Sleeping Beauty, Countable Additivity, and Rational Dilemmas

Jacob Ross
University of Southern California

Currently, the most popular views about how to update *de se* or self-locating beliefs entail the one-third solution to the Sleeping Beauty problem.[1] Another widely held view is that an agent's credences should be countably additive.[2] In what follows, I will argue that there is a deep tension between these two positions. For the assumptions that underlie the one-third solution to the Sleeping Beauty problem entail a more general principle, which I call the Generalized Thirder Principle, and there are situations in which the latter principle and the principle of Countable Additivity cannot be jointly satisfied. The most plausible response to this tension, I argue, is to accept both of these principles and to maintain that when an agent cannot satisfy them both, he or she is faced with a rational dilemma.

In writing this essay, I benefited enormously from comments from, and discussions with, Frank Arntzenius, Cian Dorr, Adam Elga, Branden Fitelson, Matthew Kotzen, David Manley, Chris Meacham, Sarah Moss, Mark Schroeder, Teddy Seidenfeld, Mike Titelbaum, Peter Vranas, Brian Weatherson, Jonathan Weisberg, Gideon Yaffe, and an anonymous referee. My greatest debt is to Kenny Easwaran, for uncountably many invaluable discussions.

1. This solution is defended in Arntzenius 2003, Dorr 2002, Draper and Pust 2008, Elga 2000, Hitchcock 2004, Horgan 2004, Monton 2002, Seminar 2008, Stalnaker 2008a, Titelbaum 2008, and Weintraub 2004.

2. See, for example, Howard 2006, Walley 1991, Weatherson 2005, and Williamson 1999. The principle of Countable Additivity traces back to Kolmogorov 1956 [1933].

Philosophical Review, Vol. 119, No. 4, 2010
DOI 10.1215/00318108-2010-010
© 2010 by Cornell University

I will begin, in the first section, by generalizing the original Sleeping Beauty problem and showing how the one-half and one-third solutions to this problem can be extended to the more general problem. In the second section, I will show that, of the two resulting principles, only the Generalized Thirder Principle (GTP) conflicts with Countable Additivity (CA). And I will show that this conflict arises even for thirders who do not accept any principle of indifference that conflicts with CA. In the third section, I will show that if one accepts any of the main arguments for the one-third solution, then one is committed to the GTP since the latter principle follows from the premises of each of these arguments. In the fourth section, I will examine, in greater detail, one kind of argument for the one-third solution, namely, Dutch book arguments. I will show that, from the premises of this kind of argument, we can derive both the GTP and the conflicting principle of CA. In the concluding section, I will consider various possible responses to the conflict between CA and the GTP, and I will argue that the most plausible response is to acknowledge the possibility of rational dilemmas.

1. Generalizing the Sleeping Beauty Problem

The Sleeping Beauty problem is a problem about how we should rationally respond to changes in what we know about our location. The basic problem is that the standard model of how we should revise our degrees of belief, namely, the model of conditionalization, assumes that changes in our beliefs proceed by way of eliminating possibilities. However, if we allow for possibilities that can be represented by centered or *de se* propositions—propositions of the form *I am now in such-and-such a location*—then it seems that changes in our self-locating information can involve adding possibilities that had formerly been ruled out. Thus, when you are driving to work, you can rule out the possibility that you could express by saying "I am now in my office," but once you are sitting at your desk, you can no longer rule out this possibility. Thus, even if conditionalization accounts for how we should respond to the elimination of self-locating possibilities, it is widely agreed that it does not account for how we should respond to all changes in self-locating information.[3] What is controversial, however, is what model we should adopt in its place. Much of this controversy has focused on the following case:

3. See Arntzenius 2003, Hitchcock 2004, Meacham 2008, and Titelbaum 2008.

> *Original Sleeping Beauty Problem.* A fully rational agent named Beauty participates in an experiment in which she falls asleep on Sunday night, and the number of times she subsequently awakens is determined by the toss of a fair coin. If the coin comes up heads, then she is awoken only on Monday morning, and after she falls asleep on Monday night, she remains asleep for the duration of the experiment. If, however, the coin comes up tails, then she will be awoken first on Monday morning and again on Tuesday morning. But while she is asleep on Monday night, all her memories of the previous day will be erased, so that when she awakens on Tuesday morning, it will seem like her first experimental awakening. The problem is to determine what her credence should be, when she awakens on Monday morning, in the proposition that the coin came up heads.

In this case, the only change that occurs in Beauty's information between Sunday evening and Monday morning is that she loses the self-locating information that it is Sunday evening and gains the self-locating information that it is either Monday morning or Tuesday morning. But Beauty knew all along that she would undergo this change in her self-locating information, regardless of whether the coin comes up heads or tails (let us call these two possibilities *Heads* and *Tails*, respectively). It may seem, therefore, that when Beauty awakens on Monday morning, she does not learn anything that should influence her opinions about the outcome of the coin toss. Thus, it may seem that since her credence in Heads should be one-half on Sunday evening, it should remain one-half on Monday morning. (Call those who endorse this view *halfers.*[4])

In spite of the prima facie plausibility of the one-half solution, a spate of arguments has been given for the view that when Beauty awakens on Monday morning, her credence in Heads should instead be one-third. (Call those who endorse this view *thirders.*) We will consider several arguments for the one-third solution in section 3. But in the meantime it will be useful to generalize the original Sleeping Beauty problem, as well as the two standard solutions to it.

Let us define a *Sleeping Beauty problem* as any problem in which a fully rational agent, Beauty, will undergo one or more mutually indistinguishable awakenings and in which the number of awakenings she will undergo is determined by the outcome of a random process. Let *S* be a partition of alternative hypotheses concerning the outcome of this process. Beauty knows the objective chance of each hypothesis in *S*, and she also knows how many times she will awaken conditional on each of these

4. See Lewis 2001 for an early defense of this position.

hypotheses, but she has no other relevant information. The problem is to determine how her credence should be divided among the hypotheses in S when she first awakens. The original Sleeping Beauty problem is the instance of this class of problems in which the random process that determines how many times Beauty awakens is a fair coin toss and in which Beauty awakens once given Heads and twice given Tails.

The most natural way to generalize the one-half solution to the original Sleeping Beauty problem is to say that in any Sleeping Beauty problem, what Beauty learns when she first awakens should have no effect on her credence in any of the hypotheses in S. Hence, since, prior to her first awakening, her credence in any given hypothesis in S should be equal to its objective chance, it should remain so when she first awakens. Let us call this the *Generalized Halfer Principle.*

The most natural way to generalize the one-third solution is as follows:

> *Generalized Thirder Principle* (*GTP*). In any Sleeping Beauty problem, defined by a partition S, upon first awakening, Beauty's credence in any given hypothesis in S should be proportional to the product of its objective chance and the number of times Beauty awakens if this hypothesis is true.[5]

Since, therefore, in the original Sleeping Beauty problem, Heads and Tails have the same objective chance, and Beauty would awaken twice as often given Tails as given Heads, it follows from the GTP that, upon first awakening, Beauty's credence in Heads should be one-half her credence in Tails. And since these two credences must sum to one, her credence in Heads should be one-third.

It will be useful to express this principle formally. For any hypothesis i in S, let $Ch(i)$ be the objective chance that hypothesis i is true, and let $N(i)$ be the number of times Beauty awakens if i is true. Let P be Beauty's credence function upon first awakening. Where i and j are any two hypotheses in S, so long as $P(j)$ and $Ch(j)$ are positive, and so long as $N(i)$ and $N(j)$ are both integers, we can state the GTP in ratio form, as follows:

$$\text{For all } i, j \in S, \frac{P(i)}{P(j)} = \frac{N(i)Ch(i)}{N(j)Ch(j)} \quad \text{whenever } P(j), Ch(j) > 0, \text{ and } N(i), N(j) \in Z$$

5. I am grateful to Chris Meacham for a very helpful discussion concerning this generalization.

To allow for the possibility that $P(j)$ or $Ch(j)$ equals zero, we need to state the principle in product form, as follows:

$$\text{For all } i, j \in S,\ P(i)\,N(j)\,Ch(j) = P(j)\,N(i)\,Ch(i) \quad \text{whenever} \quad N(i), N(j) \in Z$$

The qualification "whenever $N(i), N(j) \in Z$" indicates that the GTP applies only to pairs of hypotheses, i and j, such that the number of times Beauty will awaken conditional on each is an integer. This qualification is necessary, for if either $N(i)$ or $N(j)$ is infinite, then one of the products will be undefined. Thus, the GTP is an essentially finitistic principle.

2. The Conflict between Countable Additivity and the GTP

The principle of Countable Additivity can be stated thus:

> *Countable Additivity* (*CA*). For any set of countably many propositions, any two of which are incompatible, rationality requires that one's credences in the propositions in this set sum to one's credence in their disjunction.

The Generalized Halfer Principle does not generate any violations of CA. For according to the former principle, in any Sleeping Beauty problem defined by a partition S, when Beauty first awakens, her credences in the hypotheses in S should remain what they were on Sunday night, when all parties to the debate agree that they should be equal to the objective chances of these hypotheses. Hence, assuming objective chances are countably additive, her credences in these hypotheses will always satisfy CA. By contrast, the Generalized Thirder Principle does generate violations of CA, as can be seen from the following case:

> *Sleeping Beauty in St. Petersburg*. This is a Sleeping Beauty problem in which the random process that determines how many times Beauty awakens is a series of coin tosses in which a fair coin is tossed repeatedly until it first lands heads. The set S consists in all the alternative hypotheses concerning the length of this series of coin tosses. Thus, for each positive integer x, S will include a *finite sequence hypothesis* h_x, according to which the length of the sequence of coin tosses is x; if this hypothesis is true, Beauty will awaken 2^x times. S will also include the *infinite sequence hypothesis* h_∞, or the hypothesis that the coin never comes up heads and hence the sequence is infinite; if this hypothesis is true, Beauty will awaken exactly once.

In this problem, for any finite sequence hypothesis h_x, $N(h_x) = 2^x$, and $Ch(h_x) = 1/2^x$ since this is the objective chance that a fair coin first lands heads on the *x*th toss. As for the infinite sequence hypothesis h_∞, $N(h_\infty) = 1$ and $Ch(h_\infty) = 0$ since this is the objective chance of an infinite sequence of consecutive tails tosses. Hence, it follows from the Generalized Thirder Principle that for any two positive integers, *x* and *y*,

$$P(h_x)N(h_y)Ch(h_y) = P(h_y)N(h_x)Ch(h_x)$$

Hence, $P(h_x)(\frac{2^y}{2^y}) = P(h_y)(\frac{2^x}{2^x})$

And so $P(h_x) = P(h_y)$. Thus, when Beauty first awakens, she must have the same credence in each of the finite sequence hypotheses. Further, it follows that

$$P(h_\infty) = \frac{P(h_1)N(h_\infty)Ch(h_\infty)}{N(h_1)Ch(h_1)} = \frac{P(h_1)\cdot 1\cdot 0}{1} = 0$$

Thus, upon first awakening, Beauty must have a credence of zero in the infinite sequence hypothesis. And so her credence in the disjunction of the finite sequence hypotheses must be one. But since there is a countably infinite number of finite sequence hypotheses, and Beauty must have the same credence in each, her credences in these hypotheses cannot sum to one, and so it cannot sum to her credence in their disjunction. Thus, it follows from the Generalized Thirder Principle that Beauty is rationally required to violate Countable Additivity.[6] The two principles therefore conflict.

One might think that the conflict between the GTP and CA is not a special problem for thirders. For thirders and halfers generally agree that when Beauty awakens on Monday morning in the original Sleeping Beauty problem, she should have the same credence in the centered proposition that can be expressed "the coin landed tails, and I have awoken exactly once" as she has in the centered proposition that can be expressed "the coin landed tails, and I have awoken exactly twice." This claim is generally justified by appealing to an indifference principle. Just how this principle should be formulated is controversial, but one might

6. One might attempt to avoid this conclusion by allowing for infinitesimal probabilities. For a critique of this approach, see Hájek 2003. Moreover, while the details of the arguments would need to be modified, variants of the arguments for CA from Dutch book coherence and from conglomerability, which we will be discussing later on, can be given against credence functions that assign equal, infinitesimal credence to each element in a countably infinite partition. And so even if such credence functions were compatible with CA, they would be as problematic as ones that violate CA.

suppose that any indifference principle that would entail this particular claim would also entail the following, more general claim:

> *Sleeping Beauty Indifference* (*SBI*). In any Sleeping Beauty problem, for any hypothesis *h* in *S*, upon first awakening, Beauty should have equal credence in each of the awakening possibilities associated with *h*,

where an *awakening possibility associated with h* is a centered proposition of the form "hypothesis *h* is true, and it is now the *n*th morning of the experiment" that is consistent with what Beauty knows when she first awakens.

This principle appears to conflict with CA. For it implies that, in a Sleeping Beauty problem, where *S* consists of only a single hypothesis, *h*, and *h* is associated with infinitely many awakening possibilities, Beauty must have equal credence in each of these awakening possibilities. And so her credences in these possibilities cannot possibly sum to their disjunction, which is one.[7] One might claim, therefore, that since almost everyone agrees that upon awakening, Beauty should have equal credence in the two awakening possibilities associated with Tails, almost everyone is committed to SBI. And since CA conflicts with SBI, all parties to the debate should reject CA, regardless of whether they accept the Generalized Thirder Principle.

But there are two problems with this argument. The first is that CA does not clearly conflict with SBI. The conflict between SBI and CA arises only when the latter principle is understood to range not only over ordinary propositions but also over *centered* propositions of the form "it is now the *n*th morning of the experiment." And so one might hold that the proper lesson to be drawn from the above argument is that Countable Additivity should be understood as ranging only over uncentered propositions.

The second problem is that, in order to defend the claim that Beauty should have equal credence in the two awakening possibilities associated with Tails, one needn't appeal to SBI. For one might instead appeal to the following, weaker principle:

> *Finitistic Sleeping Beauty Indifference* (*FSBI*). In any Sleeping Beauty problem, for any hypothesis *h* in *S*, if the number of times Beauty awakens conditional on *h* is finite, then upon first awakening, Beauty should have equal credence in each of the awakening possibilities associated with *h*.

7. For a very similar argument against a more general indifference principle, see Weatherson 2005.

And this principle does not, by itself, conflict with CA.[8] Moreover, as we will see in the next section, the GTP can be justified without appealing to any infinitistic indifference principle.

3. Why Thirders Are Committed to the Generalized Thirder Principle

In this section, I will consider a variety of arguments for the one-third solution and argue that each one is best understood as resting on assumptions that entail the GTP.

3.1. The Long-Run Frequency Argument

One argument for the one-third solution, first sketched in Elga 2000, can be stated thus:

> Suppose Beauty were to know, when she awakens, that the experiment in which she is participating is to be repeated many times and that she is now awakening in one of these repetitions, though she has no idea which. If the number of repetitions were sufficiently large, then she could be confident that about one-third of them are *Heads-awakenings*—awakenings that occur on trials in which the coin lands Heads. Since, in this case, she should regard her present awakening as equally likely to be any one of the actual awakenings, and since she believes that about one-third of these are Heads-awakenings, her credence that her present awakening is a Heads-awakening should be one-third. But Beauty's credence in the proposition that her current awakening is a Heads-awakening should not depend on how many times she believes the experiment is repeated, and so if it should be one-third in the many-trials case, it should likewise be one-third in the single-trial case. And since, in the original Sleeping Beauty problem, Beauty's credence in Heads, upon first awakening, must be equal to her credence in the proposition that her present awakening is a Heads-awakening, it follows that her credence in Heads should be one-third.

8. In Weatherson 2005, 630–33, there is an argument that suggests that even FSBI may lead to violations of CA. This argument involves an example that is similar to the case of Sleeping Beauty in St. Petersburg, in that it involves a finite but unbounded number of indistinguishable experiences. This argument, however, relies on an assumption about how the Principal Principle is to be applied. This assumption is analogous to what I call "First Awakening Admissibility" in section 3.2, and this assumption would be rejected by many halfers. Thus, Weatherson's argument does not show that FSBI, on its own, conflicts with CA.

In this argument, the following two principles appear to be implicit:

> *Long-Run Frequency.* In any Sleeping Beauty problem, if Beauty knew that the experiment she is undergoing was to be repeated many times, holding constant the objective chance of each of the possible outcomes of the experiment, as well as the number of times she awakens conditional on each outcome, then her credence in any hypothesis h, in S, should be proportional to the expected long-run frequency of awakenings that occur on trials of the experiment with the outcome predicted by h.
>
> *Repetition Independence.* In any Sleeping Beauty problem, the credence Beauty should have in any hypothesis h, in S, should not depend on the number of times that she believes the experiment is to be repeated.

On the basis of these two assumptions, we can derive the GTP. For if a Sleeping Beauty experiment were repeated many times, then for any hypothesis h, in S, the expected long-run frequency of awakenings that occur on trials in which the experiment has the outcome predicted by h would be proportional to the product of the objective chance of this outcome and the number of times Beauty awakens conditional on this outcome. Thus, this expected frequency will be proportional to the product of $Ch(h)$ and $N(h)$. Hence, it follows from Long-Run Frequency that, if Beauty knows that the experiment will be repeated many times, then she should satisfy the GTP. And so it follows, from Repetition Independence, that she should satisfy the GTP even if the experiment is performed only once.

3.2. The Argument from the Principal Principle

A second argument for the one-third solution, also first presented in Elga 2000, can be stated thus:

> There are two ways in which the Sleeping Beauty experiment could be run: the experimenter could toss the coin on Sunday night, before Beauty first awakens, or he could toss it on Monday night. Since Beauty will awaken on Monday regardless, this difference should not affect Beauty's credences on Monday morning. And so, without loss of generality, we can assume the coin is tossed on Monday night. Now suppose that on Monday morning, Beauty awakens at 9 a.m., and at 9:01 she is told that it is Monday. In this case, at 9:01, Beauty will know that the coin has not yet been tossed. Hence, she will know that none of her experiences can have been affected by the outcome of the coin toss. And so she will know that she has no evidence bearing on the outcome of the coin toss, beyond her knowledge of the objective chances of the two outcomes. And so her

credence in Heads should then be equal to its known objective chance, which is one-half. But assuming, in accordance with FSBI, that Beauty initially has equal credence in the two awakening possibilities associated with Tails, it follows that she will have credence one-half in Heads, after learning that it is Monday, only if, when she first awakens, her credence in Heads is one-third.

This argument can be formulated more explicitly on the basis of the following four assumptions, together with FSBI:

Restricted Principal Principle. If a fully rational agent knows the objective chance of a hypothesis h, and he or she knows that he or she possesses no information that is inadmissible in relation to h, then his or her credence in h must be equal to its known objective chance.[9]

First Awakening Admissibility. In any Sleeping Beauty problem, if, after first awakening, Beauty were to learn only that she has awoken exactly once, then she would know that she has no information that is inadmissible with respect to any hypothesis in S.

Finite Additivity. If an agent is fully rational, then for any set of finitely many propositions, any two of which are incompatible, his or her credence in the propositions in this set must sum to his or her credence in their disjunction.

Restricted Conditionalization. If an agent is fully rational, and the only change that occurs in his or her epistemic situation is that he or she learns some centered or uncentered proposition e, in which he or she antecedently had positive credence, then he or she must update his or her credence in any uncentered proposition h, by conditionalizing on e. Hence, if P_1 and P_2 represent his or her credence functions before and after learning e,

$$P_2(h) = P_1(h \mid e) = \frac{P_1(h \,\&\, e)}{P_1(e)} \text{ whenever } P_1(e) > 0$$

Now let P represent Beauty's credence function in the original Sleeping Beauty problem upon first awakening. And let P^+ represent what her credence function would be one minute later if she were then to learn

9. This principle is "restricted" because a condition for its application is that the agent not only possesses no inadmissible evidence, but also *knows* that he or she possesses no such evidence. Here I am grateful to Brian Weatherson for a very illuminating discussion.

that she had awoken only once. When Beauty first awakens, her credence will be divided among the following three possibilities:

- T1: The coin landed tails, and it is now the first morning of the experiment;
- T2: The coin landed tails, and it is now the second morning of the experiment;
- H1: The coin landed heads, and it is now the first morning of the experiment.

From the Restricted Principal Principle, together with First Awakening Admissibility, it follows that her credence in Heads at 9:01 must be equal to its known objective chance, which is one-half. Hence, $P^+(\text{H1}) = P^+(\text{Heads}) = 1/2$. Further, since the only change in Beauty's epistemic situation between 9:00 and 9:01 is that she learns $(\text{H1} \vee \text{T1})$, and since her prior credence in this disjunction is greater than zero, we can apply Restricted Conditionalization. Thus, $P^+(\text{H1}) = P(\text{H} \mid \text{H1} \vee \text{T1})$. Therefore, $P(\text{H1} \mid \text{H1} \vee \text{T1}) = 1/2$, which implies that $P(H1) = P(T1)$. And it follows from FSBI that $P(\text{T1}) = P(\text{T2})$. Hence, $P(\text{T1}) = P(\text{T2}) = P(\text{H1})$. And by Finite Additivity, these three credences must sum to one. Therefore, $P(\text{Heads}) = P(\text{H1}) = 1/3$.

I will now show that the five premises that underlie this argument entail the GTP. In any Sleeping Beauty problem defined by a partition S, let P represent Beauty's credence function upon first awakening and let P^+ represent the credence function she would have one minute later if she were then to learn that she had awoken only once. For any positive integer x, let w_x be the centered proposition that can be expressed thus: "It is now the *x*th morning of the experiment." Now from the Restricted Principal Principle, together with First Awakening Admissibility, it follows that when Beauty learns that she has awoken only once, her credence in each hypothesis in S must be equal to its objective chance. That is,

$$\text{For all } h \in S,\ P^+(h) = Ch(h)$$

Now let us assume, for the time being, that $P(w_1) > 0$. We can therefore apply Restricted Conditionalization, as follows:

$$\text{For all } h \in S,\ P^+(h) = P(h \mid w_1) = \frac{P(h \,\&\, w_1)}{P(w_1)}$$

Hence, from the last two equations,

$$\text{For all } h \in S,\ Ch(h) = \frac{P(h \,\&\, w_1)}{P(w_1)}$$

By Finite Additivity, for any hypothesis h, in S, Beauty's credences in the various awakening possibilities associated with h must sum to her credence in h. And, from FSBI, so long as the number $N(h)$ of awakening possibilities associated with h is finite, Beauty must have equal credence in each of these awakening possibilities. It follows that her credence in h must be $N(h)$ times her credence in the first of these awakening possibilities, $(h \,\&\, w_1)$ (that is, the centered proposition that can be expressed "h is true and it is now the first morning of the experiment"). That is,

$$\text{For all } h \in S,\ P(h) = N(h)\,P(h \,\&\, w_1) \text{ whenever } N(h) \in Z$$

From the last two equations,

$$\text{For all } h \in S,\ P(h) = N(h)\,Ch(h)\,P(w_1) \text{ whenever } N(h) \in Z$$

This entails

$$\text{For all } i, j \in S,\ P(i)\,N(j)\,Ch(j)\,P(w_1) = P(j)\,N(i)\,Ch(i)\,P(w_1) \quad \text{whenever } N(h) \in Z$$

And since we are now assuming that $P(w_1) > 0$, this entails the GTP.

Suppose, on the other hand, that $P(w_1) = 0$. It follows that for any hypothesis h, in S, $P(h \,\&\, w_1) = 0$. Hence, so long as the number of awakening possibilities associated with h is finite, FSBI entails that the probability of each of these possibilities must be zero. And so it follows, by Finite Additivity, that the probability of h must be zero. Thus, assuming $P(w_1) = 0$, we can conclude that for any hypothesis h, in S, $P(h) = 0$ whenever $N(h) \in Z$. And in this case, the GTP will be trivially satisfied. Hence, whatever we assume about the value of $P(w_1)$, the premises of the Principal Principle argument entail the GTP.

3.3. The Hypothetical Priors Argument

Another standard argument proceeds as follows:

> Consider a variant of the original Sleeping Beauty problem in which, in addition to the three awakening possibilities in the original problem (T1, T2, and H1), there were a fourth:

H2: The coin landed heads, and it is now the second morning of the experiment.

Call this the *equal number case* since it involves an equal number of awakening possibilities associated with each outcome of the coin toss. In this case, Beauty should regard the two awakening possibilities associated with Heads as equally probable, just as she should regard the two possibilities associated with Tails as equally probable. And since, in this case, what she knows about her self-location is the same, conditional on Heads, as it is conditional on Tails, it follows that her self-locating information does not favor either Heads or Tails. Hence, her credences in Heads and in Tails should each remain one-half. Consequently, she should have equal credence (one-quarter) in each of the four awakening possibilities.

But the only difference between the two cases, with respect to the awakening possibilities that are consistent with Beauty's evidence when she awakens, is that the equal number case includes H2. Consequently, if, after awakening in the equal number case, Beauty were to learn that H2 is false, then she would have the same relevant knowledge—and should thus have the same credence in Heads—as she has when she first awakens in the original case. But in the equal number case, she begins with credence one-quarter in each of T1, T2, H1, and H2. And so if she were then to learn that H2 is false, she would end up with one-third credence in Heads. So her credence in Heads should likewise be one-third when she first awakens in the original Sleeping Beauty case.

Variations of this argument can be found in Dorr 2002, Arntzenius 2003, and Horgan 2004, 2008. We can axiomatize this argument on the basis of the same premises as the Principal Principle Argument, except replacing the controversial assumptions of First Awakening Admissibility with the following two assumptions:

Equal Number Admissibility. In any Sleeping Beauty problem in which Beauty awakens the same number of times regardless of which hypothesis in *S* is true, upon first awakening, Beauty knows that she has no information that is inadmissible with respect to any hypothesis in *S*.

Hypothetical Priors. For any two Sleeping Beauty problems, SB_1 and SB_2, that are defined by the same partition *S*, if the set of awakening possibilities in SB_1 is a subset of those in SB_2, then immediately after first awakening in SB_1, Beauty's credences in the hypotheses in *S* should be the same as they would be if she had begun with the credences she has upon first awakening in SB_2, and she then acquired the additional self-locating information she has in SB_1.[10]

To see how these premises entail the one-third solution to the original Sleeping Beauty problem, let P_E represent Beauty's credences upon first awakening in the equal number case and let P_E^+ represent the credence function she would have if she began with P_E and then learned that ¬H2. From Restricted Conditionalization, we can infer

$$P_E^+(\mathrm{H1}) = P_E(\mathrm{H1} \mid \neg\mathrm{H2}) = \frac{P_E(\mathrm{H1} \,\&\, \neg\mathrm{H2})}{P_E(\neg\mathrm{H2})} = \frac{P_E(\mathrm{H1})}{P_E(\neg\mathrm{H2})}$$

And from the Principal Principle, together with Equal Number Admissibility, it follows that $P_E(\mathrm{Heads}) = 1/2$. And since, in the equal number case, there are two awakening possibilities associated with Heads (H1 and H2), it follows from FSBI that when Beauty awakens, she must have equal credence in each. Since they must sum to her credence in Heads, which is one-half, it follows that $P_E(\mathrm{H1}) = P_E(\mathrm{H2}) = 1/4$. Hence, $P_E(\neg\mathrm{H2}) = 3/4$. Plugging these values into the above equation,

$$P_E^+(\mathrm{H1}) = \frac{1/4}{3/4} = \frac{1}{3}$$

But since the set of awakening possibilities in the original Sleeping Beauty problem is a subset of those in the equal number case, and since the additional information Beauty possesses in the former case is ¬H2, it follows from Hypothetical Priors that $P(\mathrm{H1}) = P_E^+(\mathrm{H1}) = 1/3$. Hence, upon awakening in the original Beauty problem, Beauty's credence in Heads must be one-third.

I will now show how the GTP can be derived from the premises of the Hypothetical Priors argument. Let SB be an arbitrary Sleeping Beauty problem defined by a partition, S, and let SB_E be a variant of SB in which Beauty will awaken exactly once regardless of which hypothesis in S is true. Let P and P_E represent Beauty's credences, upon first awakening, in SB and SB_E, respectively. And let P^+ represent the credences Beauty would have if she began with P and then learned that she has awoken only once.

Since, in SB_E, Beauty will awaken exactly once no matter what, it follows from the Principal Principle, together with Equal Number

10. For passages indicating that Arntzenius, Dorr, and Horgan are committed to these two principles, see Arntzenius 2003, 363–64; Dorr 2002, 294–95; and Horgan 2008, 158.

Admissibility, that

$$\text{For all } h \in S,\ P_E(h) = Ch(h)$$

Since the set of awakening possibilities in SB_E is a subset of those in SB, and the only additional self-locating information Beauty has in SB_E is that she has awoken only once, it follows from Hypothetical Priors that

$$\text{For all } h \in S,\ P_E(h) = P^+(h)$$

Hence, from the last two equations,

$$\text{For all } h \in S,\ P^+(h) = Ch(h)$$

As we saw in the previous section, from the above equation, along with Restricted Conditionalization, Finite Additivity, and FSBI, we can derive the GTP.

3.4. The De Re *Conditionalization Argument*

Robert Stalnaker has offered an argument for the one-third solution that is similar to the one just considered, in that it claims that when Beauty awakens on Monday morning, she updates by eliminating a possibility associated with Heads. However, on Stalnaker's view, the possibility she eliminates should be thought of not as a centered proposition but rather as a *de re* proposition concerning the time at which she awakens. This enables him to understand the change in Beauty's credences between Sunday evening and Monday morning as occurring by a simple process of conditionalization. The argument can be stated thus:[11]

> In order to understand the change in Beauty's credences between Sunday evening and Monday morning, Beauty must be represented as learning, on Monday morning, some uncentered proposition in whose negation

11. See Stalnaker 2008a, 59–64, as well as Stalnaker 2008b. For a more detailed analysis of this argument and of how it commits Stalnaker to the GTP, see Ross 2010. Weatherson n.d. offers an alternative interpretation of Stalnaker's argument and provides a very interesting discussion of its connection with the GTP. Whereas, on my reading of Stalnaker, Beauty arrives at her Monday morning credences from her Sunday evening credences via conditionalization, on Weatherson's interpretation, Beauty arrives at these credences by first *recalibrating* on Monday morning, assigning credences to propositions she was initially unable to represent, and then conditionalizing. Thus, on Weatherson's reading, Stalnaker's argument is a version of the hypothetical priors argument, discussed in section 3.3.

> she had positive credence on Sunday evening. In particular, where t is the time at which she awakens on Monday morning, she must be represented as learning, of t, that it is a time at which she awakens, and hence as conditionalizing on this information. Now if Beauty antecedently knew when t is located, then she would antecedently know that she awakens at t, and so she could not learn this on Monday morning. Beauty must therefore be represented as not yet knowing, on Sunday evening, when t is. What she does know on Sunday evening, however, is that she awakens twice as often given Tails as she does given Heads. Hence, her prior credence in the proposition that she awakens at t should be twice as great conditional on Tails as it is conditional on Heads. But on Sunday evening, Beauty should have equal credence in Heads and in Tails. And so, on Monday morning, after conditionalizing on a proposition that she regarded as twice as probable given Tails as given Heads, her credence in Heads should shift to one-third.

This argument can be formalized on the basis of the Restricted Principal Principle, together with the following three premises:

Prior Admissibility. In any Sleeping Beauty problem, on the evening before the first experimental awakening, Beauty knows that she has no information that is inadmissible in relation to any hypothesis in S.

De Re Conditionalization. In any Sleeping Beauty problem, where t is the time at which Beauty first awakens, for any hypotheses h, in S, Beauty's credence in h, upon first awakening, should be equal to the credence she would have in h on the previous evening conditional on the proposition that she awakens at t, assuming that on the previous evening Beauty had no information about the location of t.

Proportionality. In any Sleeping Beauty problem, where t is the time at which Beauty first awakens, if, on the evening before the first awakening, Beauty had no information about the location of t, then her credence that she awakens at t, conditional on any hypothesis h, in S, should be proportional to the number of times she awakens given h.

In order to state this principle formally, let t be the time of Beauty's first awakening and let A be the proposition that she awakens at t. Let P represent her credences when she first awakens and let P^{-} represent the credences she would have on the previous evening, assuming that at that time she had no information about the location of t. The principle of

Proportionality can thus be stated as follows:

$$\text{For all } i, j \in S, P^{-}(A \mid i) N(j) = P^{-}(A \mid j) N(i) \text{ whenever } N(i), N(j) \in Z$$

How the one-third solution can be derived from the above premises is fairly clear from the informal exposition of the argument given above. And so we may proceed directly to the derivation of the GTP. It follows from *De Re* Conditionalization that, for any hypothesis h, in S, $P(h) = P^{-}(h \mid A)$. Hence,

$$\text{For all } i, j \in S, P(i) P^{-}(j \mid A) = P(j) P^{-}(i \mid A)$$

This entails

$$\text{For all } i, j \in S, P(i) P^{-}(A \mid j) P^{-}(j) = P(j) P^{-}(A \mid i) P^{-}(i)$$

And by the Principal Principle, together with Prior Admissibility, the prior probabilities of i and j must be equal to their objective chances. Hence,

$$\text{For all } i, j \in S, P(i) P^{-}(A \mid j) Ch(j) = P(j) P^{-}(A \mid i) Ch(i)$$

Hence, from Proportionality,

$$\text{For all } i, j \in S, P(i) N(j) Ch(j) = P(j) N(i) Ch(i) \text{ whenever } N(i), N(j) \in Z$$

3.5. The Diachronic Dutch Book Argument

A final argument for the one-third solution is the diachronic Dutch book argument. Several authors have given diachronic Dutch book arguments against the one-half solution.[12] That is, they have argued that if Beauty's credence in Heads, upon awakening, is one-half, then she will be vulnerable to a *Dutch book strategy*, meaning that she can be offered a sequence of bets, each of which she regards as fair, such that her taking all of these bets would result in a sure loss. But we can show, more generally, that if Beauty's credence in Heads, upon awakening, is anything other than one-third, then she will be vulnerable to such a strategy. The key to this argument is that if a bookie offers Beauty a bet on Heads (or on Tails) each time she awakens, then the number of times he offers her this bet

12. See, for example, Arntzenius 2002, Hitchcock 2004, and Draper and Pust 2008.

will depend on whether Heads or Tails is true, and so it will depend on whether the bets in question are winning or losing bets. We can show that the bookie is able to exploit this asymmetry in constructing a Dutch book strategy for Beauty, unless her credence in Heads upon awakening is one-third. The argument requires the following premises, in addition to the Restricted Principal Principle:

> *Prior Admissibility.* In any Sleeping Beauty problem, on the evening before the first experimental awakening, Beauty knows that she has no information that is inadmissible in relation to any hypothesis in *S*.
>
> *Dutch Book Strategy Invulnerability.* If an agent is fully rational, then he or she is not vulnerable to any legitimate Dutch book strategy.
>
> *Every Awakening Legitimacy.* In the context of a Sleeping Beauty problem, a Dutch book strategy is legitimate so long as it involves offering Beauty the same bets every time she awakens, and so long as, in other respects, the protocol followed by the bookie is not affected by any information that Beauty lacks.
>
> *Awakening Independence.* In any Sleeping Beauty problem, since Beauty's awakenings are all mutually indistinguishable, her credences will be the same each time she awakens.
>
> *Fairness.* If a fully rational agent is offered a bet whose causal expected utility is zero, then he or she will regard this bet as fair.[13]

Let x be Beauty's credence in Heads upon awakening on the first morning of the experiment. From Awakening Independence, it follows that her credence in Heads is x each time she awakens. Assume, first, that x is greater than one-third. In this case, the Dutch book strategy to which Beauty is vulnerable will consist in offering Beauty a bet on Tails on the evening before the first experimental awakening, and then a bet on Heads each time she awakens during the experiment. The first bet, on Tails, will cost \$2, and will pay \$4 if the coin comes up tails. If we define the *payoff* of a bet as what it pays minus what it costs, then the payoff of this bet will be \$2 if Tails and −\$2 if Heads. Then, every time she awakens, she is offered a follow-up bet on Heads with a payoff of $\$(3-3x)$ if Heads and $-\$3x$ if Tails.

13. As Arntzenius 2002 shows, the Dutch book argument for the one-third solution requires that expected utilities be understood causally rather than evidentially.

From the Restricted Principal Principle, together with Prior Admissibility, Beauty's credence in Tails, on the evening before the first awakening, must be one-half, and so the causal expected utility of the bet on Tails will be zero. And so it follows, from Fairness, that she will regard this bet as fair. And since, each time she awakens, her credence in Heads is x, she will likewise regard the follow-up bets on Heads as fair. Now if the coin comes up heads, then the bet on Tails will have a payoff of $-\$2$. And she will be awoken exactly once, and so she will be offered only one bet on Heads, and its payoff will be $\$(3 - 3x)$. But since we are assuming that x is greater than one-third, this payoff must be less than \$2. And so the net payoff of the two bets will be negative. If, on the other hand, the coin comes up tails, then the payoff of the bet on Tails will be \$2. And she will be awoken twice, and so on two occasions she will be offered a bet on Heads. Each will have a payoff of $-\$3x$, for a total of $-\$6x$. And since x is greater than one-third, the combined loss on the two Heads bets will exceed the win on the Tails bet.

Now assume that x is less than $1/3$. In this case, the Dutch book strategy will consist in first offering Beauty a bet on Heads on the evening before the first awakening, with a payoff of \$2 if Heads and $-\$2$ if Tails. Then, every time she awakens, she will be offered a bet on Tails with a payoff of $\$3x$ if Tails and $\$(3x - 3)$ if Heads. It can easily be verified that Beauty will regard each of these bets as fair, and yet together they result in a sure loss.

It follows from Every Awakening Legitimacy that the Dutch book strategy that figures in the above argument is legitimate, and hence that if Beauty's credence in Heads, upon awakening, is anything other than one-third, then she is vulnerable to a legitimate Dutch book strategy. And so we can infer, from Dutch Book Strategy Invulnerability, that unless Beauty's credence in Heads is one-third, she is irrational.

On the basis of the same premises, we can derive the GTP. In any Sleeping Beauty problem defined by a partition S, suppose that when Beauty first awakens she violates the Generalized Thirder Principle. It follows from Awakening Independence that she has the same credences every time she awakens; let P represent these credences. Since she violates the GTP, there must be two hypotheses, i and j, belonging to S, such that $N(i)$ and $N(j)$ are integers, and either

$$P(i)N(j)Ch(j) < P(j)N(i)Ch(i)$$

or

$$P(i)\,N(j)\,Ch(j) > P(j)\,N(i)\,Ch(i)$$

If the former inequality obtains, let A be hypothesis i and let B be hypothesis j; otherwise, let A be hypothesis j and let B be hypothesis i. Thus,

$$P(A)\,N(B)\,Ch(B) < P(B)\,N(A)\,Ch(A)$$

and it follows that $P(B) > 0$. It also follows that

$$\frac{P(A)\,Ch(B)}{N(A)} < \frac{P(B)\,Ch(A)}{N(B)}$$

And hence it follows that for some positive value of Δ,

$$P(A)\left(\frac{Ch(B)}{N(A)} + \Delta\right) = P(B)\left(\frac{Ch(A)}{N(B)} - \Delta\right) (^{*})$$

We can now show that Beauty will be vulnerable to the following Dutch book strategy. On the evening before the first awakening, she is offered two bets. The first is a bet on the disjunction of A and B and the second is a bet on A that is conditional on the disjunction of A and B. Furthermore, every time she awakens, she is offered a follow-up bet on B that is conditional on the disjunction of A and B. The following table indicates the payoffs of each of these three bets, depending on whether A is true, or B is true, or neither A nor B is true.

	If A	If B	If neither A nor B
Bet 1	$\$(1 - Ch(A \text{ or } B))\Delta$	$\$(1 - Ch(A \text{ or } B))\Delta$	$-\$Ch(A \text{ or } B)\Delta$
Bet 2	$\$Ch(B)$	$-\$Ch(A)$	\$0
Follow-up Bets	$-\$(Ch(B)/N(A) + \Delta)$	$\$(Ch(A)/N(B) - \Delta)$	\$0

From the Restricted Principal Principle, together with Prior Admissibility, it follows that, on the evening before the first awakening, Beauty's credences in A and in B and in their disjunction must be equal to the objective chances of these propositions. Hence, by Fairness, she will regard the first two bets as fair. And, given equation $(^{*})$, every time she awakens and is offered a follow-up bet, she will likewise regard this bet as fair.

But the combined payoffs of the bets she is offered is sure to be negative. For suppose A is true. In this case, there will be $N(A)$ follow-up bets, and so we can infer from the above payoff matrix that the total payoff will be

$$\$((1 - Ch(A \text{ or } B))\Delta + Ch(B) - N(A)(Ch(B)/N(A) + \Delta))$$
$$= \$(1 - N(A) - Ch(A \text{ or } B))\Delta$$

Since we know that $P(B) > 0$, it follows that $Ch(B) > 0$, and hence that $Ch(A \text{ or } B) > 0$. Consequently, since $N(A)$ must be at least one, $1 - N(A) - Ch(A \text{ or } B)$ must be negative. And so the total payoff must be negative. Suppose, next, that B is true. In this case, there will be $N(B)$ follow-up bets, and so the total payoff will be

$$\$((1 - Ch(A \text{ or } B))\Delta - Ch(A) + N(B)(Ch(A)/N(B) - \Delta))$$
$$= \$(1 - N(B) - Ch(A \text{ or } B))\Delta$$

And since $N(B)$ must be at least one, the total payoff must be negative. Suppose, finally, that neither A nor B is true. In this case, the payoff of the first bet will be $-\$Ch(A \text{ or } B)\Delta$, and all the other bets will be called off. And so, again, the net payoff will be negative. Thus, in each case, Beauty will suffer a loss if she accepts all these bets.

The strategy described above is legitimate, according to Every Awakening Legitimacy. And so we may conclude that if Beauty violates the GTP, then she is vulnerable to a legitimate Dutch book strategy, and hence, by Dutch Book Strategy Invulnerability, that she is irrational.

It seems, therefore, that the one-third solution and the GTP are supported by a diachronic Dutch book argument. The situation, however, as we will see in the next section, is more complicated. And much of this complication arises from the conflict between the GTP and CA.

4. The Relevance of Dutch Book Arguments

In this section, I will begin by arguing that the diachronic Dutch book arguments for the one-third solution and for the GTP are much stronger than the Dutch book argument for the one-half solution. I will then argue, however, that all of these arguments are questionable, and that none is as strong as the diachronic Dutch book argument for CA. I will conclude with a discussion of *synchronic* Dutch book arguments, and I will argue that such arguments likewise support CA over the GTP.

4.1. Why Diachronic Dutch Book Arguments Support the Thirders' View over the Halfers'

There are some who would reject the diachronic Dutch book argument for the one-third solution considered in the previous section, on the ground that it relies on the following premise:

> *Every Awakening Legitimacy*. In Sleeping Beauty contexts, a diachronic Dutch book strategy is legitimate so long as the same bets are offered to Beauty every time she awakens and so long as, in other respects, the protocol followed by the bookie is not affected by any information that Beauty lacks.

This is a premise that one could reject with some plausibility. For one might plausibly hold that a legitimate Dutch book strategy cannot exploit *any* information that the agent lacks. But if, in the original Sleeping Beauty problem, the bookie offers Beauty a bet on Heads (or on Tails) every time she awakens, then the number of bets that are offered to Beauty will depend on whether Heads or Tails is true. Hence, since Beauty doesn't know whether Heads or Tails is true, a Dutch book strategy in which Beauty is offered a bet on Heads (or on Tails) every time she awakens would appear to exploit information that Beauty lacks. And so one might plausibly argue that such a strategy is illegitimate, and hence that Every Awakening Legitimacy is false.

One might instead hold that in the context of a Sleeping Beauty problem, a legitimate Dutch book strategy is one in which the number of times a given bet is offered to Beauty does not depend on how many times Beauty awakens or on any other information that Beauty lacks. This would be true, for example, if Beauty were offered a fixed set of bets on one and only one occasion when she awakens. Hence, in place of Every Awakening Legitimacy, one might propose the following:

> *Single Awakening Legitimacy*. In Sleeping Beauty contexts, a Dutch book strategy is legitimate so long as, regardless of how many times Beauty awakens, she is offered a fixed set of bets on exactly one occasion when she awakens and so long as, in other respects, the protocol followed by the bookie is not affected by any information that Beauty lacks.

And this assumption, one might hold, supports not the one-third solution but rather the one-half solution.[14] For suppose that when Beauty awakens, her credence in Heads, which we may label x, is less than

14. See Hitchcock 2004.

one-half. In this case, she will be vulnerable to the following strategy. On the evening before the first awakening, she is offered a bet that costs \$.5 and pays \$1 if Heads. And on exactly one occasion when she awakens, she is offered a follow-up bet that costs $\$(1 - x)$ and pays \$1 if Tails. She will regard both these bets as fair, and yet their total payoff will be negative. Suppose, on the other hand, that x is greater than one-half. In this case, Beauty will be vulnerable to a parallel strategy in which, on the evening before the first awakening, she is offered a bet that costs \$.5 and pays \$1 if Tails, and on exactly one occasion when she awakens, she is offered a follow-up bet that costs $\$x$ and pays \$1 if Heads. Once again, she will regard both bets as fair, and their total payoff will be negative. Hence, it would seem that, assuming Single Awakening Legitimacy, the only way Beauty can avoid being vulnerable to a legitimate Dutch book strategy is for her credence in Heads upon awakening to be one-half.

One might think, therefore, that the Dutch book argument for the one-third solution is not very strong. For one might think that the assumption of Every Awakening Legitimacy is no more plausible than the assumption of Single Awakening Legitimacy, and one might think that the latter supports not the one-third solution but rather the one-half solution. But this would be a mistake. Properly understood, Single Awakening Legitimacy, like Every Awakening Legitimacy, supports the one-third solution. For what the above argument shows is that if Beauty is to avoid being vulnerable to a Dutch book strategy in which she is offered a follow-up bet on exactly one occasion when she awakens, then *when she is offered this follow-up bet* her credence in Heads must be one-half. But this does not entail that *when she awakens* her credence in Heads must be one-half. For, in the scenario in which Beauty is offered a follow-up bet on exactly one occasion when she awakens, Beauty's being offered such a bet on a given awakening will provide evidence for Heads. And so, if her credence in Heads is to be one-half upon being offered this bet, it must be less than one-half upon awakening.[15]

Indeed, in the scenario in which Beauty is guaranteed to be offered a follow-up bet only once, and in which she is aware of this protocol, we can show that her credence in Heads, immediately after

15. In Arntzenius 2002, and in Bradley and Leitgeb 2006, it is argued that in Sleeping Beauty problems, Beauty's betting odds may come apart from her credences. Here I am making the different point that even if her betting odds should match the credences she has upon being offered the bet, they may differ from the credences she had prior to being offered the bet.

awakening, must be one-third.[16] Let P be Beauty's credence function immediately after awakening. And let B be the centered proposition that can be expressed thus: "I am offered a bet on this awakening." If Heads is true, then Beauty will be offered a bet on the only occasion when she awakens. And so, when she awakens, her credence that she will be offered a bet on this awakening, conditional on Heads, must be one. That is, $P(B \mid \text{Heads}) = 1$. But if Tails is true, then Beauty will be offered a bet on only one of the two occasions when she awakens. Hence, by a plausible principle of indifference, when Beauty awakens, her credence that she will be offered a bet on this awakening, conditional on Tails, should be one-half.[17] That is, $P(B \mid \text{Tails}) = 1/2$. And we have seen that if Beauty is to avoid being vulnerable to a Dutch book strategy of the kind we are now considering, then upon being offered the follow-up bet, her credence in Heads must be one-half. And so $P(\text{Heads} \mid B) = 1/2$. Therefore, by Bayes' Theorem,

$$\frac{P(\mathit{Heads})\,P(B \mid \mathit{Heads})}{P(B)}$$

$$= \frac{P(\mathit{Heads})\,P(B \mid \mathit{Heads})}{P(\mathit{Heads})\,P(B \mid \mathit{Heads}) + (1 - P(\mathit{Heads}))\,P(B \mid \mathit{Tails})} = \frac{1}{2}$$

And from the above equation, together with the known values of $P(B \mid \text{Heads})$ and $P(B \mid \text{Tails})$, we can infer that $P(\text{Heads}) = 1/3$. Therefore, assuming Single Awakening Legitimacy, the only way for Beauty to avoid being vulnerable to a legitimate Dutch book strategy is for her credence in Heads, upon awakening, to be one-third.

The above argument can be generalized to show that, assuming Single Awakening Legitimacy, if Beauty is to avoid being vulnerable to a legitimate Dutch book strategy, she must satisfy the GTP. For it follows from Single Awakening Legitimacy that in order to avoid being vulnerable to such a strategy, Beauty's credences in the hypotheses in S, upon awakening *and being offered a follow-up bet*, must be equal to her credences in these hypotheses on the evening before the first experimental awakening—for otherwise she will have different betting odds

16. If Beauty is unaware of the causal setup that determines how many times she is offered a follow-up bet, and hence she is unaware that she is guaranteed to be offered such a bet only once, then we cannot derive this result. But in this case, the bookie will be exploiting information that Beauty lacks.

17. We can obtain this result from FSBI so long as we define the partition S not as {Heads, Tails} but rather as {Heads, Tails and the follow-up bet is offered on Monday, Tails and the follow-up bet is offered on Tuesday}.

on the two occasions, which a bookie will be able to exploit. And by the Principal Principle, together with Prior Admissibility, her initial credences in these hypotheses must equal their objective chances. Hence, where P represents Beauty's credences upon awakening, and B is the centered proposition that can be expressed *I am offered a bet on the present awakening*, it follows that for any hypothesis h, in S, $P(h \mid B) = Ch(h)$. And where w_1 is the centered proposition that can be expressed *I have awoken only once*, it follows from a plausible principle of indifference that, for any h in S, $P(B \mid h) = P(w_1 \mid h)$. Hence, by Bayes' Theorem, $P(h \mid w_1) = P(h \mid B)$. Therefore, $P(h \mid w_1) = Ch(h)$. And from this conclusion, as we saw in section 3.2, we can derive the GTP. Thus, the GTP can be derived on the basis of either Every Awakening Legitimacy or Single Awakening Legitimacy.

4.2. Why Diachronic Dutch Book Arguments Support CA over the GTP

But there's trouble. For the very same premises that underlie the two Dutch book arguments we have considered for the GTP give rise to Dutch book arguments for the conflicting principle of Countable Additivity. The basic problem is that if Beauty violates CA, then her credences will be *nonconglomerable*, in the sense that her unconditional credence in some proposition will exceed the upper bound for her credence in this proposition conditional on every element in some countable partition.[18] And in virtue of having nonconglomerable credences, she will be vulnerable to a legitimate Dutch book strategy, on either of the two conceptions of such a strategy we have considered.

For an illustration of this problem, consider the following case:

> *Magical Mystery Number.* This is a Sleeping Beauty problem in which the number of times Beauty awakens is determined by a random variable A, the possible values of which coincide with the positive integers. For any positive integer n, $Ch(A = n) = 1/2^n$, and the number of times Beauty awakens, given that $A = n$, is 2^n. B is another random variable, independent of A, such that for any positive integer n, $Ch(B = n)$ likewise equals $1/2^n$. While Beauty is sleeping, before her first experimental awakening, a fair coin is tossed, and if it comes up heads, the value of A is written onto

18. The connection between countable additivity and conglomerability was first discussed in de Finetti 1972. See also Schervish, Seidenfeld, and Kadane 1984.

> a piece of paper, but if it comes up tails, then the number that is written onto the piece of paper is the value of B.[19]

Let X be whatever number is written onto the piece of paper. Let Heads and Tails denote the propositions that the coin comes up heads, and tails, respectively. And let P represent Beauty's credences when she first awakens. Since the number of times Beauty awakens is independent, both of the value of B and of the outcome of the coin toss, we may infer that when she first awakens, she knows that she has no information that is inadmissible in relation to either of these outcomes. Hence, by the Restricted Principal Principle, $P(\text{Heads}) = Ch(\text{Heads}) = 1/2$, and for any positive integer value of n, $P(B = n) = Ch(B = n) = 1/2^n$. And since the value of B is independent of the outcome of the coin toss,

$$P(B = n \mid \text{Tails}) = 1/2^x \text{ for all } n \in Z+ \ (^{**})$$

By contrast, the number of times Beauty awakens does depend on the value of A and it is inversely proportional to the objective chance of this value. Hence, it follows from the GTP that Beauty must have equal credence in each of the possible values of A. And so, by finite additivity, her credence in each must be zero. Consequently,

$$P(A = n \mid \textit{Heads}) = 0 \text{ for all } n \in Z+ \ (^{***})$$

But Beauty knows that if the coin comes up tails, then $X = B$, whereas if it comes up heads, then $X = A$. Hence, from (**) and (***) we can infer that for any positive integer value of n, $P(X = n \mid \text{Tails}) = 1/2^x$ and $P(X = n \mid \text{Heads}) = 0$, from which it follows that $P(\text{Heads} \mid X = n) = 0$. And so Beauty's credence in Heads, conditional on each of the possible values of X, will be significantly less than her unconditional credence in Heads, which is one-half. And so her credences will be nonconglomerable.

As a result, she will be vulnerable to a Dutch book strategy. For since, upon first awakening, Beauty's credence in Heads is one-half, she will regard $1 as a fair price to pay for a bet that returns $2 if Heads. But after she learns the value of X, her credence in Heads will be zero, and so she will regard $2 (or at least any quantity less than $2) as a fair price

19. This example is based on one discussed in de Finetti 1972, 205; and in Howson 2008, 12. Note that the Magical Mystery Number case is very similar to that of Sleeping Beauty in St. Petersburg, except that in the former case the value of the random variable that determines how many times Beauty awakens is guaranteed to be finite. This stipulation is not essential to the argument but simplifies matters considerably.

to pay for a bet that returns $2 if Tails. And yet accepting both bets would result in a sure loss.

The Dutch book strategy to which Beauty will be vulnerable can be spelled out in either of two ways. In the first version of the strategy, which will count as legitimate according to Single Awakening Legitimacy, Beauty is offered the sequence of two bets just described on exactly one of the mornings when she awakens. And in the second version of the strategy, which will count as legitimate according to Every Awakening Legitimacy, she is offered this sequence of bets every time she awakens. Either way, if she accepts all the bets she is offered, she will suffer a sure loss. Thus, in Magical Mystery Number, if Beauty satisfies the GTP, then she will be vulnerable to a legitimate Dutch book strategy on either of the two conceptions we have considered. Note, further, that since the number of times Beauty awakens is guaranteed to be finite, she will be offered only finitely many bets regardless of which version of the strategy is employed. And so neither version can be rejected on the ground that it involves offering Beauty infinitely many bets.

Thus, if we accept either Every Awakening or Single Awakening Legitimacy, together with Diachronic Dutch Book Invulnerability, then we can infer that rationality requires that one satisfy both CA and the GTP, and hence that there are circumstances in which agents are under conflicting rational requirements. Therefore, we must either accept the possibility of rational dilemmas, or else we must reject the premises that underlie the Dutch book arguments for the one-third solution.

Moreover, there is independent reason to be skeptical of the latter premises. For if we accept either Every or Single Awakening Legitimacy, then we must claim that a Dutch book strategy can be legitimate even when the agent loses track of his or her location between the time the first bet is offered and the time when subsequent bets are offered. But there is reason to doubt that such a strategy is legitimate. The basic problem is this. We know that one can rationally violate reflection in relation to a proposition when one rationally expects to lose track of information that is relevant to this proposition.[20] For example, let p be the proposition that it will snow on New Year's Day. Suppose that one's initial credence in p was .5 (since that is the known frequency of snow on New Year's Day), but that after reading the weather report one's credence increases to .7. Suppose, however, that one rationally expects to forget the weather report without gaining any other relevant information in the interim. In

20. See Christensen 1991.

this case, one might rationally expect one's credence in *p* to return to .5. Hence, one's current credence in *p*, conditional on one's future credence in *p* being .5, will exceed .5, and so one will violate reflection. And because one violates reflection, one will be vulnerable to a Dutch book strategy.[21] Note, further, that in order to be vulnerable to this Dutch book strategy, it suffices that one *expects* to lose track of relevant information: it doesn't matter whether such a loss actually occurs. But since the expectation that one will lose track of relevant information can be perfectly rational, it follows that vulnerability to a diachronic Dutch book, in a context in which one expects to lose track of relevant information, does not indicate that one is irrational.

Hence, in order to defend the claim that the Dutch book strategies that figure in the arguments for the one-third solution and the GTP are legitimate, one would have to maintain that, in Sleeping Beauty problems, where Beauty expects to lose track of her location between Sunday evening and Monday morning, she will not thereby lose any *relevant* information since the only information she will lose is self-locating. But such a move would itself undermine the legitimacy of the Dutch book strategies that figure in the arguments under consideration. For it has been shown elsewhere that, if one holds that, in Sleeping Beauty contexts, changes in self-locating information are not relevant to uncentered propositions (such as Heads), then one is committed to claiming that in such contexts, one can rationally violate reflection on Sunday evening.[22] Hence, one is committed to claiming that in such contexts, a fully rational agent can be vulnerable to a Dutch book strategy in which the first bets are offered on Sunday evening and subsequent bets are offered after the agent loses track of his or her location.

21. This is shown in van Fraassen 1984.

22. This is illustrated in the "black and white room" version of the Sleeping Beauty problem, discussed in Meacham 2008, 263, a case Meacham attributes to David Manley. In this case, whether Beauty awakens once or twice is determined by the toss of a nickel. If $\text{Tails}_{\text{nickel}}$, then Beauty will awaken twice: in a black room on Monday, and in a white room on Tuesday. But if $\text{Heads}_{\text{nickel}}$, then Beauty will awaken only once, on Monday, in a room whose color depends on the toss of a dime ($\text{Heads}_{\text{dime}}$: black, $\text{Tails}_{\text{dime}}$: white). According to the view that only uncentered information is relevant to uncentered propositions, if Beauty awakens in a white room, the only information she learns that is relevant to $\text{Heads}_{\text{nickel}}$ is that $\sim(\text{Heads}_{\text{nickel}}\ \&\ \text{Heads}_{\text{dime}})$. And if she awakens in a black room, the only relevant information she learns is that $\sim(\text{Heads}_{\text{nickel}}\ \&\ \text{Tails}_{\text{dime}})$. Either way, she will be excluding a possibility in which $\text{Heads}_{\text{nickel}}$ without excluding any possibility in which $\text{Tails}_{\text{nickel}}$, and so her credence in $\text{Heads}_{\text{nickel}}$ will decrease. And since she can anticipate this on Sunday evening, she will then violate reflection.

Thus, whatever one may think about the relevance of self-locating information to uncentered propositions, it is very difficult to maintain that an agent is irrational whenever he or she is vulnerable to a Dutch book strategy in which bets are offered over a period during which he or she expects to lose track of such information. And since Single and Every Awakening Legitimacy both imply that vulnerability to such strategies indicates irrationality, we should be skeptical of these principles.

This ground for skepticism does not, however, extend to the following principle:

> *Learning Legitimacy.* A diachronic Dutch book strategy is legitimate if, between the times when the first and last bets are offered to the agent, the only relevant change in his or her epistemic situation is that he or she gains additional information.

This principle does not support the Dutch book arguments for the one-third solution or for the GTP, since these arguments involve Beauty's being offered a sequence of bets over a period in which she loses track of the time. This principle does, however, support the diachronic Dutch book argument for CA. For, as we have seen, if an agent violates CA, then he or she will be vulnerable to a Dutch book strategy that involves offering him or her bets at two times between which he or she simply learns which proposition within some partition is true. It seems, therefore, that CA is supported by a diachronic Dutch book argument that is more compelling than any such argument for the one-third solution or for the GTP.

4.3. Why Synchronic Dutch Book Arguments Support CA over the GTP

The relevance of *synchronic* Dutch book arguments to the conflict between CA and the GTP is somewhat controversial. Since there is a well-known synchronic Dutch book argument for CA,[23] but none for the one-third solution or for the GTP, considerations of synchronic Dutch book coherence would seem to favor CA. But in the present context, this claim is contentious since some very prominent thirders have maintained that the synchronic Dutch book argument for CA has no force. According to Arntzenius, Elga, and Hawthorne 2004, 262, this Dutch book argument is undermined by the following case:

> *Satan's Apple* (synchronic version). Satan has cut a delicious apple into infinitely many pieces, labeled by the natural numbers. For each piece of

23. See Bartha 2004 and Williamson 1999.

> the apple, Eve must choose whether to take it or decline it, and she must make all these choices simultaneously. If she takes merely finitely many of the pieces, then she suffers no penalty. But if she takes infinitely many of the pieces, then she is expelled from the Garden for her greed. Either way, she gets to eat whatever pieces she has taken.

For any given piece *a*, of the apple, no matter which other pieces Eve chooses to take, she would be better off taking these other pieces together with *a* than taking these other pieces alone. Thus, for any given piece *a*, the option of taking *a* dominates the option of declining *a* with respect to the partition that consists in the various possible combinations of other pieces that Eve might take. What this shows is that we must not accept an unrestricted dominance principle, according to which one is rationally required to choose a given option whenever there is some partition with respect to which it dominates every alternative. For such a principle would have the absurd consequence that Eve is rationally required to take each and every piece of the apple, thereby being expelled from the Garden. Hence, "the first lesson of Satan's Apple," according to Arntzenius, Elga, and Hawthorne 2004, 264, is that in "infinite cases, rationality does not require one to choose one's dominant options."

Moreover, they argue that the synchronic Dutch book argument for CA implicitly assumes that rationality does require taking one's dominant options in infinite cases since this argument assumes that one is rationally required to accept a given bet *b* whenever, no matter which combination of other bets one accepts within some infinite set of bets, the expected value of taking these other bets together with *b* exceeds the expected value of taking these other bets alone. Hence, they conclude the synchronic Dutch book argument for CA rests on an illegitimate extension of dominance reasoning to infinite cases.

But this is the wrong lesson to draw from the case of Satan's Apple. For the dominance reasoning that yields the unacceptable conclusion that Eve is required to take every piece of the apple is flawed for reasons that have nothing to do with the infinite size of the relevant partition. It is generally agreed that one is rationally required to choose the dominant option, with respect to a given partition, only when the elements of this partition are *independent*, in some sense, of one's choice. According to causal decision theorists (among whom Arntzenius, Elga, and Hawthorne number themselves), the relevant kind of independence is *causal independence*.[24] And if one restricts dominance reasoning to

24. See Joyce 1999.

partitions whose elements are causally independent of one's choice, then this is ground enough for rejecting the dominance reasoning in the case of Satan's Apple. For in this case, for any given piece *a* of the apple, the partition with respect to which taking *a* dominates declining *a* is the partition among the possible combinations of other pieces Eve might take. But which such combination Eve will choose is not causally independent of her choice concerning *a*.

The reason for the lack of causal independence is this. In order to make a given choice rationally, one needs a conception of the relevant causal role of this choice. Now if one is faced with a set of simultaneous decision problems, where the relevant causal role of one's choice concerning any one of these problems is independent of one's choices concerning the others, then one can form an adequate conception of the relevant causal role of one's choice concerning any one of these problems without knowing what one's other choices will be. For example, in the kind of choice situation that figures in standard synchronic Dutch book arguments, an agent is presented with a set *S* of bets, and he or she must simultaneously choose, for each bet in *S*, whether to accept it or decline it. Here, the only relevant causal role of his or her choice concerning any one of these bets *b* is in determining the payoff from that bet: if he or she declines *b*, the payoff from *b* will be $0, whereas if he or she accepts *b*, then the payoff will be either positive or negative, depending on whether *b* is a winning or losing bet. And the agent's payoff from *b* is not affected by his or her choices concerning any of the other bets in *S*. Hence, the relevant causal role of his or her choice concerning *b* is unaffected by his or her choices concerning the other bets. Therefore, the agent can form an adequate conception of the causal role of accepting any given bet without knowing what other bets he or she will accept. And so he or she can rationally make his or her choices concerning the various bets independently.

But the situation is very different in the case of Satan's Apple. Here, the most significant effect of Eve's choices concerning the pieces of the apple is that these choices together determine whether she is expelled from the Garden: for she will be expelled if and only if she takes infinitely many pieces. Hence, if she chooses to take infinitely many pieces, then for each of the pieces she chooses to take, her choosing to take this piece will be part of what causes her to be expelled from the Garden, whereas if she chooses to take only finitely many pieces, then for each piece she chooses to take, her choosing to take this piece will have no such causal role. Hence, the relevant causal role of Eve's choice concerning any one of the pieces of the apple depends on what other pieces she takes. And so Eve cannot form an adequate conception of this role without knowing what her other choices will be. Consequently, Eve cannot rationally choose whether to take the various pieces independently. If she is rational, she must make these choices together, by way of choosing a *total strategy* that specifies which pieces of the apple to take and which to decline.

But if she makes her choices in this way, then her choices concerning the individual pieces of the apple will not be causally independent of one another: for any piece *a*, had Eve made a different choice concerning *a*, she would have done so by way of choosing a different total strategy that might well have involved making different choices concerning the other pieces of the apple. Therefore, for any piece *a*, the partition of combinations of choices Eve might make concerning the other pieces of the apple will not be causally independent of her choice concerning *a*. But it is precisely this partition in relation to which taking *a* dominates declining *a*. Therefore, since taking *a* does not dominate declining *a* with respect to a partition that is causally independent of this choice, it follows that, so long as we restrict the dominance principle in the manner that causal decision theory requires, we can block the inference that Eve is rationally required to take each piece of the apple.

Thus, in the case of Satan's Apple, the failure of dominance can be explained purely in terms of the failure of causal independence. Accordingly, this case does not show that "in infinite cases, rationality does not require one to choose one's dominant options" (Arntzenius, Elga, and Hawthorne 2004, 264).[25] And so this case does not undermine the Dutch book argument for CA.

5. Conclusion: A Case for Rational Dilemmas

I argued, in section 3, that the main arguments for the one-third solution support the GTP. And I argued, in section 4, that the strongest Dutch book arguments, both diachronic and synchronic, support not the GTP

25. There are other cases, however, such as the two-envelope paradox, which may call into question the applicability of dominance reasoning in infinite cases; see Dietrich and List 2005. These cases, however, assume that a fully rational agent can have unbounded utilities. And this is not a legitimate assumption in the context of the debate over the validity of infinite Dutch book arguments. For, as McGee 1999 demonstrates, if an agent has unbounded utilities, then he or she will be vulnerable to a synchronic Dutch book involving infinitely many bets. Hence, anyone who accepts the validity of infinite Dutch book arguments will deny that a rational agent can have unbounded utilities.

but rather CA. In light of these arguments, how are we to respond to the conflict between the two principles?

One possibility is to reject the GTP. If one does this, then one must claim that each of the arguments we have considered for the GTP has at least one false premise. Since several of the arguments for the GTP appeal, implicitly or explicitly, to a finitistic principle of indifference, one could reject these arguments so long as one denies that such indifference principles apply even in finitistic cases. One might claim instead that when Beauty awakens, she could rationally have more credence in *Tails and Monday* than in *Tails and Tuesday*, or vice versa, or else one might claim that she should not have precise credences in these possibilities at all.[26] Similarly, since several arguments for the GTP appeal to a principle of conditionalization that applies to centered evidence, one could reject these arguments if one held that conditionalization properly applies only to uncentered evidence.[27] These moves, however, involve considerable costs.[28] More generally, given the wide range of arguments for the one-third solution, and given that several of these arguments appear to rest on plausible premises, many people would prefer to avoid having to reject all these arguments.

A second possibility is to reject CA. But such a move will likewise involve considerable costs. First, as I have argued, if one rejects CA, then one must deny that diachronic Dutch book arguments are valid within contexts in which the only relevant change in the agent's situation is that he or she gains information. And if one denies this, then it may be difficult to plausibly endorse *any* diachronic Dutch book arguments, including the argument for conditionalization. Second, if one rejects CA, then one must deny that synchronic Dutch book arguments are valid when they involve infinitely many bets. And, as I have argued, it is hard to justify restricting Dutch book arguments to finite contexts. And so if one denies that they are valid in infinite contexts, then it may be difficult to plausibly endorse *any* synchronic Dutch book arguments, including the arguments

26. For a defense of the latter view, see Weatherson 2005. See also Arntzenius 2002.

27. See Meacham 2008, where Meacham argues that our credences in uncentered propositions should be affected only by changes in our uncentered evidence. Mere changes in centered evidence should affect only our credences in centered propositions.

28. For a discussion of the problems that arise if one denies that one should conditionalize on centered evidence, see Bostrom 2002 and Dorr n.d. And for a discussion of problems that arise if one denies that one should have equal credences, or precise credences, in the awakening possibilities associated with a given hypothesis, see Elga 2004 and White forthcoming, respectively.

for *finite* additivity, normalization, and nonnegativity.[29] Hence, rejecting the Dutch book argument for CA may require rejecting the Dutch book arguments for all the standard Bayesian principles. Third, as noted earlier, if one rejects CA, then one must reject the principle of Conglomerability, a principle that has a great deal of plausibility.[30] And finally, if one rejects CA, then one must renounce the important results that depend on this principle in a number of domains, including statistics, confirmation theory, and the philosophy of science.[31]

Fortunately, there is a third alternative. We can accept both CA and the GTP, for these two principles are not logically incompatible. The two principles conflict, in the sense that there are situations in which it is impossible to satisfy both. But so long as we grant that agents can be under conflicting rational requirements, and hence that agents can face rational dilemmas, we can accept both CA and the GTP.

Thus, given the various motivations for CA and for the GTP, the conflict between these principles provides some ground for accepting the possibility of rational dilemmas.[32] It also provides insight into how such dilemmas are possible. For there is an important difference between the motivations for CA and the motivations for the GTP. The arguments for CA aim to show that one must satisfy CA on pain of incoherence: the Dutch book arguments aim to show that if someone violates CA, then he or she evaluates bets in an incoherent manner, and the argument from Conglomerability aims to show that if someone violates CA, then there is an incoherence between his or her conditional and unconditional credences. By contrast, the strongest arguments for the GTP aim to show not that an agent who violates this principle is incoherent, but rather that such an agent fails to respond appropriately to his or her evidence. Thus, the argument from the Principal Principle aims to show that in any Sleeping Beauty problem, if one violates the GTP, then upon learning that one has awoken only once, one must either fail to respond appropriately to this knowledge or else fail to respond appropriately to one's knowledge of objective chances. The *De Re* Conditionalization argument aims to show that if one violates the GTP, then one must fail to respond

29. See Skyrms 1984.

30. For an important recent discussion and defense of conglomerability, see Easwaran 2008.

31. See Bartoszynski and Niewiadonska-Bugaj 1996, Earman 1992, and Kelly 1996.

32. For other arguments for the existence of rational dilemmas, see Priest 2002. And for closely related arguments, see Christensen 2007.

appropriately to the *de re* information one acquires upon awakening. And the Frequency and Hypothetical Priors arguments both aim to show that if one violates the GTP, then one's credences must differ from those one would have in a situation in which one's evidence is not relevantly different. And, plausibly, such a disparity indicates a failure to respond appropriately to one's evidence.

Suppose we are persuaded by these arguments, and we conclude internal coherence requires satisfying CA, while responding appropriately to one's evidence requires satisfying the GTP. In this case, the lesson to be drawn from the conflict between CA and the GTP is that there are contexts in which evidential considerations and considerations of coherence pull in opposing directions. And if we grant that full rationality requires both internal coherence and responding appropriately to one's evidence, then we should conclude that there are contexts in which full rationality is impossible, as in these contexts, to quote Henry Sidgwick, "reason is divided against itself."[33]

References

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

———. 2003. "Some Problems for Conditionalization and Reflection." *Journal of Philosophy* 100: 356–70.

Arntzenius, Frank, Adam Elga, and John Hawthorne. 2004. "Bayesianism, Infinite Decisions, and Binding." *Mind* 113: 251–83.

Bartha, Paul. 2004. "Countable Additivity and the de Finetti Lottery." *British Journal for the Philosophy of Science* 55: 301–21.

Bartoszynski, Robert, and Magdalena Niewiadonska-Bugaj. 1996. *Probability and Statistical Inference.* New York: J. Wiley.

Bostrom, Nick. 2002. "Self-Locating Belief in Big Worlds: Cosmology's Missing Link to Observation." *Journal of Philosophy* 99: 607–23.

Bradley, Darren, and Hannes Leitgeb. 2006. "When Betting Odds and Credences Come Apart: More Worries for Dutch Book Arguments." *Analysis* 66: 119–27.

Christensen, David. 1991. "Clever Bookies and Coherent Beliefs." *Philosophical Review* 100: 229–47.

———. 2007. "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals." *Oxford Studies in Epistemology* 2: 3–31.

De Finetti, Bruno. 1972. *Probability, Induction and Statistics.* New York: J. Wiley.

Dietrich, Franz, and Christian List. 2005. "The Two-Envelope Paradox: An Axiomatic Approach." *Mind* 114: 239–48.

33. Sidgwick 1907, 508.

Dorr, Cian. 2002. "Sleeping Beauty: In Defence of Elga." *Analysis* 62: 292–96.

———. n.d. "A Challenge for Halfers." Unpublished manuscript.

Draper, Kai, and Joel Pust. 2008. "Diachronic Dutch Books and Sleeping Beauty." *Synthese* 164: 281–87.

Earman, John. 1992. *Bayes or Bust?* Cambridge, MA: MIT Press.

Easwaran, Kenny. 2008. "Foundations of Conditional Probability." PhD diss., University of California, Berkeley.

Elga, Adam. 2000. "Self-Locating Belief and the Sleeping Beauty Problem." *Analysis* 60: 143–47.

———. 2004. "Defeating Dr. Evil with Self-Locating Belief." *Philosophy and Phenomenological Research* 69: 383–96.

Hájek, Alan. 2003. "What Conditional Probabilities Could Not Be." *Synthese* 137: 273–323.

Hitchcock, Christopher. 2004. "Beauty and the Bets." *Synthese* 139: 405–20.

Horgan, Terry. 2004. "Sleeping Beauty Awakened: New Odds at the Dawn of the New Day." *Analysis* 64: 10–21.

———. 2008. "Synchronic Bayesian Updating and the Sleeping Beauty Problem: Reply to Pust." *Synthese* 160: 155–59.

Howard, John. 2006. "Countable Additivity and the Foundations of Bayesian Statistics." *Theory and Decision* 60: 127–35.

Howson, Colin. 2008. "De Finetti, Countable Additivity, Consistency and Coherence." *British Journal for the Philosophy of Science* 59: 1–23.

Joyce, James. 1999. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.

Kelly, Kevin. 1996. *The Logic of Reliable Inquiry*. New York: Oxford University Press.

Kolmogorov, Andrey. 1956 [1933]. *Foundations of the Theory of Probability*, trans. and ed. N. Morrison. 2nd English ed. New York: Chelsea.

Lewis, David. 2001. "Sleeping Beauty: Reply to Elga." *Analysis* 61: 171–76.

McGee, Vann. 1999. "An Airtight Dutch Book." *Analysis* 59: 257–65.

Meacham, Christopher. 2008. "Sleeping Beauty and the Dynamics of *De Se* Beliefs." *Philosophical Studies* 138: 245–69.

Monton, Bradley. 2002. "Sleeping Beauty and the Forgetful Bayesian." *Analysis* 62: 47–53.

Priest, Graham. 2002. "Rational Dilemmas." *Analysis* 62: 11–16.

Ross, Jacob. 2010. "Stalnaker's Sleeping Beauty Argument." Home page for Jacob Ross at the Web site for the University of Southern California. www.rcf.usc.edu/~jacobmro/.

Schervish, Mark, Teddy Seidenfeld, and Joseph Kadane. 1984. "The Extent of Nonconglomerability of Finitely Additive Measures." *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 66: 205–26.

Seminar, Oscar. 2008. "An Objectivist Argument for Thirdism." *Analysis* 68: 149–55.

Sidgwick, Henry. 1907. *The Methods of Ethics.* London: MacMillan.

Skyrms, Brian. 1984. *Pragmatics and Empiricism.* New Haven: Yale University Press.

Stalnaker, Robert. 2008a. *Our Knowledge of the Internal World.* Oxford: Oxford University Press.

———. 2008b. "Comment on 'Stalnaker on Sleeping Beauty'." *BrainPains.* http://brainpains.com/wordpress/?p=54 (accessed October 12, 2009).

Titelbaum, Michael. 2008. "The Relevance of Self-Locating Beliefs." *Philosophical Review* 117: 555–606.

van Fraassen, Bas. 1984. "Belief and the Will." *Journal of Philosophy* 81: 235–56.

Walley, Peter. 1991. *Statistical Reasoning with Imprecise Probabilities.* London: Chapman and Hall.

Weatherson, Brian. 2005. "Should We Respond to Evil with Indifference?" *Philosophy and Phenomenological Research* 70: 613–35.

———. n.d. "Stalnaker on Sleeping Beauty." Unpublished manuscript.

Weintraub, Ruth. 2004. "Sleeping Beauty: A Simple Solution." *Analysis* 64: 8–10.

White, Roger. Forthcoming. "Evidential Symmetry and Mushy Credences." *Oxford Studies in Epistemology.*

Williamson, Jon. 1999. "Countable Additivity and Subjective Probability." *British Journal for the Philosophy of Science* 50: 401–16.