchical partition structure, over the preliminary possibilities consistent with the non-bracketed information, that (a) is strongly symmetrical, (b) justifies assignments of preliminary probabilities to some of the cells in the partition, on grounds of strong-symmetry based evidential indifference, and (c) enables the use of some additional, highly plausible (or even self-evident) indifference claim(s) to justify assignments of preliminary probabilities to the remaining cells of the partition. Then apply generalized conditionalization, conditionalizing on one's bracketed information.[7]

Table 1

1. HEADS
    - 1.a. MON
        - 1.a.i. R_1
        - 1.a.ii. R_2
        - 1.a.iii. R_4
    - 1.b TUES
        - 1.b.i. R_2
        - 1.b.ii. R_3
        - 1.b.iii. R_4
2. TAILS
    - 2.a. MON
        - 2.a.i. R_1
        - 2.a.ii. R_2
        - 2.a.iii. R_3
    - 2.b. TUES
        - 2.b.i. R_1
        - 2.b.ii. R_3
        - 2.b.iii. R_4

Table 2

1.	HEADS (1/2)		
	1.a.	MON	(1/4)
		1.a.i.	R_1
		1.a.ii.	R_2
		1.a.iii.	R_4
	1.b	TUES	(1/4)
		1.b.i.	R_2
		1.b.ii.	R_3
		1.b.iii.	R_4
2.	TAILS	(1/2)	
	2.a.	MON	(1/4)
		2.a.i.	R_1
		2.a.ii.	R_2
		2.a.iii.	R_3
	2.b.	TUES	(1/4)
		2.b.i.	R_1
		2.b.ii.	R_3
		2.b.iii.	R_4

Table 3

1.	HEADS (1/2)		
	1.a.	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	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.	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.	TUES	(1/4)
		2.b.i.	R_1 (1/12)
		2.b.ii.	R_3 (1/12)
		2.b.iii.	R_4 (1/12)

3. An Alternative Preliminary Probability Distribution

Pust's first objection is that there is another way to assign preliminary probabilities, involving a different hierarchical partition structure, that seems at least as justified as the assignment of preliminary probabilities in Table 1—an assignment which would lead, by generalized conditionalization, to a probability of ½ for HEADS. The alternative partition structure and the alternative probability assignment are given in Table 4. Pust says this regarding Table 4:

> Importantly, the probabilities assigned at each level of the hierarchical partition in Table [4] appear *at least as justified* as those assigned at each level in the alternative hierarchical partition in Table [3]. The two tables agree at the first stage, assigning the same probabilities to the highest-level partition over {HEADS, TAILS}. Table [4] reverses the ordering of Table [3] with respect to the subsequent two levels in the hierarchical structure. On evidential symmetry grounds it assigns equal probabilities to the elements of the {R_1, R_2, R_3, R_4} sub-partition, and then, on the basis of the Elga-Lewis restricted principle of indifference, assigns equal probabilities to each day with the rule-based sub-partition.
>
> This shows, I believe, that Horgan and Mahtani have provided no reason to regard the preliminary probabilities in Table [3] as *more reasonable* than the preliminary probabilities in Table [4]…. Hence, Horgan and Mahtani's argument fails to be dialectically compelling at the first step, there being no clear reason to prefer the preliminary probability [assignment] required by their case for the thirder position to that apparently favorable to the halfer position. (p. 13)

This objection prompts the following elaboration of the Horgan/Mahtani argument. The relevant kind of "strong symmetry," to which they appealed as justification for the probabilities assigned to the top level and the middle level of the hierarchical partition structure in Table 1, should be understood as involving both a top-down aspect and a bottom-up aspect. In particular, as regards the cells in the middle level, the partition structure exhibits symmetry both (a) with respect to the cells above (those in the top level), and (b) with respect to the cells below (those in the bottom level). In general, one cannot safely appeal to symmetry in a hierarchical partition structure, as grounds for treating the cells within a given level as being equally probable, unless that level exhibits both top-down and bottom-up symmetry.

Table 4

1.	HEADS	(1/2)		
	1.a.	R_1	(1/8)	
		1.a.i.	MON	(1/8)
	1.b.	R_2	(1/8)	
		1.b.i.	MON	(1/16)
		1.b.ii.	TUES	(1/16)
	1.c.	R_3	(1/8)	
		1.c.i.	TUES	(1/8)
	1.d.	R_4	(1/8)	
		1.d.i.	MON	(1/16)
		1.d.ii.	TUES	(1/16)
2.	TAILS	(1/2)		
	2.a.	R_1	(1/8)	
		2.a.i.	MON	(1/16)
		2.a.ii.	TUES	(1/16)
	2.b.	R_2	(1/8)	
		2.b.i.	MON	(1/16)
	2.c.	R_3	(1/8)	
		2.c.i.	MON	(1/16)
		2.c.ii.	TUES	(1/16)
	2.d.	R_4	(1/8)	
		1.d.i.	TUES	(1/8)

Table 5

1.	HEADS		(1/2)	
	1.a.	R_1	(1/12)	
		1.a.i.	MON	(1/12)
	1.b.	R_2	(1/6)	
		1.b.i.	MON	(1/12)
		1.b.ii.	TUES	(1/12)
	1.c.	R_3	(1/12)	
		1.c.i.	TUES	(1/12)
	1.d.	R_4	(1/6)	
		1.d.i.	MON	(1/12)
		1.d.ii.	TUES	(1/12)
2.	TAILS	(1/2)		
	2.a.	R_1	(1/6)	
		2.a.i.	MON	(1/12)
		2.a.ii.	TUES	(1/12)
	2.b.	R_2	(1/12)	
		2.b.i.	MON	(1/12)
	2.c.	R_3	(1/6)	
		2.c.i.	MON	(1/12)
		2.c.ii.	TUES	(1/12)
	2.d.	R_4	(1/12)	
		1.d.i.	TUES	(1/12)

The rationale that Pust offers for the probability assignment in Table 4 does not fit this justificatory format. Although the partition structure in Table 4 does exhibit top-down symmetry with respect to cells in the middle level, it fails to exhibit bottom-up symmetry for that level. Assigning equal probabilities to those cells, some of which have more sub-possibilities than others, would therefore *not* be a safe invocation of evidential symmetry. Indeed, I claim that doing so would be outright *mistaken*. Since this partition structure is not strongly symmetrical (in the relevant sense), the correct way to assign preliminary probabilities to cells in the structure is to read them off from Table 3. The result is given in Table 5.

Two further points deserve emphasis. First, the probability assignment in Table 5 conforms to the Elga-Lewis indifference principle, which requires indistinguishable "centered" possibilities associated with a given "uncentered" possibility all to have the same epistemic probability. I need not, and do not, question this highly constrained indifference principle.[8]

Second, although I have here elaborated somewhat upon what Horgan and Mahtani mean by "strong symmetry," I have not defined this notion (and they did not try to do so in the other paper). Although seeking out an adequate, sufficiently general, definition of strongly symmetrical hierarchical partition structures is certainly a philosophically worthwhile task, such a definition is not required for present purposes. It's enough that one can recognize strong symmetry, or its absence, when one sees it.[9] Strong symmetry in the space of preliminary epistemic possibilities, when present, constitutes a useful form of evidential indifference that can ground—or partially ground, perhaps in combination with other forms of evidential indifference, and/or with available information about matters of objective chance—the assignment of preliminary probabilities; and preliminary probabilities that are thus grounded can then be subjected to generalized conditionalization. (I do offer a proposed definition of strong symmetry in Appendix 1, plus some additional observations about this notion in Appendix 2.)

4. Giving Chance Its Due

Pust's second objection is that the Horgan/Mahtani strong-symmetery based account "cannot generalize, as it must, to…cases in which the *chances* of HEADS and TAILS are not equal" (p. 13). He writes:

> Horgan and Mahtani claim that "in light of the fact that the space of candidate-possibilities possesses the strongly symmetrical partition structure exhibited by Table 1, Beauty concludes that the evidence provided by the information I* is indifferent with respect to the two outer cells HEADS and TAILS… She therefore assigns preliminary probabilities of ½ each to HEADS and TAILS" (338). *This* rationale for the assignment of preliminary probabilities to HEADS and TAILS makes no mention of chance. Surely, however, known chances must constrain in some way the assignment of preliminary probabilities to HEADS and TAILS, no matter how such credence is subsequently further divided amongst further sub-partitions and sub-sub-partitions. (p. 14)

Fair enough. I agree that Horgan and Mahtani were too cavalier in not mentioning the known chances of HEADS and TAILS, and in presenting their argument as though they think that the assignment of preliminary probabilities of ½ each to HEADS and TAILS is warranted solely on the basis of the strong symmetry of the partition structure in Table 1. So another elaboration is in order. What they should have said is this: because of the strong symmetry of that partition structure, HEADS and TAILS can safely and correctly be assigned preliminary epistemic probabilities that are equal to their known objective chances.[10]

Given this needed emendation, their argument does indeed generalize—as Pust rightly insists that it must. Suppose, for instance, that one knows that the coin is biased 3-to-1 in favor of TAILS. Then the appropriate preliminary probabilities are given in Table 6. And, when one applies generalized conditionalization on the bracketed information R_1, one obtains an epistemic probability for HEADS of 4/19.

Table 6

1.	HEADS (1/4)			
	1.a.	MON	(1/8)	
		1.a.i.	R_1	(1/24)
		1.a.ii.	R_2	(1/24)
		1.a.iii.	R_4	(1/24)
	1.b	TUES	(1/8)	
		1.b.i.	R_2	(1/24)
		1.b.ii.	R_3	(1/24)
		1.b.iii.	R_4	(1/24)
2.	TAILS	(3/4)		
	2.a.	MON	(3/8)	
		2.a.i.	R_1	(1/8)
		2.a.ii.	R_2	(1/8)
		2.a.iii.	R_3	(1/8)
	2.b.	TUES	(3/8)	
		2.b.i.	R_1	(1/8)
		2.b.ii.	R_3	(1/8)
		2.b.iii.	R_4	(1/8)

A further point is very important to emphasize, regarding the relation between known chances and epistemic probabilities. Sometimes, because one's pertinent evidence exhibits certain kinds of asymmetries, these asymmetries will generate epistemic probabilities that deviate from known chances. Here is an example. I am staying at a hotel in your city, and we agree in a telephone conversation to meet for breakfast tomorrow at 8:00 a.m. in the hotel restaurant. I know—and I reliably inform you—that the hotel has two elevators, that the management has apologetically informed me that one of the elevators operates more slowly than the other one, and that they didn't say which one. (Thus, although we both know that the two elevators do not have equal *chances* of arriving first at a given floor when the call-button is pushed, nonetheless our evidence is *indifferent* with respect to which of them is more likely to arrive first at any given floor after the call-button is pushed on that floor.) We agree—and you know I will keep the agreement—on the following contingencies. Tomorrow morning I will flip a fair coin just before I press the call-button for the elevator. If the coin comes up Heads and the left elevator picks me up, then I will say to you in the morning, "I'm buying breakfast." If the coin comes up Tails and the left elevator picks me up, I will say to you in the morning, "You're buying breakfast." If the right elevator picks me up, then (regardless of the outcome of the coin toss) I will say to you in the morning, "I'm buying breakfast." When you arrive at the hotel restaurant the next morning I'm already there, and I say to you "I'm buying breakfast."

What is your epistemic probability, in this situation, for HEADS? Well, after you and I made our agreement last evening, you were able to assign the following probabilities to the cells in the following hierarchical partition structure (letting 'L' symbolize 'The left elevator picks him up', 'R' symbolize 'The right elevator picks him up', and 'B' symbolize 'He says that he's buying breakfast'):

HEADS (1/2)

L & B (1/4)

R & ~B (1/4)

TAILS (1/2)

L & B (1/4)

R & B (1/4)

These prior probabilities, in turn, entail the following prior conditional probabilities: P(HEADS|B) = 1/3, P(TAILS|B) = 2/3. So, having just heard me say "I'm buying breakfast," you should conditionalize on this new information and set your current epistemic probability for HEADS at 1/3—even though the known chance of HEADS is ½.

The moral of this example is clear, and should be borne firmly in mind: when one's prior (or bracketed) evidence generates a hierarchical partition structure that (a) has main cases with prior (or preliminary) epistemic probabilities that are equal to known chances, and (b) has sub-cases (or sub-sub-cases, etc.) whose prior (or preliminary) epistemic probabilities result from evidential indifference, then conditionalizing on new (or bracketed) information that asymmetrically excludes some of the sub-cases (or sub-sub-cases, etc.) is apt to yield epistemic probabilities for the main cases that deviate from their known chances. This moral will figure importantly below.

5. An Alternative Bracketing

Pust's next objection is that the Horgan/Mahtani strong-symmetry based account "cannot generalize, as it must,…to *alternative bracketings* of R_1." He writes:

> Horgan and Mahtani recognize that there are a variety of ways in which one might subtract or "bracket" information from one's current total information. So, if it is not to yield inconsistent results, generalized conditionalization on the credence distribution produced by one way of bracketing a given piece of information from one's total knowledge must yield *the same* updated probabilities as any other way of bracketing that information or any other piece of information from one's total knowledge. As R_1 can be "weakly bracketed" [i.e., bracketed in such a way that the residual unbracketed information includes a disjunctive statement with R_1 as one disjunct] in any number of ways, preliminary probability distributions on such alternative bracketings must, when subjected to generalized conditionalization, all yield the same result as the assignment relative to the bracketing to which Horgan and Mahtani appeal. (p. 15)

I fully agree, and I maintain that the Horgan/Mahtani strong-symmetry based account *does* generalize in this way, as indeed it must if it is sound. But Pust goes on to say the following, with reference to Table 7:

> Consider, for example, the hierarchical partition structure in Table [7], produced by weakly bracketing R_1 and including in I* the three-way disjunction of R_1, R_2, and R_3 instead of the

> logically weaker four-way disjunction of R_1, R_2, R_3, and R_4. There is no way to justify the thirder result by generalized conditionalization on preliminary probabilities over Table [7] *if* those preliminary probabilities are assigned to the sub-partitions and sub-sub-partitions by appeal to symmetry *within* a given partition, as it appears they must be…. Rather, given generalized conditionalization on R_1, such a rationale would yield results which contradict those Horgan and Mahtani defend for Table [3]…. We may conclude, I think, that Horgan and Mahtani have failed to provide a defensible and coherent account of how preliminary probabilities are to be assigned and updated via generalized conditionalization. (pp. 15-16)

My response is as follows. First, Table 7 does not exhibit the kind of strong symmetry that makes for safe and correct appeals to evidential indifference in assigning preliminary probabilities to various cells in the given hierarchical partition structure. For, although the mid-level cells 1.a, 1.b, 2.a, and 2.b do exhibit top-down symmetry within the partition structure, they fail to exhibit bottom-up symmetry. (Cell 2.a has three sub-cells, whereas cells 1.a, 1.b, and 2.b each have only two sub-cells.) Second, because the partition structure in Table 7 lacks the pertinent kind of strong symmetry, it is not even safe to assume that the main cells in the structure—viz., HEADS and TAILS—have preliminary epistemic probabilities that are identical to their known chances. Third (and bearing in mind these first two points), a safe and acceptable way to assign preliminary probabilities to the various cells in Table 7 would be to apply the Horgan/Mahtani method to Table 7 itself. (As it were, do some pre-preliminary bracketing in order to obtain a strongly symmetrical pre-preliminary hierarchical partition structure whose cells are assigned the suitable pre-preliminary probabilities, and then do generalized conditionalization to obtain the right preliminary probabilities for the cells in Table 7.)

Table 7

1. HEADS
    - 1.a. MON
        - 1.a.i. R_1
        - 1.a.ii. R_2
    - 1.b. TUES
        - 1.b.ii. R_2
        - 1.b.ii. R_3
2. TAILS
    - 2.a. MON
        - 2.a.i. R_1
        - 2.a.ii. R_2
        - 2.a.iii. R_3
    - 2.b. TUES
        - 2.b.i. R_1
        - 2.b.ii. R_3

Table 8

1.	HEADS (4/9)			
	1.a.	MON	(2/9)	
		1.a.i.	R_1	(1/9)
		1.a.ii.	R_2	(1/9)
	1.b.	TUES	(2/9)	
		1.b.ii.	R_2	(1/9)
		1.b.ii.	R_3	(1/9)
2.	TAILS	(5/9)		
	2.a.	MON	(3/9)	
		2.a.i.	R_1	(1/9)
		2.a.ii.	R_2	(1/9)
		2.a.iii.	R_3	(1/9)
	2.b.	TUES	(2/9)	
		2.b.i.	R_1	(1/9)
		2.b.ii.	R_3	(1/9)

Here is how to do that. Weakly bracket the information (R_1 v R_2 v R_3), while retaining as unbracketed the logically weaker information (R_1 v R_2 v R_3 v R_4). Then take steps 2-4 of the five-step reasoning described in section 2 above; this yields Table 3, a strongly symmetrical hierarchical partition structure with the appropriate pre-preliminary probabilities filled in. Then, using the pre-preliminary probabilities in Table 3, do generalized conditionalization on the information ~R_4, to obtain the correct preliminary probabilities for the cells in Table 7. The result is given in Table 8. One can now use the preliminary probabilities in Table 8 to do generalized conditionalization a second time, now conditionalizing on the information R_1. And the result, as it should be, is this: P(HEADS) = 1/3, P(TAILS) = 2/3. So the alternative bracketing used to generate Table 7 makes no trouble for the Horgan/Mahtani method, and provides no reason to think that this method fails to generalize to alternative bracketings.

Two further points deserve emphasis. First, the Horgan/Mahtani method crucially involves choosing a bracketing that will generate a *strongly symmetrical* hierarchical partition structure. Strong symmetry is needed in order to for the method to get safely started—in order to have the pertinent kind of rationale (a) for assigning preliminary probabilities, to cells with known chances, that do not deviate from those known chances, and (b) for assigning preliminary probabilities to various other cells by appeal to symmetry-based evidential indifference, and (c) for enabling the use of one (or several) highly plausible additional evidential-indifference claim(s)—in the case of the argument in section 2, the NDWD claim—in assigning the remaining probabilities.

Second, the moral that emerged at the end of Section 4 applies directly to Table 7. When the probabilities in Table 3 are deployed, in using generalized conditionalization on ~R_4 as a basis for determining the correct preliminary probabilities in Table 7, this results in the *asymmetrical exclusion* of various sub-sub-cases in Table 7. Because of that, the correct preliminary probabilities for HEADS and TAILS in Table 7 end up deviating from the known chances—as indicated in Table 8.

6. Rationality and Knowledge of One's Own Current Epistemic Probabilities

Pust's final objection is this:

> [G]iven (a) a conception of epistemic probability with which Horgan and Mahtani seek to make their argument compatible, and (b) a highly plausible principle of rationality, generalized

conditionalization is coherent *only when* it is equivalent to a *possible* diachronic update, whereas the instance of generalized conditionalizeion required by Horgan and Mahtani's new argument is *not* equivalent to any such possible diachronic update, being essentially synchronic. (p. 18)

A conception of epistemic probability with which we seek to make our argument compatible is one that incorporates the following two constraints, each of which Pust has employed elsewhere against other certain thirder (and against some halfer arguments too):

> The Epistemic Possiblity Constraint (EPC): Epistemic probabilities are equal to possible rational degrees of belief. (p. 18)
>
> The Temporal Indexical Constraint (TIC): No temporally indexical claim can in appear in temporally disjoint credence functions. (p. 21)

Pust further elaborates his objection this way:

> Horgan and Mahtani's argument appeals only to an essentially synchronic instance of generalized conditionalization, one which appeals only to candidate epistemic possibilities which Beauty can grasp in her current situation upon awakening, and so satisfies TIC.
>
> However, while such an essentially synchronic instance of generalized conditionalization is required to satisfy TIC, no such instance of essentially *synchronic* conditionalization can, given a plausible constraint on rationality independent of the Sleeping Beauty problem, respect EPC. The plausible constraint is that a perfectly rational agent will have complete knowledge of her own *current* credences. It follows that it is impossible to bracket out *only* a non-indexical partition of her current knowledge and apply Generalized Conditionalization in the way required by Horgan and Mahtani's argument for 1/3.
>
> In virtue of her perfect rationality, Beauty will have, at every moment, perfect confidence in her (then) current credences. So, in her epistemic situation upon awakening, $P(P_{now}(R_1) = 1) = 1$. Given that the instance of generalized conditionalization at issue is essentially synchronic, it follows that Beauty cannot "update" on the knowledge which must be excluded from her total knowledge, *I*, in order to yield *I**. After all, if $P\_(P_{now}(R_1) = 1) = 0$ then the preliminary conditional probability for *any* proposition conditional on R_1 & $P_{now}(R_1) = 1$ must be undefined and such "updating" is incoherent. (p. 22-23)

This objection prompts the following two-part elaboration of the Horgan-Mahtani approach. First, when one brackets the information R_1 for purposes of assigning preliminary probabilities, one also should bracket one's knowledge that one currently assigns epistemic probability 1 to R_1, plus one's current knowledge that one currently assigns probability 1 to the statement "I currently assign probability 1 to the statement 'I currently assign probability 1 to R_1'," etc. Second, one should apply generalized conditionalization to one's strongest *pertinent* bracketed evidence—where bracketed evidence counts as pertinent, in the operative sense, only if it is *compatible with one's preliminary probabilities*.

The second condition is extremely plausible. For, it is a direct consequence of the natural and intuitive rationale for conditionalization (both conventional conditionalization and generalized conditionalization). As a prelude to articulating this rationale, a definition will be useful. When a rational agent possesses evidence E, the agent thereby also possesses various higher-order items of information concerning the first-person epistemic status of E—e.g., "I know that E," "I assign epistemic probability 1 to E," and the like. Let a higher-order item of information about the first-person epistemic status of E be *directly possession-dependent* upon E just in case (1) possession of E by a fully rational agent would guarantee possession of that higher-order information-item by that agent, and (2) non-possession of E by a fully rational agent would guarantee non-possession of this higher-order item by that agent.

The rationale for the second above-mentioned condition can now be formulated as follows. While holding in abeyance a given item of evidence E that one currently possesses (where E might be either newly acquired evidence or "old evidence"), one inquires about the current epistemic probability of a statement S, P(S). One reasons that P(S) should be equal to the conditional epistemic probability of S given E, P_(S|E), that would obtain for a rational agent who (a) lacks the evidence E, (b) thereby also lacks any higher-order information-items concerning the epistemic status of E that are strongly possession-dependent upon E itself, but (c) otherwise possesses all and only the relevant evidence vis-à-vis S that one actually possesses oneself. So, for purposes of applying conditionalization, the strongest *relevant* evidence that one actually possesses, but would not be possessed by the envisioned rational agent in the envisioned epistemic situation, is E itself—not, say, "E and I know that E," or "E and my epistemic probability for E is 1," or the like. It would not make sense to treat any of these latter statements as expressing one's to-be-conditionalized-upon evidence, because even though they all express information that one *actually* possesses (by virtue of actually possessing the evidence E), each of them is a statement that that a rational agent in the envisioned counterfactual situation *would know to be false* (by virtue of *not* possessing evidence E, in that situation).[11] This is so both for conventional conditionalization and for generalized conditionalization.

The point is reinforced by the fact that many applications of conditionalization that are commonly regarded as deploying prior probabilities really deploy preliminary probabilities that are not literally prior probabilities at all. Consider, for example, the version of the Monty Hall problem I described in section 2, in which Monty *first* opens one of the doors other than the contestant's chosen door *and then* tells the contestant that he (Monty) knew where the prize was and deliberately opened a door that was both unchosen and prizeless. (Suppose that the contestant has chosen door 3 and Monty has opened door 2.) The contestant can *only then* ascertain pertinent conditional probabilities as a basis for conditionalization—which means that these will be preliminary probabilities that are not literally prior probabilities. If it were really true that in this case, conditionalizion requires first assigning preliminary probability zero to "I know that Monty opens door 2" and then trying to conditionalize using the putative preliminary conditional probabilities

> the preliminary probability that the prize is behind door 1 given that Monty opens door 2 and I know that Monty opens door 2

and

> the preliminary probability that the prize is behind door 3 given that Monty opens door 2 and I know that Monty opens door 2,

then the attempt to conditionalize would require using fractions with denominator zero. But surely the contestant *can* conditionalize in this case. And surely the pertinent preliminary conditional probabilities are not those just mentioned, but rather these:

> the preliminary probability that the prize is behind door 1 given that Monty opens door 2,
> the preliminary probability that the prize is behind door 3 given that Monty opens door 2.

So although Pust is right that Beauty, in order to apply generalized conditionalization to Table 7, must begin by bracketing not just R_1 but also her knowledge that her epistemic probability for R_1 is 3 (and her knowledge that she knows that her epistemic probability for R_1 is 1, etc.), the fact remains that her strongest *pertinent* bracketed evidence vis-à-vis HEADS and TAILS is just R_1 itself. Thus there is nothing incoherent about the reasoning that Horgan and Mahtani recommend, and this reasoning does not violate either Pust's Epistemic Possibility Constraint, or his Temporal Indexical Constraint, or his contention that a perfectly rational agent will have complete knowledge of her own current epistemic probabilities.

7. The Horgan/Mahtani Argument Re-Stated

Let me know reformulate the Horgan/Mahtani argument for the 1/3 answer to the Sleeping Beauty problem, in a way that incorporates the various aspects of elaboration introduced above. Beauty, having been awakened today by the experimenters, brackets her knowledge of R_1, plus her knowledge that she assigns epistemic probability 1 to R_1, plus her knowledge that she assigns epistemic probability 1 to 'I assign epistemic probability 1 to R_1', etc. She considers a body of pertinent evidence I* that excludes this bracketed knowledge, that includes the logically weaker information (R_1 v R_2 v R_3 v R_4), and also includes all the other pertinent evidence she actually possesses. She now contemplates the hierarchical partition structure in Table 1, and asks herself how to assign preliminary probabilities, relative to the information I*, to the cells in this structure She notes several aspects of symmetry in the structure: (1) the cells in the middle level exhibit top-down symmetry, because each of the top cells HEADS and TAILS has the same pair of sub-cases, {MON, TUES}; and (2) the cells in the middle level exhibit bottom-up symmetry consisting of these facts: (a) each of the middle cells has as sub-cases exactly three of the four possibilities in the partition { R_1, R_2, R_3, R_4}, (b) each of the middle cells has as sub-cases two twice-awakening rules and one once-awakening rule, and (c) each of the four possibilities in { R_1, R_2, R_3, R_4} occurs in the bottom level exactly three times. Beauty rightly thinks to herself, "This hierarchical partition structure exhibits the kind of strong symmetry that (i) makes it epistemically safe to equate the preliminary probabilities of HEADS and TAILS with their known chances of ½ each, (ii) makes it epistemically safe to treat MON and TUES as equally likely sub-cases of HEADS and as equally likely sub-cases of TAILS (and therefore, given (i), makes it epistemically safe to assign preliminary probability ¼ to each of the middle cells), and (iii) makes it epistemically safe to follow the same procedure, for each of the four middle cells, in how one assigns preliminary probabilities to its three sub-sub-cells." On the basis of this rationale, Beauty rightly assigns the preliminary probabilities that are indicated in Table 2.

Concerning the remaining cells in Table 2 (the sub-sub-cells), Beauty rightly thinks this to herself: "Each of the four possible situations depicted in the middle level of the partition structure holds fixed the presumed referent of 'today' (and also holds fixed the presumed outcome of the coin flip), and then subdivides into sub-sub-cases involving three possibly-operative rules. Since those subcases do not differ among themselves concerning what happens today, but instead only differ among themselves concerning what happens on the other day during the sleep experiment, there is no difference among the subcases that makes a difference to their comparative likelihood; rather, they are equally probable." On the basis of this rationale, Beauty assigns preliminary probabilities of 1/12 each to the sub-sub-cells in Table 2, thereby obtaining the preliminary probabilities that are indicated in Table 3.

Beauty appreciates that generalized conditionalization is a sound way to reason about epistemic probability. One can correctly apply generalized conditionalization after first bracketing virtually any portion of one's total available evidence that one likes; and, if one chooses to *weakly* bracket some portion of one's evidence, then one can include in the unbracketed information I* virtually any disjunctive statement at all that includes the unbracketed information as one disjunct. Beauty also appreciates, however, that great care is needed in assigning preliminary probabilities to the cells in hierarchical partition structures that result from bracketing. Unless the resulting partition structure is *strongly* symmetrical, in general it will not be epistemically safe to treat symmetries as a basis for evidential indifference; and often it will not even be epistemically safe, as regards cells in the partition structure that have known chances, to assume that the epistemic probabilities of those cells are equal to their known chances.

Beauty has been careful in just this way. She has deliberately deployed a mode of weak bracketing that does yield a strongly symmetrical hierarchical partition structure, and she has safely assigned to the cells in that structure the preliminary probabilities indicated in Table 3.

Beauty also appreciates that a perfectly rational agent will have complete knowledge of her own current epistemic probabilities, and that she herself has such knowledge. She appreciates too that appropriate applications of generalized conditionalization on old evidence should deploy not one's *total* bracketed information, but rather that portion of one's total bracketed information that is consistent with the preliminary probabilities that accrue to one's unbracketed information. For, only that portion of her total bracketed information counts as pertinent bracketed *evidence*. (This is a consequence of the underlying rationale for both conventional conditionalization and generalized conditionalization, as she understands.) She realizes that because of this, her strongest pertinent bracketed evidence is R_1—not, say, "R_1, and I know R_1, and my epistemic probability for R_1 is 1."

So, using the preliminary probabilities in Table 3, she applies generalized conditionalization to the bracketed information R_1. She correctly—and soundly—concludes that her epistemic probability for HEADS is 1/3.

8. Making Thirdism Intuitively Plausible

One reason why halfism exerts a strong intuitive tug is that Beauty is completely unsurprised when she finds herself being awakened by the experimenters with no memory of a prior awakening. She already knew, back on Sunday, that this would happen. This can make it seem that her newly acquired, essentially indexical, knowledge expressible as "I was awakened today by the experimenters" is not relevant new evidence that bears on the epistemic probability of HEADS. Another reason for the intuitive tug of halfism is a tendency to think that unless one knows the *actual* outcome of an event whose possible outcomes have known chances, then the epistemic probabilities of those outcomes are identical to the known chances. And probably these two (mistaken) tendencies reinforce one another, resulting in a net intuitive tug that is especially strong.

One thing that I myself want, from a proposed solution to a vexing probability puzzle, is that this solution engages my intuitions: it makes the recommended answer *seem right intuitively*. For me, this first happened when I hit upon the argument described in the third paragraph of Section 1 above. I said to myself, "Ah, the information that I was awakened today by the experimenters excludes one sub-case of HEADS—one possibility about *today*—that is consistent with the conditions governing the experiment that I learned about back on Sunday: viz., the sub-case in which the coin-toss lands Heads, today is Tuesday, and I'm now in a dreamless sleep." In the moment when this thought occurred to me, I converted from halfism to thirdism. I stand by this reasoning, and for some people (like me) such reasoning is a way of finding thirdism intuitively correct. But I admit that it is least *weird* to assign a preliminary probability of ¼ to a statement entailing that one is presently unconscious. And Pust, as I said at the outset, claims that doing so makes no sense.

The Horgan/Mahtani argument for thirdism is more complicated. But perhaps it can provide a different way to engage one's intuitions. Here is my suggestion for how to make this happen.
First, bear in mind the question "Why and how does the information expressible as "I was awakened today by the experimenters" generate an epistemic probability for HEADS that differs from the known

chance of ½? Second, focus on the coin-flip/elevator example in Section 4 above, and on the reason why the information that the other person is buying breakfast generates an epistemic probability of 1/3 for HEADS. (The key is asymmetrical sub-case exclusion.) Then, while focusing on the preliminary probabilities in Table 3 (which themselves should seem intuitively correct), take note of this fact: when one starts with the preliminary probabilities in Table 3, and then uses generalized conditionalization on R_1 as a basis for determining the epistemic probability of HEADS, the result is not only (a) that the information R_1 effects asymmetrical sub-sub-case exclusion on Table 3 (and thereby effects the asymmetrical exclusion of the entire sub-case 1.b), but also (b) that this exclusion-asymmetry arises *because of* the way that the sub-sub-cases in Table 3 are themselves determined by the information "I was awakened today by the experimenters." This latter information is the reason why R_1 occurs twice as a sub-sub-case under TAILS, while only occurring once as a sub-sub-case under HEADS—and thus is the underlying reason why generalized conditionalization on R_1 effects an asymmetrical form of sub-case exclusion that renders the epistemic probability of HEADS different from its known chance of ½.

Appendix 1: Strongly Symmetrical Hierarchical Partition Structures

I will propose a definition of strong symmetry for hierarchical partition structures, and will comment on its applicability to the partition structure in Table 1 above. First, let a *hierarchical partition structure*, relative to a space of epistemic possibilities and fully distinct partitions P_1 P_2,…P_n of that space, be a set $\{X_1, X_2, \ldots, X_n\}$ of sets of cells (respectively, the set of first-level cells, the set of second-level cells, …, the set of nth-level cells) such that $n > 1$ and

(1) for each i such that $1 \leq i < n$,

(a) each cell in X_i is filled by exactly one member of a partition P_i,

(b) each cell in X_i has one or more sub-cells in X_{i+1},

(d) for each cell in X_i and each element e in partition P_{i+1}, e fills at most one sub-cell of X_i,

(e) every sub-cell of X_i is in X_{i+1}, and

(2) each cell in X_n has no sub-cells and is filled by exactly one member of partition P_n.

I suggest the following definition. A hierarchical partition structure $\{X_1, X_2, \ldots, X_n\}$ is *strongly symmetrical* just in case:

for each i such that $1 \leq i < n$,

(1) every cell in X_i has exactly as many sub-cells as every other cell in X_i,

(2) every member of partition P_i occurs exactly as many times within cells in X_i as does every other member of P_i,

(3) for every two members z and w of the partition P_i, the number of cells in X_{i-1} containing a sub-cell of X_i in which z occurs is the same as the number of cells in X_{i+1} containing a sub-cell of X_i in which w occurs, and

(4) if $x_{i-1,1}$, $x_{i-l,2}$, …, $x_{i-1,r}$ are the successive cells in X_{i-1} and $Y_{i,1}$, $Y_{i,2}$, …, $Y_{i,r}$ are the successive sets of sub-cells in X_i of $x_{i-1,1}$, $x_{i-l,2}$, …, $x_{i-1,r}$ respectively, then each member of the partition P_i occurs within exactly the same number of sets from $\{Y_{i,1}, Y_{i,2}, \ldots, Y_{i,r}\}$ as does every other member of P_i.

The partition structure shared by Tables 1- 3 and 6 above satisfies this definition of strong symmetry (whereas the partition structures in the other Tables above do not). Clause 1 is satisfied by the second level of the partition structure, because each first-level cell has the same number of sub-cells (viz., 2), and each second-level cell has the same number of sub-cells (viz., 3). Clause 2 is satisfied because MON and TUES both occur the same number of times within the second level (viz., 2), and each of R_1, R_2, R_3, and R_4 occurs the same number of times within the third level (viz., 3). And clause 4 is satisfied because MON and TUES both occur in the same number of sub-cell sets attached to level-1 cells (viz., 2 such sub-cell sets), and each of R_1, R_2, R_3, and R_4 occurs within the same number of sub-cell sets attached to level-2 cells (viz., 3 such sub-cell sets).[12]

Appendix 2: Full Symmetry, Symmetrical Symmetry Breaking, and Strong Symmetry

Let a *downward path* through an n-level hierarchical partition structure be a set consisting of an element of partition P_1 that fills some level-1 cell in that structure, plus an element of partition P_2 that fills some level-2 sub-cell of that level-1 cell, plus …, plus an element that fills some level-n sub-cell of that level-(n-1) cell. A hierarchical partition structure is *fully* symmetrical just in case the following holds: for every combination of one element each from P_1, P_2, …, P_n, there is a corresponding downward path through the structure.

The hierarchical partition structure in Table 1 above is not fully symmetrical; rather, it contains only 12 downward paths, whereas a fully symmetrical one would have 16 downward paths—one for each combination of one element each from {HEADS, TAILS}, {MON, TUES}, and {R_1, R_2, R_3, R_4}. Table 1 thus exhibits a certain kind of symmetry breaking—a breaking of full symmetry. Nonetheless, the manner in which full symmetry is broken is itself symmetrical, in the following way: each of the 4 level-2 possibilities has as sub-possibilities exactly 3 elements of {R_1, R_2, R_3, R_4}, and each of the four elements of {R_1, R_2, R_3, R_4} is excluded exactly once as a sub-possibility of some level-2 possibility. Although strong symmetry is a more inclusive category than full symmetry, a hallmark of strong symmetry is that it can only deviate from full symmetry by exhibiting a form of symmetry breaking that is itself symmetrical. Strong symmetry—either full symmetry or symmetrically broken full symmetry—is the feature of a hierarchical partition structure that makes for epistemic safety both (a) in equating epistemic probabilities with known chances, and (b) in assigning certain other epistemic probabilities on grounds of symmetry-based evidential indifference.

References

Arntzenius, F. (2003). Some problems for conditionalization and reflection. *Journal of Philosophy, 100,* 356-370

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,* 143-147.

Hajek, A. (2003). What Conditional Probability Could Not Be. *Synthese*, *137*, 273-323.

Horgan, T. (2000). The Two-Envelope Paradox, Nonstandard Expected Utility, and the Intensionality of Probability. *Nous, 34,* 578-602.

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 Pus. *Synthese, 160,* 155-159.

Horgan, T. and Mahtani, A. (2013). Generalized conditionalization and the sleeping beauty problem. *Erkenntnis, 78,* 333-351.

Pust, J. (2008). Horgan on sleeping beauty. *Synthese, 160,* 97-101.

Pust, J. (2012). Conditionalization and essentially indexical credence. *Journal of Philosophy, 109,* 295-315.

Pust, J. (2013). Sleeping beauty, evidential support and indexical knowledge: reply to Horgan. *Synthese, 190,* 1489-1501.

Pust, J. (In press). Beauty and conditionalization: reply to Horgan and Mahtani. *Erkenntnis*.

Titelbaum, M. (2013). Ten Reasons to Care About the Sleeping Beauty Problem. *Philosophy Compass* 8/11: 1003-1017.

---

[1] Here and throughout, I use the expression 'epistemic probability' rather than the popular term 'credence'. This is because 'credence' is usually glossed either as degree of belief or as rational degree of belief, and I believe that epistemic probabilities are not degrees of belief. (I also believe that there are no such psychological states as degrees of belief.) Nothing in this paper turns on these claims, by my believing them does motivate me to eschew the term 'credence'.

[2] This reply to Pust's objection is neutral about the nature of epistemic probability. In particular, it is consistent with the widely held view that epistemic probability is so-called "credence"—i.e., quantitative *degree of belief* (or quantitative *rational* degree of belief), on a zero-to-one scale. If epistemic probability is construed this way, then one's own preliminary probabilities are naturally thought of as the degrees of belief that would be possessed by a perfectly rational agent whose pertinent evidence exactly matches the pertinent *unbracketed* evidence that one possesses oneself. My point is that such an agent need not be *oneself*. In an earlier essay (CITATION DELETED FOR BLIND REVIEW), I reply to Pust's objection a different way: I appeal to my own preferred construal of epistemic probability as *quantitative degree of evidential support*. On that construal, I argue, it is just irrelevant that Beauty could not be in an epistemic situation in which her total relevant evidence coincides with the unbracketed portion of her actual relevant evidence.

[3] Pust (2012) argues that this claim is a consequence of each of the three major accounts of indexical thought in the philosophical literature. In a similar vein, Pust (2013) argues that the above-described argument for the 1/3 answer to the Sleeping Beauty problem cannot be vindicated by construing epistemic probability (as I do) as quantitative degree of evidential support, because "the most plausible account of quantitative degree of support, when conjoined with any of the three major accounts of indexical though in such a way as to plausibly constrain rational credence, contradicts essential elements of Horgan's argument" (p. 1489). Addressing these arguments in detail would take a paper in itself. But, at least as regards the argument in Pust (2013), I think the upshot would be the following. Any

adequate account of indexical thought must respect the fact that 'believes that…' contexts and epistemic-probability contexts both are intensional. In particular, if two co-referring time-denoting terms are such that one of them is temporally indexical and the other is not, then no adequate account of indexical thought can entail that the two terms are inter-substitutable *salva veritate* within belief contexts and within epistemic-probability contexts. But I think that the argument in Pust (2013) implicitly assumes that the evidential-support relation is not intensional in this way, and instead allows such substitutions. This I would strongly deny. (For more on the intensionality of epistemic probability, with application to the two-envelope paradox, see Horgan 2000.)

[4] The value and importance of structuring the argument into five distinct stages, and of adding the present section's subsequent commentary on the argument as so structured, became clear to me from an anonymous referee's remarks on an earlier draft of the present paper.

[5] It would be better to say that Beauty assigns these preliminary probabilities on the basis of the combination of (i) the strong symmetry of the hierarchical partition-structure, and (ii) what she knows about the objective chances of HEADS and TAILS. Roughly, she knows that HEADS and TAILS each have an objective chance of ½. More accurately, she (a) she knows that HEADS and TAILS each have an objective chance *very close* to ½, and (b) her evidence is indifferent about whether, if the chances are slightly different, the chance of HEADS is slightly greater than the chance of TAILS or vice versa. (For simplicity, hereafter I will ignore the more accurate fact and I will acquiesce in the commonly made assumption that the known chance of an ordinary coin coming up heads when flipped is ½, and likewise for its coming up tails.) I return to this theme in section 4 below.

[6] I am using the locution 'essentially indexical possibility', rather than the recently popular locution 'centered possibility', because so-called centered possibilities are often described as though they are *metaphysical* possibilities with "designated centers." (If one also conceives of a metaphysical possibility as a set of metaphysically possible worlds, then the corresponding centered possibility will be a set that results from the first one by designating one and the same center for each of the metaphysically possible worlds in the original set. I.e., it will be a set of (metaphysically-possible-and-with-designated-center) worlds, all with the same designated center.) It seems to me, however, that a temporally indexical self-locational term like 'today' should be viewed as a context-dependent rigid designator: on the day of its usage, it rigidly designates that very day and no other. Thus, on the day of its usage, there are no metaphysically possible worlds in which today is any other day than the day currently rigidly designated by 'today'. So the essentially indexical possibilities in question in the Sleeping Beauty problem cannot all be rightly regarded as being metaphysical-possibilities-with-designated-centers. Rather, they are epistemic possibilities *not all of which are metaphysically possible*. Nothing in the current paper turns on these claims, but my believing them does motivate me to eschew talk of "centered possibilities."

[7] The recipe just described is readily adaptable to various other puzzles in the philosophical literature about probability and essentially indexical possibilities—for instance, to the Vishnu/Brahma self-duplication scenarios discussed in section V of Arntzenius (2003). Halfers about the Sleeping Beauty problem are apt to balk at some of Arntzenius's claims about the epistemic probabilities that accrue to the outcomes of coin flips in these scenarios; the Horgan-Mahtani recipe can be deployed to bolster those claims.

[8] Although I would prefer to formulate it in terms of essentially-indexical possibilities and non-indexical possibilities. Cf. note 6.

[9] And in any event, a proposed definition would have to be tested for adequacy by assessing it for conformity with pre-theoretic intuitive judgments about scenarios that do—and scenarios that do not—exhibit the kind of evidential symmetry that safely makes for indifference-based assignment of probabilities. Proposed theoretical definitions, for philosophically interesting concepts that figure in pre-theoretical reasoning, are accepted or rejected on *abductive* grounds—with the data for abduction consisting largely of people's pre-theoretic intuitive judgments about the applicability, or non-applicability, of the pertinent concept in concrete scenarios. In this connection, a referee points out that Carnap, in seeking to develop a systematic theory of confirmation, initially considered a confirmation function c-dagger that was strongly symmetrical, but later gave it up because it did not vindicate what he called "learning from experience." The confirmation functions he subsequently employed—the function c*, and more complicated successors—did not exhibit strong symmetry. The referee remarks that "allowing for learning from experience is exactly the kind of conformity with pre-theoretic intuitive judgments that the author should be looking for." By way of brief response, let me say the following. First, I certainly do not mean to suggest that evidential indifference is always grounded in strong symmetry—or even that the only evidential-indifference considerations operative in the Horgan-Mahtani argument are considerations of strong symmetry. (Step 4 invokes the further indifference claim I labeled NDMD.) Second, unlike Carnap, who ambitiously sought a systematic confirmation theory that would assign epistemic probabilities to virtually any proposition an agent might consider, I myself believe that principled epistemic probabilities arise only very rarely, and only under quite special circumstances—e.g., when one knows the objective chance of a proposition and one can rightly treat this known chance as its epistemic probability, or when one confronts a partition of n possibilities (for known n) over which one's evidence is indifferent (in which case each cell in the partition has epistemic probability 1/n). For the most part, learning from experience is not a phenomenon that justifies the assignment of specific epistemic probabilities at all. Rather, normally such information-gain only justifies certain kinds of qualitiative judgments of epistemic likelihood—e.g., qualitative non-comparative judgments of the form "Proposition p is highly probable," "Proposition p is improbable," etc., and qualitative comparative judgments of the form "Proposition p is more probable than

proposition q," "Propositions p and q are equally probable," etc. So in my view, the principal kinds of intuitive pre-theoretic judgments about the evidential import of symmetry that are pertinent to epistemic probability are intuitive judgments that arise in those relatively rare circumstances in which there is a principled basis for assigning any epistemic probabilities at all.

[10] On some construals of chance, a possible outcome that initially has a chance other than either one or zero takes on either a chance of 1 upon occurring or a chance of 0 upon failing to occur. But the epistemic probability of an outcome can remain equal to its "known chance" even if one knows that by now the outcome has either come about or failed to come about—provided that one's evidence is indifferent about the outcome.

[11] It might be argued, furthermore, that a rational agent in the envisioned epistemic situation could not have a conditional probability for S that is conditional on any one of these latter statements anyhow, because the fraction-formula expressing such a putative conditional probability would have a denominator that is equal to 0, and hence this fraction-formula would be undefined. (Pust himself argues that way, in the final sentence of the lately quoted passage.) However, Alan Hajek (2003) makes a strong case against treating the familiar ratio formula $P(A|B) = P(A\&B)/P(B)$ as a definition of the notion of conditional probability (which is what Kolmogorov did in his classic axiomatization of probability theory), and Hajek also makes a strong case for the meaningfulness of certain conditional probabilities where the condition-statement has probability zero. But I do not think that Hajek's arguments are directly pertinent here, because (as just argued in the text) there is independent motivation for saying that Beauty should apply generalized conditionalization by invoking preliminary conditional probabilities involving R_1 as the condition-statement—a statement whose preliminary probability is ¼—rather trying to apply generalized conditionalization by invoking putative preliminary conditional probabilities involving $[R_1 \& (P_{now}(R_1) = 1)]$ as the condition-statement—a statement whose preliminary probability is zero.

[12] Acknowledgements.