
The Sleeping Beauty Controversy

Peter Winkler

Abstract. In 2000, Adam Elga [12] posed the following problem:

Some researchers are going to put you to sleep. During the two days that your sleep will last, they will briefly wake you up either once or twice, depending on the toss of a fair coin (Heads: once; Tails: twice). After each waking, they will put you back to sleep with a drug that makes you forget that waking. When you are first awakened, to what degree ought you believe that the outcome of the coin toss is Heads?

This may seem like a simple question about conditional probability, but 100 or so articles (including thousands of pages in major philosophy journals) have been devoted to it. Herein is an attempt to summarize the main arguments and to determine what, if anything, has been learned.

1. INTRODUCTION. Consider the following problem, restated here in the third person (Sleeping Beauty, or SB):

Sleeping Beauty agrees to the following experiment. On Sunday, she is put to sleep, and a fair coin is flipped. If it comes up Heads, she is awakened on Monday morning; if Tails, she is awakened on Monday morning and again on Tuesday morning. In all cases, she is not told the day of the week, is put back to sleep shortly after, and will have no memory of any Monday or Tuesday awakenings.

When Sleeping Beauty is awakened on Monday or Tuesday, what—to her—is the probability that the coin came up Heads?

I choose to use the term “probability” above since mathematicians are accustomed to assigning numerical values to that quantity, as computed or estimated by a particular person about a particular event in a particular circumstance. Below, I will assume that this probability is the degree of “credence” that the person *should have*.

The issue is not that the Sleeping Beauty problem is undecidable (that is, not solvable from the axioms of set theory); indeed, it seems that nearly everyone discussing the problem has a strong opinion about its answer. The arguers fall into camps and subcamps, some claiming that the answer comes down to the nature of probability, the meaning of consciousness, evidential versus causal decision theory, one-world versus many-world quantum mechanics, conditioning versus updating—or exactly how the problem is phrased.

Those who believe that the answer is $1/2$ have been dubbed “halfers,” while those who go with $1/3$ are called “thirders.” If you believe either answer could be correct—depending, perhaps, on interpretation or phrasing—you are a “dualist”; if you think there is no correct answer because the problem cannot be well posed, you’ll be called an “objector.” There are subcamps (see below), and a few folks who say they simply don’t know. Perhaps there should even be a category for those who don’t care. But, hey, you’ve read this far, right?

<http://dx.doi.org/10.4169/amer.math.monthly.124.7.579>
MSC: Primary 60A99, Secondary 00A30; 00A05

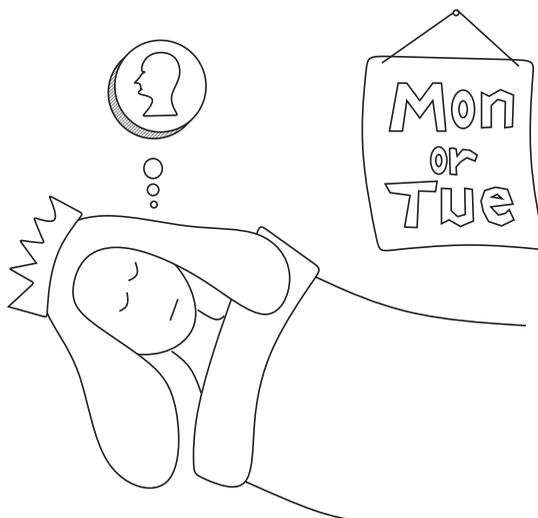


Figure 1. Sleeping Beauty innocently dreaming about the coin.

Below, I will present the main arguments of the halfers and thirders, temporarily omitting references in order to maintain brevity. Then I will review the problem's history and go back over the arguments in reverse order, looking for enlightenment.

Information. Say the halfers: Before SB is put to sleep on Sunday, her credence that the fair coin will come up Heads is inarguably $1/2$. She knows she will be awakened, so when she inevitably is, *she has no new information*, and therefore, her credence in Heads cannot have changed.

Reflection. Similarly: On Sunday, the thirder SB *knows* that, on Monday, she will give credence $1/3$ to Heads. But then she should already have credence $1/3$ on Sunday, which is absurd.

Repetition. Say the thirders: Repeat the experiment 100 times; then SB will be awakened about 150 times, 50 of them to Heads. So SB's probability of Heads, on awakening, must be $1/3$.

Gambling. Say the thirders: Ask SB, upon each awakening, if she's willing to have \$3 deducted from her bank account if the coin landed Heads, provided that \$2 is *added* to her account if the coin landed Tails. (She understands that if it's Tuesday morning and she accepted the bet on Monday, her holdings may already have changed—but that should have no influence on today's decision.)

As a thirder, SB should accept the bet: Her expectation is $\frac{1}{3}(-\$3) + \frac{2}{3}(+\$2) > 0$. And she's right to do so: Over the course of the experiment, she ends up \$4 ahead if the coin landed Tails and only \$3 behind if it landed Heads. A corresponding calculation by the halfer SB would lead her to refuse the bet.

Symmetry. Suppose that 15 minutes after her Monday awakening, SB will be told what day it is. If she hears it's Monday, her probability of Heads is $1/2$. Suppose instead SB will be told the result of the coin flip. Then, if it's Tails, her probability that it's Monday is $1/2$. The first implies $\mathbb{P}(\text{Heads and Monday}) = \mathbb{P}(\text{Tails and Monday})$; the second, that $\mathbb{P}(\text{Tails and Monday}) = \mathbb{P}(\text{Tails and Tuesday})$. These three probabilities represent exhaustive and mutually exclusive events, thus each is equal to $1/3$.

2. HISTORY. The above symmetry argument appeared in Elga's seminal paper [12]. Elga extracted the Sleeping Beauty problem (named by Robert Stalnaker) from Example 5 of Michele Piccione et al. [31], one of many papers in a volume of *Games and Economic Behavior* dedicated to a decision-theoretic problem known as "The Absent-Minded Driver."¹ Thus began a storm of arguments, papers, and blog comments, drawing in philosophers, mathematicians, and even physicists, then (seemingly) everyone.

Philosophers are after much bigger game than the Sleeping Beauty problem itself. How does one determine the credence that should be given to a particular proposition in given circumstances, and how should that credence be updated with new information or the passage of time? Sleeping Beauty is a demanding test for any theory that addresses these questions, incorporating loss of consciousness, loss of memory, and absence of time indication. Many philosophers acknowledge the success of Kolmogorov's probability axioms [24], but question whether they are equipped for handling propositions like "I am thirsty" whose truth-values change with time and perspective.

The debate in the philosophy community is thus concerned not just with the "correct" answer to Sleeping Beauty, but with which (if any) answer is implied by a given theory. Some, like Michael Titelbaum (see, e.g., [41, 42, 43] but not [40]) doubt that there is a good reason to choose any current framework over the rest in an effort to settle the issue. But even if you agree with this sentiment, you might ask which answer you *want* your theory to provide, and here the evidence (to me) is that the thirder position has emerged as the dominant view. This conclusion is not arrived at by counting papers as if they were votes since, of course, papers are supposed to present new ideas. Many published papers are attacks on the thirder position from various other camps, many others rebuttals. But anti-thirders seem to see themselves as fighting the establishment; concedes halfer Joel Pust [34]: "Most of those writing on the SB problem have argued that one-third is the correct answer."

First to challenge Elga was David Lewis [25], rebutted by Cian Dorr [10]. Frank Arntzenius started as an objector [1] but became a thirder [2]. Rachel Briggs [5] advanced the idea that causal² decision theorists (like her) should be thirders, and evidential decision theorists should be halfers, refuted by Vincent Conitzer [7]. Peter J. Lewis [26] argued that the many-world view of quantum physics implies halfism, rebutted by Alistair Wilson [47]. John Pittard, a halfer himself, argues in [33] that halfers must endorse robust perspectivalism.³ More challenges and rebuttals are cited below in connection with particular arguments.

There have been many attempts to reconcile thirders and halfers by saying both are right (see Jacob Ross [36], who reaches a "rational dilemma," or Berry Groisman [15]), or both wrong (Nick Bostrom [4]); there is even (see Namjoong Kim [23]) a camp (the "lessers") who hold that the answer is less than 1/2 but maybe not 1/3. Jessi Cisewski et al. [6] claim that both sides, and all values between, are supportable depending

¹In its simplest form, the absent-minded driver wants to take the second highway exit to get home but can't distinguish the second exit from the first and knows that he will not remember whether he already passed an exit. Preferring to miss both exits rather than get off at the first, he reluctantly decides to stay on the highway, but when he gets to an exit, he reconsiders. In considering randomized algorithms for the driver, most writers assumed what would be the thirder view—causing, according, e.g., to Wolfgang Schwarz [37], paradoxical conclusions.

²More about causal versus evidential decision theory will be found in the "Gambling Revisited" section below.

³Perspectivalism says that two agents can rationally disagree about a proposition even though each gives the other's argument the same weight as her own. Pittard uses "robust" to emphasize that in his SB variation, the arguments do not merely have equal weight; they are identical.

on whether one believes SB’s total knowledge—including degree of indigestion—is assumed to be exactly the same Monday and Tuesday.⁴ Pradeep Mutalik [30], who writes for the excellent online magazine *QUANTA*, believes the answer depends on whether the question asks about “the coin associated with this experiment” or “the coin associated with this awakening.” But most mathematicians would not, I think, accept the notion that equivalent events can have different probabilities.

To see why the thirders are leading, I will revisit Elga’s symmetry argument, which looks a lot like a proof of the thirder position. Is it? We mathematicians like to think that an argument either is or isn’t a proof, but this applies only to formal proofs (which, like honest politicians, are much talked about but rarely seen). In real life, much more is demanded of proofs of theorems that have counterintuitive consequences. I will therefore return in reverse order to the earlier arguments, ending with the halfers’ compelling information argument. Perhaps the consequences of the thirder position can be made more comfortable.

3. REVISITS.

Symmetry revisited. Elga’s symmetry argument retains the honor of being the most popular target of nonthirders and has been attacked on every conceivable front. Here is Elga’s argument in symbolic form.

At a moment of SB’s awakening there are three possible relevant states: Heads (H) and Monday (M); Tails (T) and Tuesday (U); T and M. If SB is to be told (on every awakening) the day of the week and is now told “Monday,” then, by coin-flip symmetry,

$$\mathbb{P}(H|M) = \mathbb{P}(T|M).$$

If SB is to be told (on every awakening) the state of the coin and is now told “Tails,” then, this time by indistinguishability of the Monday and Tuesday awakenings,

$$\mathbb{P}(M|T) = \mathbb{P}(U|T).$$

Therefore,

$$\mathbb{P}(H \wedge M) = \mathbb{P}(H|M) \cdot \mathbb{P}(M) = \mathbb{P}(T|M) \cdot \mathbb{P}(M) = \mathbb{P}(T \wedge M)$$

and

$$\mathbb{P}(T \wedge M) = \mathbb{P}(M|T) \cdot \mathbb{P}(T) = \mathbb{P}(U|T) \cdot \mathbb{P}(T) = \mathbb{P}(T \wedge U),$$

so

$$\mathbb{P}(H \wedge M) = \mathbb{P}(T \wedge M) = \mathbb{P}(T \wedge U).$$

Since these events are mutually exclusive and exhaustive, their probabilities sum to 1, so each has probability 1/3 and in particular

$$\mathbb{P}(H) = \mathbb{P}(H \wedge M) = 1/3.$$

⁴An issue with [6] is that its equation (1) gives the probability that an event is witnessed at least once during the experiment, but what is required for SB’s conditioning is that the event is witnessed at time t .

Some halfers, including Lewis [25], dispute the claim that $\mathbb{P}(H|M) = 1/2$ —a hard position to maintain when you consider that the experimenters don't need to flip the coin until Monday night. (How can SB's credence in Heads be other than 1/2 if she knows the coin hasn't been flipped yet?) Others (called "double-halfers") concede that $\mathbb{P}(H|M) = 1/2$, maintaining that SB's credence in Heads doesn't change when she hears it's Monday. How, then, do they deal with the math? Mikael Cozic [9] suggests that conditioning is not the right way to modify SB's credence; Ioannis Mariolis [27] claims there are two kinds of "it is Monday" events, one of which he calls "Monday*"; Joseph Halpern [16] claims a "difference between the probability of heads conditional on it being Monday versus the probability of heads conditional on learning that it is Monday";⁵ Patrick Hawley [17] does not accept that SB should be uncertain of the day!⁶ Roger White [46] doesn't know what's wrong with Elga's argument but claims that a natural generalization of it has an unacceptable consequence. (Terry Horgan [20] is happy with the generalization but argues that its consequence is not only acceptable but demonstrably correct.) As far as I can tell, every published attack on Elga's argument has been riposted, except perhaps for the just-published [6], whose quarrel with Elga is that he *assumes* that when SB is awakened, the events "it is Monday" and "it is Tuesday" are exclusive and exhaustive.

To help clarify the main camps, here's a rephrasing (some might disagree) of Sleeping Beauty: Alice, Bob, and Charlie have each taken a new sleeping pill. In hospital experiments, half the subjects slept through the night (as intended by the pill's creators), but the other half woke up once in the middle of the night then returned to sleep and woke up in the morning with no memory of the night awakening.

Alice wakes up in the middle of the night, and her credence that the pill has worked drops to zero (no argument there).

Bob wakes up in the morning; his credence that the pill worked remains at 1/2 (thirders and double-halfers would agree, but the Lewisians would not).

Charlie, who has blackout shades in his bedroom, wakes up not knowing whether it is morning. According to the thirders, his credence in the efficacy of the pill is 1/3 until he raises the shades, at which point it rises to 1/2 or drops to zero. If he desperately wants the pill to have worked, you might think Charlie would be happy to see the morning sun—but would he? The double-halfers believe that the sun does not change Charlie's credence that the pill worked.

Further symmetry arguments have been advanced. Modulo details, Dorr [10] and Arntzenius [2] suppose that regardless of the coin flip, SB will be awakened on both days, but in the event of Heads, SB will be told "Heads and Tuesday" 15 minutes after her Tuesday awakening. Then, for the first 15 minutes that SB is awake, "Heads and Monday," "Heads and Tuesday," "Tails and Monday," and "Tails and Tuesday" are equiprobable by symmetry. After 15 minutes, when SB doesn't get the "Heads and Tuesday" signal, she eliminates that option, and the other three possibilities must remain equiprobable.

Other mathematical arguments, including those in Jeffrey Rosenthal's [35] in *The Mathematical Intelligencer* and Titelbaum's 2008 paper [41], have been offered. These as well are plausible to a mathematician—except, of course, for possible unintuitive consequences of the thirder position, which will be addressed below.

⁵ which would indeed be the case, were SB not told in advance, in Elga's argument, that she would learn what day it is.

⁶ It is not my intention to make fun of philosophy here, quite the contrary: I love it that philosophers question *everything*. Somebody has to!

Gambling revisited. Both halfers and thirders have attempted to employ “Dutch books” to discredit the opposition; some of their arguments can be found in Christopher Hitchcock [18] (great title). A Dutch book is a sequence of “fair” bets with a guaranteed negative outcome—of course, this ought not to be possible, so the argument is that, if a Dutch book can be made, the probabilities upon which the fairness of the bets relies must not be correct. (Yes, regardless of SB’s credences, she can always work out her best strategy for the whole experiment and bet accordingly. But her credences ought to be reliable guides.)

Some Dutch book (and other decision-theoretic) arguments present situations in which SB’s best choice at a given awakening might depend on her decision at a previous or forthcoming awakening. In such cases, causal and evidential decision theory might differ concerning her rational action. A causal decision theorist believes SB’s actions should be based only on what they cause (in particular, *not* on her decisions at other awakenings) while an evidential decision theorist is permitted to use, e.g., her judgment that her decisions at all awakenings will probably be the same. Readers familiar with the Prisoner’s Dilemma (see, e.g., Steven Tadelis [39]) might imagine a case where the prisoners are identical twins who don’t necessarily give a fig for one another but have historically always made the same decisions in identical situations. If the twins are evidential decision theorists, they will escape the dilemma, each reasoning that whatever he does, his twin will probably do the same. If they are causal decision theorists, too bad!

Conitzer [8] regards as suspect any decision-theoretic argument in which SB’s decisions are not “additive,” therefore independent of decisions at other awakenings. Conitzer shows that the thirder SB plays additive games optimally but does not believe that this settles the SB problem conclusively.

Repetition revisited. To the argument that $1/3$ of awakenings are Heads, the halfers reply that one should count experiments—that is, coin flips—not awakenings. Being asked twice upon Tails, they say, does not change the result of the fair coin. But, somehow, the halfer’s retort seems to lose some force if SB is asked not about the coin but about her degree of credence that it is Tuesday. Of course, the halfer answer is $1/4$ while the thirders claim $1/3$; so what? But when the question is about the awakening, not (directly) about the coin, it seems more natural to count awakenings.

Reflection revisited. The idea that if you know that tomorrow you will think “X” then you should already think “X” today is known to many philosophers as Bastiaan van Fraassen’s reflection principle [13, 14]. The reflection principle is for the most part a special case of the optional stopping theorem (see, e.g., Richard Durrett [11]), which implies that, if whenever a learning process (more generally, a martingale) is stopped the probability of a given event is the same, then the event’s *a priori* probability must be that same value. But the stopping algorithm must be implementable. This will not be the case, even when the stopping time is fixed on the clock, if SB doesn’t know what time it is. (See, e.g., [38]—another great title.)

In fact, failure of Van Fraassen’s reflection principle in such a case is a fairly ordinary occurrence, and no memory loss is required for a demonstration. In Fred’s town, if school is to be suspended on account of snow, a loud siren blast is heard at 7:00 a.m. sharp. Fred wakes up much earlier, estimating a probability of $1/2$ that a snow day will be declared but has no watch or clock. As time passes and he doesn’t hear the blast, his estimate of the snow day probability will go down, reaching perhaps $1/3$ (that magic number!⁷) at 6:59 a.m. Of course, this number will suddenly jump to 1 at

⁷The argument for $1/3$ is that at 6:59 Fred’s *a priori* probability that 7:00 a.m. has passed by is (almost)

7:00 a.m. if it is a snow day, otherwise it will continue decreasing, reaching 0 when Fred is sure it's past 7:00 a.m. But the point is, Fred *knows* that his estimate of the snow day probability will be lower at 6:59 than it is now.

The optional stopping time theorem is not contradicted because, for Fred, "6:59" is not a legitimate stopping time. If instead Fred considers his snow-day-probability-estimate at the point when it becomes light enough to see his dresser—a genuine stopping time—he finds that its expected value is $1/2$.

Similarly, "Monday" is not a legitimate stopping time for SB unless she is to be told the day of the week. When she is told it's Monday, the theorem applies and we correctly conclude that her credence in Heads remains at $1/2$.

Information revisited. How does a thirder handle the argument that, upon awakening, SB has no new information to justify changing her Sunday appraisal? Elga himself, as well as Hitchcock, Monton [29], and Vaidman and Saunders [44], concede the lack of information but believe SB's credence in Heads nonetheless changes. Robert Aumann, Sergiu Hart, and Motty Perry [3], giving a thirder argument years before Elga [12] in connection with *The Absent-Minded Driver*, take this view as well.

Other thirders, especially lately, have argued effectively that to be conscious at a given moment, even if you don't "know what time it is," can constitute genuine information, justifying updating your degree of credence. Included in this list are Arntzenius [2] (after his conversion to thirdism), Dorr [10], Horgan [19], Karlander and Spectre [21], and Weintraub [45].

There are two obvious objections to this contention. One is that, if you weren't conscious, you wouldn't have the information. True, but so what? The *Sleeping Beauty* problem itself provides an example: If SB wakes up on Tuesday and is told the day, she undoubtedly has information and can use it to conclude Tails. Yet, if she is not awakened on Tuesday, she may never know that the coin came up Heads. Similarly, Alice (in the sleeping pill example) learns upon awakening in the middle of the night that the sleeping pill did not work but would not have received the contrary information.

More persuasive is the idea that for your consciousness to provide information, you must know *when* you are conscious. But this calls into question the arbitrariness of how, as well as whether, time is measured. Must a moment be labeled to be a moment? Suppose that in SB's experimental bedroom is an LED device that tells the number of days since the Chicxulub impact. She awakens and reads the number 24,120,373,498. She has no idea what day of the week that represents, but she does know that she is conscious on that day and didn't know before that she would be.⁸

Seen this way, it would be surprising if SB got no information when awakened. Once it is conceded that being conscious at a time when SB might not have been conscious can justify modifying a credence, it follows that being conscious at what *might* be a time when she might not have been conscious also can justify modifying a credence.

$1/2$, now reduced to $1/3$ by the condition of the siren not having gone off yet. Thus "before 7:00 a.m., snow day," "before 7:00 a.m., no snow day," and "after 7:00 a.m., no snow day" are equiprobable. Sounds like SB, no? But (as Arntzenius points out in [2]) Fred can't know what his *a priori* time probability distribution will be at a given time, otherwise he could use that information to tell time!

⁸At this point in the article, readers should not be surprised to learn there are frameworks—see, e.g., [6, 16, 28]—in which the presence of any time-measuring device changes the answer to the *Sleeping Beauty* problem. Others, e.g., Kierland et al. [22], say that unknown time is not really the issue to begin with; the Monday/Tuesday factor can be replaced by a random signal, such as the color of the room in which SB is awakened.

4. CONCLUSIONS. Sleeping Beauty is indeed a beauty of a problem, and I am under no illusions that controversy about its solution will ever entirely disappear. But the extensive thought and discussion devoted to Sleeping Beauty has not been in vain; the literature suggests, at least to this writer, the following points.

- The Sleeping Beauty problem touches many fascinating issues in philosophy but, to the extent that there is agreement about what is asked, is also a mathematical question to which many think the straightforward answer is $1/3$.
- Some consequences of the $1/3$ answer appear surprising at first, but upon scrutiny, seem (for some) increasingly intuitive. In particular, being conscious at a given moment may constitute legitimate information, even if—and in some cases, especially if—the moment’s time label is not known.
- Kolmogorov’s axioms—in particular, their treatment of conditioning—have held up quite well and together with their implications (e.g., the optional stopping time theorem) may indeed have gained some admirers. Andrey Nikolayevich, were he conscious, would be justly proud. But no comprehensive theory of credence and how it is updated has been agreed upon.

Finally, there is ample confirmation here that philosophers and mathematicians have a lot to gain by talking to one another. The former, for example, are reminded that replacing a principle by a theorem can help achieve clarity, the latter that the path from words to numbers can have hidden twists.

ACKNOWLEDGMENTS. The author is greatly indebted to stimulating email conversations with Spencer Backman, Jeff Friedman, Jon Hart, Doug McIlroy, Sam Miner, Jo Ellis-Monaghan, Pradeep Mutalik, Jay Pantone, Tom Zaslavsky, Olga Zhaxybayeva, and especially the very patient Mike Titelbaum. Of course, none of the above can be assumed to endorse this article. Research supported by NSF grant DMS-0901475.

REFERENCES

1. F. Arntzenius, Reflections on Sleeping Beauty, *Analysis* **62** no. 1 (2002) 53–62.
2. ———, Some problems for conditionalization and reflection, *J. Philosophy* **199** no. 7 (2003) 356–370.
3. R. J. Aumann, S. Hart, M. Perry, The forgetful passenger, *Games Econom. Behav.* **20** (1997) 117–120.
4. N. Bostrom, Sleeping Beauty and self-location: A hybrid model, *Synthese* **157** no. 1 (2007) 59–78.
5. R. Briggs, Putting a value on Beauty, *Oxford Studies in Epistemology* **3** (2010) 3–34.
6. J. Cisewski, J. B. Kadane, M. J. Schervish, T. Seidenfeld, R. Stern, Sleeping Beauty’s credences, *Philos. Sci.* **83** no. 3 (2016) 324–347.
7. V. Conitzer, A Dutch book against Sleeping Beauties who are evidential decision theorists, *Synthese* **192** no. 9 (2015) 2887–2899.
8. ———, Can rational choice guide us to correct *de se* beliefs? *Synthese* **192** no. 9 (2015) 4107–4119.
9. M. Cozic, Imaging and Sleeping Beauty: A case for double-halfers, *Internat. J. Approx. Reason.* **52** (2011) 137–143.
10. C. Dorr, Sleeping Beauty: In defence of Elga, *Analysis* **62** (2002) 292–296.
11. R. Durrett, *Probability: Theory and Examples*. Wadsworth, Pacific Grove, CA, 1991.
12. A. Elga, Self-locating belief and the Sleeping Beauty problem, *Analysis* **60** no. 2 (2000) 143–147.
13. B. C. van Fraassen, Belief and the will, *J. Philosophy* **81** (1984) 235–256.
14. ———, Belief and the problem of Ulysses and the Sirens, *Philos. Stud.* **77** (1995) 7–37.
15. B. Groisman, The end of Sleeping Beauty’s nightmare, *British J. Philos. Sci.* **59** (2008) 409–416.
16. J. Y. Halpern, Sleeping Beauty reconsidered: Conditioning and reflection in asynchronous systems, *Oxford Studies in Epistemology* **3** (2005) 111–142.
17. P. Hawley, Inertia, optimism and Beauty, *Noûs* **1** no. 1 (2013) 85–103.
18. C. Hitchcock, Beauty and the bets, *Synthese* **139** no. 3 (2004) 405–420.
19. T. Horgan, Sleeping Beauty awakened: New odds at the dawn of the new day, *Analysis* **64** (2004) 10–21.
20. ———, Synchronic Bayesian updating and the generalized Sleeping Beauty problem, *Analysis* **67** no. 1 (Jan., 2007) 50–59.
21. K. Karlander, L. Spectre, Sleeping Beauty meets Monday, *Synthese* **174** (2010) 397–412.

22. B. Kierland, B. Monton, Minimizing inaccuracy for self-locating beliefs, *Philosophy and Phenomenological Research* **70** (2005) 384–395.
23. N. Kim, Sleeping Beauty and shifted Jeffrey conditionalization, *Synthese* **168** (2009) 295–312.
24. A. N. Kolmogorov, *Foundations of the Theory of Probability*. Ed. N. Morrison. Chelsea, New York, 1950.
25. D. Lewis, Sleeping Beauty: Reply to Elga, *Analysis* **61** (2001) 171–176.
26. P. J. Lewis, Quantum Sleeping Beauty, *Analysis* **67** no. 293, 59–65.
27. I. Mariolis, Revealing the beauty behind the Sleeping Beauty problem (2014), <http://arxiv.org/pdf/1409.3803.pdf>.
28. C. J. G. Meacham, Sleeping Beauty and the dynamics of *de se* beliefs, *Philos. Stud.* **138** no. 2 (2008) 245–269.
29. B. Monton, Sleeping Beauty and the forgetful Bayesian, *Analysis* **62** no. 1 (Jan. 2002) 47–53.
30. P. Mutalik, Why Sleeping Beauty is lost in time, *QUANTA* (March 31, 2016), <http://quantamagazine.org/20160331-why-sleeping-beauty-is-lost-in-time>.
31. M. Piccione, A. Rubinstein, On the interpretation of decision problems with imperfect recall, *Games Econom. Behav.* **20** (1997) 3–24.
32. ———, The absent-minded driver’s paradox: Synthesis and responses, *Games Econom. Behav.* **20** (1997) 121–130.
33. J. Pittard, When Beauties disagree: Why halfers should affirm robust perspectivalism. *Oxford Studies in Epistemology* **5** (forthcoming).
34. J. Pust, Horgan on Sleeping Beauty, *Synthese* **160** (2008) 97–101.
35. J. S. Rosenthal, A mathematical analysis of the Sleeping Beauty problem, *Math. Intelligencer* **31** (2009) 32–37.
36. J. Ross, Sleeping Beauty, countable additivity, and rational dilemmas, *Philos. Rev.* **119** (2010) 411–447.
37. W. Schwarz, The Absentminded Driver: No paradox for halfers, (draft, February 11, 2008), <http://www.umsu.de/words/driver.pdf>.
38. T. Seidenfeld, M. J. Schervish, J. B. Kadane, Stopping to reflect, *J. Philosophy* **101** no. 6 (2004) 315–322.
39. S. Tadelis, *Game Theory*, Princeton Univ. Press, Princeton NJ 2013.
40. M. Teitelbaum, *Sleeping Beauty* (Disney Princess ed.). Little Golden Books, New York, 2004.
41. M. G. Titelbaum, The relevance of self-locating beliefs, *Philos. Rev.* **177**, 555–605.
42. ———, Ten reasons to care about the Sleeping Beauty problem, *Philos. Compass* **8** no. 11 (2013) 1003–1017.
43. ———, Self-locating credences, in *The Oxford Handbook of Probability and Philosophy*. Ed. A. Hajek, C. Hitchcock. Oxford Univ. Press (forthcoming).
44. L. Vaidman, S. Saunders, On Sleeping Beauty controversy (2009), <http://philsciarchive.pitt.edu/archive/00000324>.
45. R. Weintraub, Sleeping Beauty: A simple solution, *Analysis* **64** (2004) 8–10.
46. R. White, The generalized Sleeping Beauty problem: A challenge for thirders, *Analysis* **66** (2006) 114–119.
47. A. Wilson, Everettian confirmation and Sleeping Beauty, *British J. Philos. Sci.* (forthcoming).

PETER WINKLER is William Morrill Professor of Mathematics and Computer Science at Dartmouth College. A winner of the Lester R. Ford award for mathematical exposition, he is the author of 150 research papers, a dozen patents, three books and a portfolio of piano rags.

Dept. of Mathematics, Dartmouth College, Hanover, NH 03755-3551 USA
peter.winkler@dartmouth.edu