A mathematical analysis of the Sleeping Beauty problem

by

Jeffrey S. Rosenthal*

(August 2008; revised March 2009.)

1. Introduction.

The *Sleeping Beauty problem* (Elga, 2000; see also Piccione and Rubinstein, 1997) is a philosophical dilemma related to conditional probability. It may be succinctly described as follows. Sleeping Beauty is put to sleep, and a fair coin (say, a nickel) is tossed. If the nickel shows heads, then Beauty is interviewed on Monday only, while if the nickel shows tails, Beauty is interviewed on both Monday and Tuesday (and given an amnesia-inducing drug between the two interviews, so she does not remember the first interview during the second). In each interview, without access to any additional information (such as the result of the coin toss, or the existence of any previous interviews, or the day of the week), Beauty is briefly woken and is asked to assess the probability that the nickel showed heads. The question is, what probability should she assign to this?

One possible answer (e.g. Lewis, 2001; Arntzenius, 2002; Bostrom, 2007; Pust, 2008) is 1/2. After all, the coin is fair, so Beauty surely would assess the probability of heads as 1/2 *before* being put to sleep. Now, when Beauty is awakened and interviewed, she apparently does not gain any new information, since she knew in advance that she would be interviewed at least once no matter what. So, if the probability of heads was 1/2 before, it seems plausible that this probability should be unchanged during the interview, giving a final answer of 1/2. This answer seems rather intuitive (and was, in fact, the author's first reaction upon hearing of this problem). But is it correct?

Another possible answer (e.g. Elga, 2000; Dorr, 2001; Monton, 2002; Weintraub, 2004; Horgan, 2004; Neal, 2007; Titelbaum, 2008) is 1/3. One argument for this is that Beauty will be interviewed twice as often when the nickel shows tails as when it shows heads. Thus, if during each interview she makes a bet in which she will win \$1 on tails but lose \$2 on

*Department of Statistics, University of Toronto, 100 St. George Street, Room 6018, Toronto, Ontario, Canada M5S 3G3. Email: jeff@math.toronto.edu. Web: http://probability.ca/jeff/

heads, then she will break even on average (and also, by the Law of Large Numbers, break even over the long run if the experiment is repeated many times). But for this to be a fair bet, the probability of heads must be 1/3. This argument is somewhat convincing. However, one concern is that it does not appear to make explicit use of the amnesia aspect, i.e. it appears to still apply if we instead assume that Beauty is permitted to remember any previous interviews. But under that assumption, the conclusion seems incorrect, since if Beauty remembers having a previous interview, she then would immediately know that the nickel was tails with probability 1.

Some authors have accepted certain parts of both of these solutions. For example, although Arntzenius (2002) argues that Beauty should assign probability 1/2 to heads, he nevertheless agrees that "if she bets according to her degree of belief of 1/2, she can be expected to lose money against a bookie, and she and the bookie know this in advance". In other words, he feels that the question of Beauty's assigned probabilities is distinct from the question of fair betting odds, the former being 1/2 and the latter being 1/3.

A different approach is to appeal to the Principle of Indifference, which asserts that equal probabilities should be assigned to any collection of indistinguishable, mutually exclusive and exhaustive events. But in this case, it is not clear to what "collection" this principle should be applied. If we apply it to "nickel heads" and "nickel tails", we obtain an answer of 1/2. If we apply it to "nickel heads and interview Monday", "nickel tails and interview Monday", and "nickel tails and interview Tuesday", we obtain an answer of 1/3. If we apply it to "interview Monday" and "interview Tuesday" we conclude that with probability 1/2 the interview will be on Monday, and then a second application implies that the probability of heads is $(1/2)(1/2) = 1/4$. In short, the Principle of Indifference does not appear to resolve the problem satisfactorily.

To a mathematical probability theorist such as myself, such controversy is frustrating. We are being asked to compute the conditional probability that the nickel showed heads, conditional on the fact that Beauty is currently being interviewed. Conditional probabilities are well understood and should be unambiguously analysable by straightforward mathematics, using the classic formula $P(A \mid B) = \frac{P(A \text{ and } B)}{P(B)}$. So, how could this simple conditional probability problem create such controversy? The difficulty seems to be that a precise mathematical interpretation of "conditional on currently being interviewed" is unclear, thus creating an obstacle to direct mathematical calculation.

This paper attempts to reconsider the problem in such a way that precise mathematical reasoning can then be applied. After such reconsideration, we then obtain the answer 1/3 through direct calculation. It is hoped that this mathematical approach avoids most of the philosophical ambiguities inherent in some previous arguments.

2. A Subproblem: the Sleeping Peon.

Consider the following simple subproblem. We find a Peon and put him to sleep, and then flip a fair coin, say a nickel. If the nickel shows tails, we wake Peon and interview him (just once), asking him to assess the probability that the nickel showed heads. If the nickel shows heads, then we flip a second fair coin, say a dime. If the dime shows tails, we similarly wake the Peon and interview him once. If not (i.e., if the nickel and dime both show heads), then we do not bother to wake or interview Peon at all. In summary, we interview Peon once if either the nickel or the dime show tails, otherwise we interview him zero times. Hence, the overall probability that Peon is interviewed is equal to 3/4. Under these circumstances, what probability should Peon assign, upon being interviewed, to the event that the nickel showed heads?

For this subproblem, the solution seems clear. Let *Interviewed* be the event that "Peon was interviewed (at all)", and let *NickelHeads* be the event "the nickel showed heads", and similarly *DimeHeads*, etc. Then Peon is being asked to assess the probability of *NickelHeads*, conditional on knowing only that the event *Interviewed* occurred. Indeed, since this subproblem involves no amnesia or multiple interviews, *all* that Peon learns is whether or not he is interviewed at all, i.e. whether or not the event *Interviewed* occurs, so it is mathematically clear that *Interviewed* is the event Peon should condition on.

That is, Peon is being asked to compute the conditional probability $P(\mathit{NickelHeads} \,|\, \mathit{Interviewed})$. He would do so as follows:

$$P(\mathit{NickelHeads} \,|\, \mathit{Interviewed}) \;=\; \frac{P(\mathit{NickelHeads} \text{ and } \mathit{Interviewed})}{P(\mathit{Interviewed})}$$

$$=\; \frac{P(\mathit{NickelHeads} \text{ and } \mathit{DimeTails})}{P(\mathit{NickelTails} \text{ or } \mathit{DimeTails})} \;=\; \frac{1/4}{3/4} \;=\; 1/3\,.$$

Thus, for this simple subproblem, the correct probability that Peon should assign during the interview to the event that the nickel showed heads is equal to 1/3. I consider this answer to be correct and clear and unambiguous, following directly from straightforward mathematical laws of conditional probability. I shall now argue that the original Sleeping Beauty problem can essentially be reduced to this simple Peon subproblem.

3. The Original Problem Revisited.

To make use of the above Peon subproblem in analysing the original Sleeping Beauty problem, we add one additional element. We assume that in addition to the previous elements (the nickel, Sleeping Beauty herself, the amnesia-inducing drug, etc.), we also have at our disposal another fair coin, say a dime. We make use of the dime as follows. If the nickel showed tails, then the dime is simply placed so that it shows heads during Beauty's Monday interview, and then repositioned so that it shows tails during Beauty's Tuesday interview. If instead the nickel showed heads (so Beauty will only be interviewed once), then the dime is instead simply flipped once in the usual fashion at the beginning of the experiment, and allowed to show its actual flipped result (either heads or tails, with probability 1/2 each) during the one interview that will take place on Monday. Furthermore, we assume that Beauty is not allowed to see the dime at all, and might not even know of its existence.

Thus, the dime does not in any way affect or control or interfere with any aspect of the original problem. However, we shall see that the dime does permit a precise mathematical analysis of the problem.

We now reason as follows. Call an interview a "heads-interview" if it takes place while the dime shows heads. If the nickel showed tails, then there will certainly be precisely one heads-interview. If the nickel showed heads, then there will be either one or zero head-interviews, with probability 1/2 each. So, the number of heads-interviews behaves just like the total number of interviews in the Peon subproblem.

Now, if Beauty were *told* just before her interview that the dime shows heads (while still undergoing complete amnesia regarding any previous interviews), then she would learn that a heads-interview did indeed occur. This would then put her in precisely the same situation as that of the Peon in the subproblem above. Hence, just like the Peon, Beauty would then assign probability 1/3 that the nickel showed heads. In summary, *if Beauty were told that the dime showed heads*, then the correct answer to the problem would be 1/3.

Similarly, if Beauty is informed (just before her interview) that the dime shows tails, then by identical reasoning, the answer would again be 1/3. In summary, the answer would be 1/3 if Beauty could see the dime, regardless of whether the dime was currently showing heads or tails. We can write this in mathematical terms as

$$P(\mathit{NickelHeads} \,|\, \mathit{DimeHeads}) \;=\; P(\mathit{NickelHeads} \,|\, \mathit{DimeTails}) \;=\; 1/3\,.$$

where now *DimeHeads* is the event that the dime shows heads during the particular interview under consideration, i.e. that the interview was a heads-interview (and similarly *DimeTails*).

In the actual problem, we assume that Beauty *cannot* see the dime. However, we now argue that, as far as probabilities for the nickel are concerned, that fact is irrelevant, and Beauty should still assign the probability 1/3 even if she does not know what the dime shows. To see this, write $P(\mathit{NickelHeads})$ for the overall probability that Beauty should assign to the event that the nickel showed heads upon being interviewed (but now without knowing about the dime). Then it follows by the Law of Total Probability that

$$\begin{aligned} P(\mathit{NickelHeads}) &= P(\mathit{DimeHeads})\, P(\mathit{NickelHeads} \,|\, \mathit{DimeHeads}) \\ &\qquad + P(\mathit{DimeTails})\, P(\mathit{NickelHeads} \,|\, \mathit{DimeTails}) \\ &= P(\mathit{DimeHeads})\,(1/3) + P(\mathit{DimeTails})\,(1/3) \;=\; 1/3\,, \end{aligned}$$

since $P(\mathit{DimeHeads}) + P(\mathit{DimeTails}) = 1$.

Thus, the answer for this version of the problem seems unambiguously and mathematically to be 1/3. And, since the mere existence of the dime (which Beauty cannot see and has no knowledge of) cannot change Beauty's probabilities, I submit that this argument shows unambiguously that the answer to the original Sleeping Beauty problem is also 1/3.

4. Some Related Issues.

While the above completes my main argument, I now consider a few other related issues.

4.1. A slight variant: randomised Sleeping Beauty.

Consider a very slight variant of the original Sleeping Beauty problem. As before, if the nickel is tails we will interview Beauty twice, once on Monday and once on Tuesday (with amnesia). And, as before, if the nickel is heads we will interview Beauty just once. The only modification is that if the nickel is heads, then rather than necessarily interviewing Beauty on Monday, we will first flip another fair coin (say, a dime), and then conduct our (one) interview on Monday if the dime is heads, or on Tuesday if the dime is tails. (We assume, as usual, that Sleeping Beauty cannot tell what day it is.)

For this variant, if Beauty were *told* that her interview was taking place on Monday, then this would reduce precisely to the Peon subproblem above. That is, as far as Monday interviews go, if the nickel showed tails then she would certainly have precisely one, while if the nickel showed heads then she would have one only with probability 1/2 (i.e., only if the dime showed tails), otherwise zero. Furthermore, the fact that the interview is actually taking place on Monday tells her that she did indeed have one Monday interview. Thus, Beauty is in precisely the same situation as the Peon in the above subproblem. So, just as in the subproblem, the correct answer for the probability that the nickel showed heads would be 1/3.

Similarly, if Beauty were told that her interview was taking place on Tuesday, the answer would again be 1/3. (The reasoning is identical to the above, except that the roles of "heads" and "tails" for the dime are interchanged.) In summary, the answer would be 1/3 if she knew which day it was, regardless of whether that day were Monday or Tuesday.

Now, in the actual problem, Beauty is *not* told which day it is. However, by the Law of Total Probability just as before, it follows that since she would have assigned probability 1/3 upon being told either that is Monday or that it is Tuesday, she should still assign probability 1/3 even if she does not know which day it is. Thus, the answer for this variant of the problem again seems unambiguously and mathematically to be 1/3.

Now, it seems clear that this variant is probabilistically equivalent to the original Sleeping Beauty problem, since in the original problem it is not relevant whether the one interview (if the nickel shows heads) takes place on Monday or Tuesday. So, this provides another (similar) argument for why the answer to the original problem is 1/3.

4.2. Yet another variant: sleeping twins.

Consider the following variant of the Sleeping Beauty problem. Suppose there are two twins, named Beauty1 and Beauty2. We put them both to sleep (in separate, soundproof rooms), and flip a fair nickel. If the nickel shows tails, we wake and interview each of them (separately). If the nickel shows heads, we flip a dime. If the dime shows tails we interview Beauty1 only, while if the dime shows heads we interview Beauty2 only. What probability should each of them assign, upon being interviewed, to the event that the nickel showed heads?

It is clear that in this variant, the situation for Beauty1 is precisely the same as that of the Peon in the above subproblem. Hence, as in that subproblem, Beauty1 should assign probability 1/3 to the nickel showing heads. Similarly, Beauty2 should also assign probability 1/3 to the nickel showing heads.

On the other hand, if we regard Beauty1 and Beauty2 as a "unit", then together they behave (probabilistically speaking) just like Sleeping Beauty in the original problem. Indeed, the total number of times that Beauty1 and Beauty2 will be interviewed is two if the nickel is tails, and one if the nickel is heads. So, since each of Beauty1 and Beauty2 should assign the probability 1/3, this suggests that Sleeping Beauty in the original problem should also assign probability 1/3. Indeed, this argument is very similar to, and perhaps more intuitive than, the argument given in Section 3 above. However, it is not completely definitive, due to the possible confusion over conditioning on the same person being interviewed twice (in the original problem), versus two different people each being interviewed once (in this variant).

4.3. A simple argument why 1/2 must be wrong.

Another mathematical insight into the original Sleeping Beauty problem can be gained by conditioning on the day of the interview, i.e. by considering how the probabilities would change if Beauty knew which day it was. Recall that, in the original problem, Beauty is interviewed on both Monday and Tuesday if the nickel showed tails, but is interviewed on Monday alone if the nickel showed heads.

Suppose first that Beauty is informed that her interview is taking place on Monday. Then, since precisely one interview would be conducted on Monday regardless of whether the nickel showed heads or tails, she should at that point assign equal probabilities to the nickel showing heads or tails. In other words, we must have $P(\mathit{NickelHeads} \,|\, \mathit{Monday}) = 1/2$, where *Monday* is the event that "the interview is taking place on Monday".

On the other hand, suppose Beauty is informed that her interview is taking place on Tuesday. Then, since it is *impossible* to have an interview on Tuesday if the nickel shows heads, she should at that point assign probability zero to the nickel showing heads. That is, we must have $P(\mathit{NickelHeads} \,|\, \mathit{Tuesday}) = 0$.

It then follows, again by the Law of Total Probability, that

$$
\begin{aligned}
P(\mathit{NickelHeads}) &= P(\mathit{Monday})\, P(\mathit{NickelHeads} \,|\, \mathit{Monday}) + P(\mathit{Tuesday})\, P(\mathit{NickelHeads} \,|\, \mathit{Tuesday}) \\
&= P(\mathit{Monday})\,(1/2) + P(\mathit{Tuesday})\,(0) \;=\; P(\mathit{Monday})\,/\,2\,.
\end{aligned}
$$

Now, it is not clear what value Beauty should assign to $P(\mathit{Monday})$, the probability (without any additional knowledge) that her interview is in fact taking place on Monday. Is it 1/2, since she could be interviewed on either day? Or 2/3, since two of the three possible interview situations (heads-Monday, tails-Monday, tails-Tuesday) involve Monday? Or 3/4, reasoning that the probabilities of those three possible interview situations are respectively 1/2, 1/4, and 1/4, and $1/2 + 1/4 = 3/4$?

In any case, since *sometimes* interviews will take place on Tuesday, we must have $P(Tuesday) > 0$, whence $P(Monday) = 1 - P(Tuesday) < 1$, whence $P(NickelHeads) = P(Monday) / 2 < 1/2$. Hence, this argument allows us to see directly that the answer 1/2 cannot be correct.

(Of course, once we agree that $P(NickelHeads) = 1/3$ is the correct answer to the original problem, then working backwards we can conclude that $P(Monday) = 2/3$.)

4.4. Generalisation to other numbers and probabilities.

Once we accept the above reasoning, then it can also be applied to various generalisations of the original problem.

For example, if Beauty will instead be interviewed n times (with amnesia each time) if the nickel showed tails, but just once if the nickel showed heads, then it follows (by replacing the dime by an n-sided die) that the answer becomes $1/(n+1)$. The original problem corresponds to $n = 2$.

Or, if the nickel actually was not a fair coin but instead had *a priori* probability q of coming up heads (and $1-q$ of coming up tails), then the answer would become $(q/2)/(q/2 + (1-q)) = q/(2-q)$. The original problem corresponds to $q = 1/2$.

If we combine both of the above modifications simultaneously, then the answer would become $q/(n-(n-1)q)$. The original problem corresponds to the values $n = 2$ and $q = 1/2$.

Many other similar variations can be solved in a similar fashion.

5. Relation to Other Probability Puzzles.

The Sleeping Beauty problem is reminiscent of certain other well-known probability puzzles in which a conditional probability at first appears to be 1/2, but upon reflection is actually 1/3. We review two such puzzles here, and then consider their relation to the Sleeping Beauty problem.

5.1. Bertrand's Box.

The *Bertrand's Box* probability puzzle was proposed by Joseph Bertrand in 1889. It is sometimes called (in an equivalent "drawers" version) the *Three Desk Problem* or *Three Drawers Problem*, or (in an equivalent "cards" version, see e.g. Rosenthal 2006) the *Three-Card Thriller* or *Three-Card Swindle*. It can be stated as follows:

> There are three boxes. Box #1 contains two gold coins, Box #2 contains two silver coins, and Box #3 contains one gold and one silver coin. One box is chosen uniformly at random, and one coin is chosen uniformly at random from that box. Suppose the chosen coin is gold. What is the probability that the chosen box was Box #3?

In this problem, many people will reason (correctly) that observing the gold coin immediately eliminates Box #2. They will then reason that Boxes #1 and #3 must (still) be equally likely, so the probability of Box #3 must be 1/2.

However, what the question is really asking is for the *conditional* probability $P(\mathit{Box3} \,|\, \mathit{CoinGold})$. This can easily be computed (e.g. Rosenthal, 2006) as:

$$P(\mathit{Box3} \,|\, \mathit{CoinGold}) \;=\; \frac{P(\mathit{Box3},\; \mathit{CoinGold})}{P(\mathit{CoinGold})} \;=\; \frac{P(\mathit{Box3})\; P(\mathit{CoinGold} \,|\, \mathit{Box3})}{P(\mathit{CoinGold})}$$

$$=\; \frac{(1/3)(1/2)}{1/2} \;=\; 1/3, .$$

Hence, the answer is 1/3. (If that seems counter-intuitive, note that conditional on *CoinGold*, the chosen coin was equally likely to be any of the three gold coins available, only one of which is in Box #3.)

5.2. Monty Hall Problem.

Another probability puzzle is the *Monty Hall problem* (vos Savant, 1990), which may be stated as follows:

> A car is equally likely to be behind any one of three doors. You select one of the three doors (say, Door #1). The Host then reveals one non-selected door (say, Door #3) which does *not* contain the car. At this point, you choose whether to stick with your original choice (i.e. Door #1), or switch to the remaining door (i.e. Door #2). What is the probability that you will win the car if you stick with your original choice?

Most people, upon first hearing this problem, believe (vociferously!) that the car is equally likely to be behind either of the two unopened doors, so the probability of winning is 1/2 regardless of whether you stick or switch. However, in fact the probabilities of winning are 1/3 if you stick, and 2/3 if you switch.

From a conditional probability point of view, this error arises because most people intuitively compute the wrong quantity. Specifically, most people compute $P(\mathit{Car1} \,|\, \mathit{CarNot3})$, i.e. the conditional probability that the car is behind Door #1 given that it is not behind Door #3. This probability is easily computed (and intuitively seen) to be 1/2. And, this would indeed be the correct answer if we assumed that the Host just *happened* to reveal Door #3, by accident, and it just *happened* not to contain a car. (In Rosenthal 2008, this variant is called the *Monty Fall problem*.)

However, in the original Monty Hall problem, what we actually want to compute is $P(\mathit{Car1} \,|\, \mathit{Host3})$, i.e. the probability that the car is behind Door #1 given that the Host *elected* to open Door #3. This is related to the Host's motivations, so to compute this properly requires certain assumptions (which are implicit, though not explicit, in the original problem). Namely, we assume that the Host knows where the car is, and will always elect to open some door which is not the door you originally selected and which does not contain a car. Furthermore, we assume the Host will choose randomly (with probability 1/2 each) if there are two different such doors available.

With these assumptions, we can compute (e.g. Rosenthal, 2006, 2008):

$$P(\mathit{Car1} \,|\, \mathit{Host3}) \;=\; \frac{P(\mathit{Car1},\ \mathit{Host3})}{P(\mathit{Host3})} \;=\; \frac{P(\mathit{Car1})\ P(\mathit{Host3} \,|\, \mathit{Car1})}{P(\mathit{Host3})} \;=\; \frac{(1/3)(1/2)}{1/2} \;=\; 1/3, .$$

Hence, the answer is 1/3. (If that seems counter-intuitive, note that the strategy of sticking will only succeed if your original guess happened to be correct, which had probability 1/3. And since we *knew* the Host was going to open *some* door not containing the car, observing this doesn't change the probability 1/3 that we were right in the first place.)

5.3. Comparison to Sleeping Beauty.

Each of these two puzzles, like the Sleeping Beauty problem, involves computing a conditional probability which may at first seem to equal 1/2, but is in fact equal to 1/3. However, there are some subtle differences between the three scenarios.

In Bertram's Box, the erroneous answer 1/2 arises purely from a misunderstanding of conditional probability. Once the rules of conditional probability are carefully brought to bear on the problem, the answer is clear and unambiguous.

In the Monty Hall problem, there is some confusion about which conditional probability should actually be computed, and also about the implicit assumptions concerning the Host's behaviour. However, once these points are clarified, then again the rules of conditional probability can be carefully brought to bear on the problem, again providing a clear and unambiguous answer.

By contrast, with the Sleeping Beauty problem, even upon careful reflection, it remains unclear how to mathematically formulate the notion of "conditional on currently being interviewed". So, without some sort of reformulation, it seems that conditional probability cannot be brought to bear directly on Sleeping Beauty.

However, the contribution of this paper is to present a slight reformulation of the Sleeping Beauty problem (by introducing a dime, while arguing that the dime does not affect the final answer), and to show that this slight reformulation can then be analysed unambiguously by the mathematics of conditional probability.

In summary, Bertram's Box and the Monty Hall problem, when formulated clearly, provide unambiguous exercises in probability theory. By contrast, the Sleeping Beauty problem necessarily involves some sort of reformulation or philosophical analysis, though this paper attempts to keep such matters to an absolutely minimum.

6. Final Discussion.

As mentioned in the Introduction, lots of articles have previously been published about the Sleeping Beauty problem, including many which argue (as I do) that the answer is 1/3. Thus, I see the main contribution of this paper not as presenting a new result, but rather as providing a simple, mathematically-based justification for why 1/3 is correct.

Specifically, my main argument requires only (i) the mathematics of conditional probability, e.g. $P(A \mid B) = \frac{P(A \text{ and } B)}{P(B)}$, and (ii) the "axiom" that people with identical relevant information will assign probabilities identically. My argument is then quite brief (Sections 2 and 3 above), and can be summarised as:

- By (i), the Peon will assign probability 1/3 to the nickel being heads.
- Hence, by (ii), after Beauty is informed that the dime shows heads (or, similarly, tails), she will assign probability 1/3 to the nickel being heads.

- Hence, by (i), Beauty will assign probability 1/3 to the nickel being heads even before she is informed what the dime shows.
- Hence, by (ii), Beauty will assign probability 1/3 to the nickel being heads even if the dime does not exist.

So, it seems that any readers who accepts mathematical probability theory together with the above rather obvious "axiom", should be convinced that 1/3 is the correct answer.

Of course, some such readers may have already been convinced by various previously published arguments, and might find the current paper superfluous. Indeed, some of the previous arguments already contained some elements of conditional probability somewhat related to those herein (e.g. the "Technicolor Beauty" variant described in Section 2.5 of the long paper by Titelbaum, 2008, has some points of commonality). And, conversely, some readers may continue to believe that 1/2 is the correct answer, for various philosophical reasons, even after reading my argument. So, I do not expect the current paper to completely resolve the controversy.

Despite these caveats, I hope and believe that there is merit in providing a simple, short, direct argument that 1/3 is the correct answer, using solid mathematical foundations with few assumptions and little philosophical ambiguity.

Acknowledgements: I am very grateful to Gary Malinas and Calvin Normore for discussing these issues with me, and to an anonymous referee for a very helpful report.

References:

Arntzenius, F. (January 2002), Reflections on Sleeping Beauty. *Analysis* **62**, 53–62.

Bostrom, N. (July 2007), Sleeping Beauty and self-location: a hybrid model. *Synthese* **157**, 59–78.

Dorr, C. (October 2002), Sleeping Beauty: in defence of Elga. *Analysis* **62**, 292–296.

Elga, A. (April 2000), Self-locating belief and the Sleeping Beauty problem. *Analysis* **60**, 143–147.

Horgan, T. (January 2004), Sleeping Beauty awakened: New odds at the dawn of the new day. Analysis **64**, 10-21.

Lewis, D. (July 2001), Sleeping Beauty: reply to Elga. *Analysis* **61**, 171–176.

Monton, B. (January 2002), Sleeping Beauty and the forgetful Bayesian. *Analysis* **62**, 47–53.

Neal, R. M. (2006), Puzzles of anthropic reasoning resolved using full non-indexical conditioning. Technical Report No. 0607, Dept. of Statistics, University of Toronto.

Piccione, M. and Rubinstein, A. (July 1997), On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* **20(1)**, 3–24.

Pust, J. (January 2008), Horgan on Sleeping Beauty. Synthese **160**, 97-101

Rosenthal, J.S. (2006), Struck by Lightning: The Curious World of Probabilities. Joseph Henry Press, Washington, D.C.

Rosenthal, J.S. (September 2008), Monty Hall, Monty Fall, Monty Crawl. *Math Horizons*, 5–7.

Titelbaum, M.G. (2008), The Relevance of Self-Locating Belief. *Philosophical Review* **117(4)**, 555–605.

vos Savant, M. (Sept. 9, 1990), *Ask Marilyn* column. Parade Magazine, p. 22.

Weintraub, R. (January 2004), Sleeping Beauty: a simple solution. Analysis **64**, 8-10.