A Nightmare for Lewisian Halfers

Lennart B. Ackermans

June 13, 2026

Abstract

What if Sleeping Beauty has dreams? I give an argument against the halfer position that only uses principles of rationality accepted by Lewisian halfers. While dreaming, Beauty's credence in heads is 1/2. After being woken up and updating on the evidence that she is woken up today, the credence becomes less than 1/2. For Lewisian halfers, this is a nightmare. If we add two plausible assumptions about Beauty's dreaming credences, we can derive the thirder solution. I also respond to and improve upon a new type of defence of Lewisian halfing by Schwarz (2025), but argue that it fails as well.

1. Introduction

> "Is it you, my Prince? You have waited a long while."
> The Prince, charmed with these words, and much more with the manner in which they were spoken, knew not how to show his joy and gratitude; he assured her that he loved her better than he did himself. […] He was more at a loss than she, and we need not wonder at it; she had had time to think of what to say to him; for it is evident (though history says nothing of it) that the good fairy, during so long a sleep, had given her very pleasant dreams.
>
> Charles Perrault, 1697, *Contes de ma mère l'Oye*
> Translated by Charles Welsh, 1901

Once upon a time there was a great lab for experimental philosophy, renowned across all human kingdoms. One day, researchers at this lab found the test subject of their dreams. Sleeping Beauty, a beautiful princess, was bestowed the gift of supreme rationality by a fairy. The experiment, to which Beauty consented, went as follows. On Sunday, Beauty was put to sleep by the researchers. On Monday, she was woken up briefly for a chat. (In some versions of the telling, she was told it was Monday after a while.) She was then put to sleep again, and her memories of what happened on that day were erased. The researchers then tossed a coin. If the coin landed tails, the researchers would wake her up on Tuesday to have an identical chat. If the coin landed heads, she would stay asleep on Tuesday. On Sunday, Beauty was told how the experiment would go. The researchers aimed to discover what Beauty's credence was in the coin landing heads after she woke up on Monday. Unfortunately, Beauty's true credence, as revealed to the researchers on that very Monday, was lost to history. Can we recover it from basic principles of rationality?

For a seemingly simple probability puzzle, the Sleeping Beauty problem has generated a surprising amount of controversy. The *thirders*, who account for the majority of published positions, argue that the solution is 1/3 (among these are Dorr, 2002; Elga, 2000; Horgan, 2004, 2007; Kim, 2022, 2024; Titelbaum, 2008). The *halfers* claim that the solution is 1/2. Among them are so-called *double halfers* (Bostrom, 2007; Briggs, 2010; Meacham, 2008; Pust, 2012) and *Lewisian halfers* (Bradley, 2011; Lewis, 2001; Schwarz, 2025).

Lewisian halfers (unlike double halfers) maintain that diachronic Bayesian conditionalization is appropriate when updating on certain types of self-locating evidence, such as which day it is. Hence, when Beauty is told it is Monday, she should apply Bayesian conditionalization to this evidence. Thirders and Lewisian halfers agree that this requires her to increase her credence in heads. Since Lewisian halfers claim that Beauty's credence in heads is 1/2 before being told it is Monday, they are committed to it becoming greater than 1/2 afterwards. Since this result is a bitter pill to swallow, Lewisian halfing is regarded as implausible by many – while still a coherent position. My first aim in this article is to show that the situation is worse: Lewisian halfing in its original version is inconsistent.

The original motivation for Lewisian halfing is that between Sunday evening and Monday after being woken up, Beauty has not received new relevant information. Since her credence in heads on Sunday is uncontroversially 1/2, Beauty must maintain a credence of 1/2 on Monday after waking up. This reasoning uses a special case of Bayesian conditionalization that we might call the *Principle of Irrelevant Evidence*, according to which only new relevant evidence can change one's credences. For Lewisian halfers, the pull of this argument is strong enough to override any concerns about the implausible credence greater than 1/2 after Beauty is told it is Monday.

I give an argument against the halfer position that relies only on the Principle of Irrelevant Evidence, Bayesian conditionalization, and some plausible premises that Lewisian halfers are unlikely to reject. Hence, they must either accept thirding, switch to another branch of halfing such as double halfing, or provide a fundamentally new defence of Lewisian halfing.

In the central argument offered in section 2, we imagine that Beauty has a lucid dream on both Monday and Tuesday. During the dream, her relevant evidence with respect to the coin landing heads has not changed since Sunday, so her credence in heads should remain unchanged. After she is woken up, she uses Bayesian conditionalization on the evidence that the researchers wake her up today. I show that her credence in heads must now be less than 1/2. Other than Bayesian conditionalization, this argument only uses very plausible premises about Beauty's dreaming credences.

If we add two assumptions about Beauty's credences while sleeping, we can further derive the thirder position (section 3). This is a novel argument for thirding, similar in mathematical structure to existing arguments by Horgan (2004, 2007) and Milano (2022). Unlike these arguments, my argument uses only diachronic conditionalization, as opposed to synchronic conditionalization. This helps to avoid the objections by Pust (2008, 2013, 2014) against this use of synchronic updating.

My argument for thirding also bears similarity to arguments by analogy such as Arntzenius (2003), Dorr (2002), and Titelbaum (2013a). However, these arguments are more susceptible to the objection that they discuss versions of the problem that are disanalogous to the original Sleeping Beauty problem as discussed by Elga and Lewis (Bradley, 2003; Kim, 2021; Schwarz, 2025). One might similarly try to object to my argument that it is disanalogous. However, the original description does not preclude that Beauty has dreams – and she *does* have dreams in the fairy tale as told by Charles Perrault.[1] Hence, the objector who claims that Beauty does not have dreams is offering a version that is seemingly inconsistent with the original.

Nevertheless, the only plausible response for the Lewisian halfer is to claim that my version is disanalogous to the original, an objection to which I turn in section 4. I argue that this response requires rejecting evidentialism, the thesis that permissible credences are only constrained by evidence. Even so, such a defence of Lewisian halfing faces obstacles.

A fundamentally new type of defence of Lewisian halfing is offered by Schwarz (2025). However, Schwarz's argument is based on a version of the Sleeping Beauty problem distinct from those of Elga and Lewis, and it does not work for those versions. Moreover, it uses an updating principle based on maximizing *average expected accuracy* which I show has highly counterintuitive consequences.

In section 5, I improve upon Schwarz's defence of Lewisian halfing using a revised version of the average expected accuracy principle. I argue that this is the most plausible defence of Lewisian halfing, and it allows for a kind of disanalogy objection against this article's main argument. However, if we take the commitments of this type of Lewisian halfing seriously, a further plausible revision forces one to reject halfing.

The article thus refutes or casts serious doubt on all published defences of Lewisian halfing. It is now up to the Lewisians to reassess their position.

[1] As far as I am aware, Horgan (2007), a thirder, is the only author who claims that Beauty's sleep is "dreamless". But Horgan's description is not the seminal one. Schwarz (2025, p. 1085) acknowledges that Beauty "could be dreaming".

	Sunday	Monday	Tuesday	
			Heads	Tails
Early morning		Lucid dream	Lucid dream	
Afternoon	Experiment explained Put to sleep	Beauty woken up Put to sleep Memory erased	Asleep	Woken up Put to sleep

Table 1: Overview of the Sleeping Beauty experiment with dreams.

2. A refutation of traditional Lewisian halfing

A lucid dream is a dream in which you know you are dreaming. Let us suppose that Beauty has a lucid dream every night.[2] The events during the experiment are summarized in Table 1.

Given that Beauty is, as Lewis calls it, the "paragon of probabilistic rationality" (Lewis, 2001, p. 171), it is not a stretch to suppose that she is capable of rational probabilistic reasoning even during her dreams. (In the fairy tale, after all, she uses the 100 years of sleep to think about what to say to her Prince after waking up.) Hence, I assume that Beauty, while dreaming, is not incapacitated in a way that renders the ordinary norms of Bayesian rationality inapplicable (notwithstanding the possible incapacitation due to the memory loss that results from the amnesia-inducing drug). Section 4 will discuss objections against these assumptions.

Beauty's credence function on Sunday evening is given by P_-. P_* is her credence function during her dream. P is her credence function after she is woken up. P_+ is her credence function after being told it is Monday. She considers the following propositions[3] on Monday:

H_1: the coin lands heads and it's Monday,
H_2: the coin lands heads and it's Tuesday,
T_1: the coin lands tails and it's Monday,
T_2: the coin lands tails and it's Tuesday,
M: it is Monday,
H: the coin lands heads.

Her evidence includes $M \leftrightarrow H_1 \vee T_1$ and $H \leftrightarrow H_1 \vee H_2$, so she treats M and H as equivalent to these disjunctions.

2.1. Beauty's dreaming credence in heads

I argue that $P_*(H) = 1/2$.

[2] We might also imagine she has a lucid dream on random nights, and happens to have one on Monday. This does not change the analysis.

[3] I use the word "proposition" to refer to the objects of beliefs, without taking a position on what kind of objects are belief objects (e.g., sentences, traditional propositions or sets of centred worlds).

The argument uses the Principle of Irrelevant Evidence used by Lewis: only new relevant evidence can produce a change in credences. Lewis (2001) claims that Beauty has not received new evidence after waking up on Monday compared with her evidence on Sunday that is relevant to the coin toss. In particular, Lewis denies that the evidence of being woken up today is relevant. This evidence, in Lewis's analysis of the problem, is equivalent to $H_1 \vee T_1 \vee T_2$. That is, it rules out H_2.

According to Lewis, the relevant portion of Beauty's evidence after waking is the same as her evidence on Sunday, and the evidence of being woken up today is irrelevant. The Lewisian must therefore claim that Beauty's relevant evidence while dreaming is also the same as on Sunday. After all, Beauty's dreaming evidence is the same as her evidence after waking, with the exception that it does not include the evidence of being woken up today – Which Lewis takes to be irrelevant. Uncontroversially, on Sunday Beauty has $P_{-}(H) = 1/2$. Hence, it follows from Lewis's assumptions – the Principle of Irrelevant Evidence and the irrelevance of Beauty's new evidence on Monday – that Beauty's credence in heads while dreaming is $P_*(H) = P_*(\neg H) = 1/2$.

This suffices to show that *Lewis* should agree with this credence, but perhaps other Lewisian halfers will attempt to reject this line of reasoning. The Principle of Irrelevant Evidence is a special case of Bayesian conditionalization in which one's relevant evidence is unchanged. Bayesian conditionalization is known to fail in some cases where self-locating evidence is learned, forgotten, or has changed in truth value. Hence, a Lewisian halfer may attempt to object to its use in my argument above.

Such an objection would likely focus on the fact that during her dream, Beauty does not know which day it is, but knows that it is either Monday or Tuesday. This is a change in her evidential situation since Sunday. Moreover, the proposition that today is either Monday or Tuesday has changed in truth value since Sunday. This is what Bradley (2011) calls a *belief mutation*, and belief mutations are potentially problematic for Bayesian conditionalization. (Bradley himself argues that such a belief mutation is irrelevant for eternal propositions like H, which would also confirm $P_*(H) = 1/2$.)

However, only relevant changes in evidence should be considered problematic. As anyone in the debate in favour of Bayesian conditionalization agrees, cases of irrelevant belief mutations are not problematic. For example, suppose I know that it is now 12:00 at 12:00, and I assign a credence q to some eternal proposition A at 12:00. If the time is irrelevant to A, I should be able to apply Bayesian conditionalization at 12:05, even though there is a belief mutation in my belief of the current time. In particular, I should still assign a credence of q to A at 12:05 if the only change in my evidence is the belief mutation concerning the current time.

The Lewisian halfer's objection should therefore show that the evidence that it is either Monday or Tuesday, or her former evidence that it is Sunday, is relevant for Beauty's belief in the outcome of the coin toss during her dream.

But clearly it is not. Ordinarily, the current day is irrelevant to the outcome of coin tosses. Unlike after Beauty is woken up, the outcome of the coin toss bears no relation to her state while dreaming: she has a lucid dream regardless. Hence, this objection is unlikely to succeed.

2.2. Beauty's dreaming credence that it is Monday and the coin landed heads

When Beauty has a dream on Monday morning, she is unsure which day it is. Her last memories are from Sunday evening. But her memories of Monday are erased on Monday evening. So when she has a dream on Tuesday, her last memories will also be from Sunday evening. Hence, from Beauty's perspective, it is possible that it is Tuesday.

Moreover, it is also clear that Beauty considers it possible that it is Tuesday and that the coin landed heads. (Only after she is woken up does this become impossible.) Hence, we have $P_*(H_2) > 0$. Moreover, by the above, we have $P_*(H) = P_*(H_1) + P_*(H_2) = 1/2$. It follows that $P_*(H_1) < 1/2$.

Note that some halfers might want to reject the assumption that Beauty is unsure which day it is, instead requiring that Beauty is sure that it is Monday while dreaming. However, this position implies that both waking up ($\neg H_2$) and subsequently being told that it is Monday ($H_1 \vee T_1$) do not give Beauty new evidence. Hence, by the Principle of Irrelevant Evidence, Beauty's credence in heads would remain 1/2 throughout the day. This position is called steadfast halfing (Schwarz, 2025), and it conflicts with the defining characteristic of Lewisian halfing that $P_+(H) = 2/3$.

2.3. Beauty's credences after waking

When Beauty is woken up, she learns that the researchers wake her up today. In our model, she learns $H_1 \vee T_1 \vee T_2 \equiv H_1 \vee \neg H$, ruling out H_2. This information is self-locating, because no eternal propositions like H, $\neg H$, M, or $\neg M$ are ruled out. The only thing that is ruled out is that she is at a particular location (Tuesday) in the world at which the coin lands heads.

This is not a belief mutation: the truth value of the proposition that the researchers awaken her today is the same before and after being awakened. Instead, this is what Bradley (2011) calls a belief discovery. Lewisian halfers agree that belief discovery is unproblematic for Bayesian conditionalization. In fact, this is a defining characteristic of Lewisian halfing, which holds that Beauty can conditionalize on the belief discovery that it is Monday after the researchers tell her it is Monday. And so they should similarly hold that Beauty can update on $H_1 \vee \neg H$.

Hence, Beauty can apply Bayesian conditionalization as usual. Combining this with Beauty's credences during her dream as determined in the previous section, we get

$$P(H) = P_*(H \mid H_1 \vee T_1 \vee T_2) \tag{1}$$
$$= \frac{P_*(H_1)}{P_*(H_1 \vee \neg H)} \tag{2}$$
$$= \frac{P_*(H_1)}{P_*(H_1) + 1/2} \tag{3}$$

Here (1) uses Bayesian conditionalization, (2) uses the definition of conditional probability, and (3) uses that H_1 and $\neg H$ are mutually exclusive and $P_*(\neg H) = 1/2$. Finally, note that from $0 \leq x < 1/2$ it follows that $x/(x + 1/2) < 1/2$. We have $P_*(H_1) < 1/2$, so we have

$$P(H) < 1/2. \tag{4}$$

This refutes Lewisian halfing.

One possible objection, alluded to by Bradley (2003), is to claim that Bayesian conditionalization is rationally required only in case the possible pieces of evidence one may receive form a partition, that is, a set of mutually exclusive propositions of probability 1. While dreaming, it is possible for Beauty to learn $\neg H_2$, by waking up, but it is not possible to learn H_2. Since $P_*(\neg H_2) < 1$, the possible pieces of evidence Beauty may receive do not form a partition.

But such a blanket ban on Bayesian conditionalization in cases of non-partitionality appears unwarranted. Although it has been argued that cases of non-partitionality require a slight alteration to conditionalization (Schoenfield, 2017), such an alteration would not produce different results in our situation.[4] Moreover, the alternative update that would be required to save Lewisian halfing has the counterintuitive consequence that all the dreaming credences assigned to H_2 are assigned to H_1 after waking up.[5] Incidentally, an updating rule that has this consequence is the *halfer rule* (Briggs, 2010; Conitzer, 2015), which is sometimes accepted by double halfers but (rightly) rejected by Lewisian halfers.

3. An argument for thirding

We can derive thirding by adding two additional plausible assumptions.

[4] Schoenfield (2017) argues that in cases of non-partitionality one should condition on "I learn that E" instead of just E. This does not make a difference here, since Beauty knows that she learns she wakes up if and only if she wakes up.

[5] To see why such an updating rule is counterintuitive, consider a similar scenario: you are unsure which day of the month June it is, assigning all days equal credence. I toss a coin without showing the result. If it is not June 1 and it lands heads, you will be instantly killed after 5 seconds. You wait 5 seconds and survive, ruling out $H_2 \vee \ldots \vee H_{30}$. It seems you should now become quite sure of tails, as Bayesian conditionalization requires. The updating rule required to save Lewisian halfing instead leads you to assign an absurdly high credence to H_1 (heads and it's June 1) of $1/2$.

3.1. Beauty's dreaming credence that it is Monday

During Beauty's dream, she is uncertain which day it is. Symmetry considerations suggest assigning an equal credence to it being Monday and Tuesday. Hence, we have $P_*(M) = 1/2$.

We can also defend a credence of 1/2 that it is Monday using Elga's restricted principle of indifference (Elga, 2004). According to this principle, one should assign each subjectively indistinguishable location within a possible world at which one can be equal credences. Two locations are subjectively indistinguishable if one's experience in both locations is identical.

During Beauty's lucid dream, H_1 and H_2 are locations in the same possible world. Moreover, they are subjectively indistinguishable. Hence, Elga's restricted principle of indifference requires $P_*(H_1) = P_*(H_2)$. By the same argument, it requires $P_*(T_1) = P_*(T_2)$. Hence, we have $P_*(M) = P_*(H_1) + P_*(T_1) = P_*(H_2) + P_*(T_2) = P_*(\neg M)$, so $P_*(M) = 1/2$.

3.2. The independence of Monday and heads while dreaming

Which day it is and how a coin toss lands are normally probabilistically independent. During Beauty's dream, her evidence does not connect the outcome of the coin toss and the current day. She does know that she will be awakened in the future depending on the current day and the outcome of the coin toss. Clearly, however, this information does not make the toss and the day dependent before she is actually woken up.

Beauty's credences in M and H are therefore independent. So using the dreaming credences argued from above, we have $P_*(H_1) = P_*(H)P_*(M) = 1/4$.

Plugging this into equation (3) yields $P(H) = 1/3$.

4. The disanalogy objection

There is one final way in which the Lewisian halfer could attempt to resist. She might accept the derivation, but object that the derivation only applies to a version of the problem that is disanalogous to the original. While the original problem does not rule out that Beauty has dreams, there is at least one different way in which such a disanalogy objection might be carried out.

It might be objected that even ideally rational agents do not form rational beliefs during their dreams, or that they do not update based on their dreaming credences after waking up. Normal humans are incapacitated during their dreams in several ways. For example, they typically do not have access to all their memories while dreaming, and some of their memories of the dream are lost after waking up. Therefore, it might be argued, the norms of rationality only apply to agents who are in a state of being awake. We should thus imagine ideal agents to be awake at all times, or to not have credal states while dreaming, or to never update based on their dreaming credences after waking up.

It should be noted that this objection does not come cheap. Bayesian norms of rationality are often thought to apply to agents who are never in a state of being awake, such as artificial intelligence and teams of scientists. In the stated form, the objection thus seems to severely restrict the scope of Bayesian epistemology.

Whichever form the disanalogy objection takes, I will argue that the dreaming scenario is still analogous to the original problem. In the subsequent sections, I will assume that the disanalogy objection takes the form of: "the original Beauty never has dreams" (and a scenario in which she does have dreams with rational credal states is "disanalogous"). But my arguments equally apply to other types of disanalogy objections.

4.1. Objections to previous arguments by analogy

A disanalogy objection might have worked for previous arguments for thirding by analogy. For example, Dorr (2002) introduces a variant in which Beauty is definitely woken up on both Monday and Tuesday but given one of two possible amnesia-inducing drugs. The first drug, administered when the coin lands tails, has the same effect as the drug in the original version. The second drug, administered in case of heads, is weaker. The weaker drug has the same effect during the first minute after waking up on Tuesday, but memories of her Monday awakening will return after one minute. Similarly to Beauty's dreaming credences in my variant, in Dorr's variant Beauty should plausibly assign a credence of $1/4$ to all four possibilities immediately after waking up. After her memories fail to come back she can rule out H_2; so by Bayesian conditionalization, she ought to believe H to degree $1/3$.

Arguments have been offered that Dorr's case is disanalogous, which also apply to Arntzenius (2003). First, as Bradley (2003) argues, in the variant case Beauty can receive both the evidence H_2 (if her memories come back) and $\neg H_2$. In the original variant, Beauty never learns H_2, since she is not woken up on Tuesday when the coin lands heads. Bradley argues that she therefore cannot conditionalize on $\neg H_2$. The dreaming scenario is not disanalogous in this sense, since it remains impossible for Beauty to learn H_2 on Tuesday.[6]

Schwarz (2025, p. 1086) also argues that the variants by Dorr and Arntzenius are disanalogous to the original because maximizing the so-called *average expected accuracy* with respect to Beauty's Sunday credences recommends Lewisian halfing, but thirding in the variants. However, as I show in section 5.2, Schwarz's arguments apply only to a version of the Sleeping Beauty problem different from the original offered by Elga and Lewis. An improved version of Schwarz's accuracy principle can nevertheless be used to create a disanalogy objection that may work. I discuss this objection in section 5.4.

[6] As discussed in section 2.3, another version of this objection is that Beauty cannot conditionalize based on her dreaming credences.

4.2. Disanalogy due to different evidence

An argument that the dreaming scenario is disanalogous is unlikely to work if one accepts the following evidentialist principle. If an agent's relevant evidence is the same in two scenarios (real or hypothetical), then the set of rationally permissible credences is the same in both.[7] Note that the objector concedes that in the scenario in which Beauty dreams, a credence of $1/2$ after waking up is rationally impermissible. At the same time, the objector claims that a credence of $1/2$ is permissible in the scenario in which she does not dream. Hence, by the evidentialist principle, the objector must claim that Beauty's relevant evidence about the coin toss is different in both scenarios, after waking up.

In the dreaming scenario, Beauty clearly has at least all the evidence that she has in the non-dreaming scenario. So the objector must claim that her memories of the dream contain additional relevant evidence about the way in which the coin landed. It is in principle possible that dreams contain evidence: for example, someone might whisper something in your ear while sleeping, you hear it in your dream, and have good reason to believe this comes from the outside world. But this is clearly not the sort of situation that Beauty finds herself in while dreaming.

A final way in which her memories might contain relevant evidence is when the fact of having dreamt itself is associated in some way with external events. For example, suppose that the experimenters cause Beauty to dream when the coin lands heads, but not when it lands tails. In such a situation, having dreamt is evidence that the coin landed heads. Again, the situation in which Beauty finds herself in the dreaming scenario is clearly not one in which the fact of having dreamt is evidence for the way in which the coin landed.

Hence, there is no plausible sense in which Beauty's memories of having dreamt are relevant evidence for the way in which the coin lands. There is thus no plausible argument for disanalogy on evidentialist grounds.

4.3. Non-evidentialist disanalogy objections

Finally, a Lewisian halfer might try to object on non-evidentialist grounds, by rejecting the above evidentialist principle. It then becomes possible that the way and order in which an agent comes to learn and forget evidence matters for norms of rationality.

But even the non-evidentialist must point to some relevant difference that might plausibly require thirding if Beauty has dreams but halfing if she doesn't. Having rejected evidentialism, the objector relies on diachronic norms of rationality to do the work. That is, she must offer some reason that rationally updating in two stages (from Sunday to Monday dreaming and Monday dreaming to Monday awake) should lead to thirding while rationally updating once (from Sunday dreaming to Monday awake) should lead to halfing. Even if such

[7] Versions of this principle are defended by Hedden (2015) and Moss (2015).

a reason exists, a second problem immediately arises: in the dreaming scenario, Beauty could also ignore her dreams and update directly based on her Sunday credences. To make her diachronic argument for disanalogy work, it seems that the objector must accept that updating directly based on Sunday credences is rationally permissible and leads to halfing. Beauty would therefore have two permissible updating methods available that conflict: one based on her Sunday credences and one based on her dreaming credences.

The objector must hold – in order to maintain the disanalogy objection – that when dreaming credences are available, the two-stage updating method yields the most rational credences. So Beauty would need to use her dreaming credences after having dreamt (leading to thirding) but her Sunday credences when she did not dream (leading to halfing). But if the objector concedes that the two-stage updating credences are *better* from the perspective of rationality, it raises the question of why the one-stage credences are preferred in the non-dreaming scenario. In the latter scenario, Beauty is still capable of deriving the credences that she would have had if she had dreamt. So if they are indeed better, it seems that she should adopt them regardless.

A possible response, similar to Schwarz (2025), is that ideal rationality breaks down in situations where an agent forgets or potentially forgets. Hence, ideal rational credences are unavailable in both scenarios. It might thus be argued that Beauty's credences in both scenarios are in some sense "bad", such that there is no principled by-analogy reasoning available to Beauty after she wakes up.

The trouble with this type of response is that Beauty's dreaming credences as defended in section 2 are highly intuitive and plausibly ideally rational. Moreover, between dreaming and waking up, Beauty does not forget any evidence. Hence, if Beauty's dreaming credences are ideally rational, then her waking credences are also ideally rational. This does not refute the claim that no ideally rational credences exist if Beauty does not dream. However, the above question can now be repeated with more force. Even when Beauty does not dream, she can derive the ideally rational credences that she would have had if she did. Why should she not adopt them?

The above considerations may not show that any version of this disanalogy objection must fail. Instead, they aim to show that such an objection is difficult to formulate and rests on particular assumptions, such as a rejection of evidentialism or the position that the Sleeping Beauty problem is outside the scope of ideal rationality. This shifts the burden of proof back to Lewisian halfers.

5. The expected accuracy argument for Lewisian halfing

Schwarz (2025) offers a new argument for Lewisian halfing that is based on a (non-standard) diachronic accuracy maximization principle. This principle, called *average expected accuracy updating*, requires Beauty to adopt an updating rule, on Sunday, that maximizes the *average expected accuracy* (AEA) over

Variant	Memory loss on	No dreams	Dreams
Lewis (2001)	Monday night	$P(H) = 1/3$	$P_*(H) = 1/2$ $P(H) = 1/3$
Elga (2000)	Monday night Tuesday night (iff tails)	$P(H) = 3/7$	$P_*(H) = 3/5$ $P(H) = 3/7$
Schwarz (2025)	Monday night (iff tails)	$P(H) = 1/2$	$P_*(H) = P(H) = 1/2$

Table 2: Results of AEA updating for different variants of the Sleeping Beauty problem.

the occasions on which it is invoked. While Schwarz's argument has problems, I show that a revised version of this principle can be used both to defend Lewisian halfing and make a disanalogy objection. I argue that this is the strongest version of Lewisian halfing, although it appears to collapse to non-halfing if we take its commitments seriously.

Schwarz maintains that the Sleeping Beauty problem is a problem of non-ideal rationality because the threat of amnesia makes it impossible to abide by ideal norms. AEA updating (as applied to memory loss cases) is proposed as a principle of non-ideal rationality to be used when ideal rationality is inapplicable.[8]

However, Schwarz's argument rests on a formulation of the Sleeping Beauty problem that is different from the original version of the Sleeping Beauty problem as described by Elga (2000) and Lewis (2001). In Schwarz's version of the problem, Beauty's memory is erased on Monday night only when the coin lands tails. In Lewis's variant, on the other hand, her memory is erased on Monday regardless of how the coin lands: "On Monday they will awaken [Beauty] briefly. [...] Then they will subject her to memory erasure" (Lewis, 2001, p. 171). In Elga's variant, Beauty's memory is definitely erased on Monday, and additionally on Tuesday if she is awakened on that day: "after each waking, they will put you back to sleep with a drug that makes you forget that waking" (Elga, 2000, p. 143).[9]

It turns out that these details matter for the proposed accuracy principle. For each of the three variants – Lewis, Elga, and Schwarz – AEA updating recommends different credences. Table 2 gives an overview of these variants and the associated accuracy-maximizing credences.

The conditional memory loss on Monday in Schwarz's variant also invalidates my dreaming argument. If Beauty has no memory loss on heads, she can rule out H_2 during her dream, since she would have had memories of Monday in that scenario. Hence, the assumption $P_*(H) > 0$ from section 2.2 is no longer warranted. My argument therefore does not apply to Schwarz's variant.

[8] Schwarz also argues that a similar average expected accuracy updating principle that applies to personal fission cases (such as duplication) is an *ideal* principle of rationality.

[9] Note that in the dreaming version of the Elga variant, I assume that her memories while dreaming are also erased. As Elga (2000, p. 143) write: "the precise effect of the drug is to reset your belief-state to what it was just before you were put to sleep at the beginning of the experiment."

There are other influential early authors whose formulation is similar to Schwarz's, with conditional memory loss on Monday, including Dorr (2002), Draper and Pust (2008), Horgan (2004), Meacham (2008), and Titelbaum (2008). Versions with definite memory loss on Monday are used by Arntzenius (2003), Bostrom (2007), Briggs (2010), Horgan (2004), Pust (2008), Titelbaum (2013b), and Weintraub (2004), and, notably, the Lewisian halfer Bradley (2011). None of the authors who adopted a version with conditional memory loss has noted the discrepancy with the original version, so it seems safe to conclude that no one thought the distinction should make a difference to Beauty's credences.

In what follows I show how AEA updating gives different results depending on the variant of the Sleeping Beauty problem. I then argue that this principle gives counterintuitive results that can be avoided by switching to an alternative principle called *revised average expected accuracy updating* (RAEA). This principle recommends Lewisian halfing in all three variants. It also turns out that RAEA updating can be used to make a disanalogy objection against my main argument against Lewisian halfing.

5.1. Average expected accuracy updating

In Schwarz's framework, propositions are sets of centred possible worlds, which specify both what is factually true in the world and where one is located in that world. A centred world w is a triple (u, i, t) of an uncentred possible world u, an individual i, and a time t. For our purposes, H_1 and H_2 can be interpreted as centred worlds (heads, Beauty, Monday) and (heads, Beauty, Tuesday). The latter centred world, for example, describes the state in which the coin lands heads, the agent is Beauty, and it is Tuesday. In what follows, I leave out "Beauty" in centred world descriptors, so we get (heads, Monday) and (heads, Tuesday). Since states will be interpreted as centred worlds, we have H_1 = (heads, Monday) and H_2 = (heads, Tuesday).

Expected accuracy updating aims to maximize the expected accuracy of the agent's credences after the update from the perspective of the agent's credences before the update. The rule is adopted at the earlier time and applied at the moment of updating. But after being put to sleep on Sunday, Beauty undergoes what Schwarz calls "doxastic fission": there are two possible later times at which the next update will take place – Monday and Tuesday. Doxastic fission raises the question of how expected accuracy is to be calculated. For his variant of the problem, Schwarz shows that maximizing *total expected accuracy* on Sunday recommends thirding, while his preferred rule of maximizing the AEA recommends halfing.

The following definitions are needed to formulate the latter rule. An accuracy measure $V(P, w)$ is a function assigning an accuracy to a credence function P at a world w, and it must be "strictly proper".[10] For each centred world w located

[10] An accuracy function V is strictly proper if any probability measure assigns itself maximum expected accuracy. That is, for all probability measures P, $\sum_w P(w)V(P', w)$ has a unique

at the initial moment (before the update), the *doxastic successors* of w are defined as a set $s(w)$ of centred worlds at which the next update after w will be carried out. These are all the later moments at which the agent's credences are a direct causal consequence of the chosen update disposition.[11] For example, when Beauty chooses an update rule on Sunday in the Lewis variant, her initial possible world (heads, Sunday) has two successors: (heads, Monday), and (heads, Wednesday), since these are the immediate next moments on which the update rule chosen on Sunday is applied given heads.[12] On the other hand, the world (tails, Sunday) has the successors (tails, Monday) and (tails, Tuesday).

Lastly, an update disposition μ assigns a probability function $\mu(w)$ to a centred world w. The agent considers only those functions μ that are possible, which requires that chosen credence functions can only be different for worlds that her evidence can distinguish when she is in them. For example, the updating rule that Beauty chooses on Sunday must satisfy $\mu((\text{heads, Monday})) = \mu((\text{tails, Monday})) = \mu((\text{tails, Tuesday}))$.

Suppose the agent's earlier probability function P_1 is defined on a sample space Ω of centred worlds. Then the *average expected accuracy* of the updating rule μ is

$$EV(\mu) = \sum_{w \in \Omega} P_1(w) \sum_{w' \in s(w)} \frac{V(\mu(w'), w')}{|s(w)|}.$$

In case of doxastic fission of a world w, the AEA sums over the accuracy of updated probabilities in possible fission states divided by the number of successors of w, $|s(w)|$.

In Schwarz's variant, maximizing expected accuracy after waking up on Monday or Tuesday with respect to P_- (Sunday) leads to $P(H) = 1/2$. The same maximization between P and after Beauty is told which day it is leads to $P_+(H) = 2/3$ (Schwarz, 2025, pp. 1081–1084). Hence, we end up with the probabilities of Lewisian halfing.

5.2. The Lewis and Elga variants

First consider Lewis's variant. As mentioned above, the world (heads, Sunday) has the successors H_1 = (heads, Monday) and H_3 = (heads, Wednesday). The world (tails, Sunday) has the successors T_1 = (tails, Monday) and T_2 = (tails, Tuesday). To all worlds centred on Monday and Tuesday, μ must assign

maximum at $P' = P$.

[11] In Schwarz's model, it is not possible for the agent to disregard her update dispositions and choose her credences differently. As soon as the agent chooses an update disposition, each doxastic state directly following the current state is fixed by this update disposition.

[12] It might be possible to have a notion of "successors" according to which (heads, Wednesday) is not a successor of (heads, Sunday). Such a version of AEA would assign the same credences to all three variants and works similar to my proposal in section 5.4, RAEA. However, it is unclear what this notion of "successors" could be, and it raises the question of how Beauty's credences in (heads, Wednesday) are determined, if not based on her Sunday credences.

the same credence function, as Beauty cannot distinguish any of her awakenings. On Wednesday, Beauty knows that it is Wednesday, so her updating rule may assign a different credence function to Wednesday. Let P_w stand for Beauty's credence function at (heads, Wednesday). As before, P is Beauty's credence function after waking up on Monday and Tuesday. The AEA of μ on Sunday is

$$EV(\mu) = \frac{1}{2}\left(\frac{V(P, H_1) + V(P_w, H_3)}{2}\right) + \frac{1}{2}\left(\frac{V(P, T_1) + V(P, T_2)}{2}\right).$$

As expected accuracy is additive, P_w can be maximized separately, so that term can be ignored for the purposes of maximizing P.[13] Maximizing this "P-portion" of the EAE then yields $P(H_1) = P(T_1) = P(T_2) = 1/3$. So AEA updating recommends thirding in Lewis's variant.

If Beauty has lucid dreams, the successors of (heads, Sunday) are H_1 and H_2 and the successors of (tails, Sunday) are T_1 and T_2. There are no Wednesday successors, since on Wednesday Beauty always has memories of her Tuesday dreaming credences. The AEA of μ becomes

$$EV(\mu) = \frac{1}{2}\left(\frac{V(P_*, H_1)}{2} + \frac{V(P_*, H_2)}{2}\right) + \frac{1}{2}\left(\frac{V(P_*, T_1)}{2} + \frac{V(P_*, T_2)}{2}\right).$$

This is maximized when $P_*(H_1) = P_*(H_2) = P_*(T_1) = P_*(T_2) = 1/4$.[14] Hence, we get the dreaming credences required by the argument for thirding in section 3.

When there is no doxastic fission, (average) expected accuracy updating corresponds to Bayesian conditionalization.[15] So it does not matter whether, between dreaming and waking up, Beauty uses plain Bayesian conditionalization or AEA updating. As in section 3, we get $P(H) = 1/3$.

Turning to Elga's variant, the successors of (heads, Sunday) are H_1 and H_3, while the successors of (tails, Sunday) are T_1, T_2, and T_3 = (tails, Wednesday). The P-portion of the AEA is thus

$$\frac{1}{2}\left(\frac{V(P, H_1)}{2}\right) + \frac{1}{2}\left(\frac{V(P, T_1) + V(P, T_2)}{3}\right).$$

Maximizing this gives $P(H_1) = 3/7$ and $P(T_1) = P(T_2) = 2/7$. As in Lewis's variant, this conclusion is unchanged when Beauty has dreams. We then get $P_*(H_1) = P_*(H_2) = 3/10$ and the same post-waking credences.[16]

---

[13] Note that maximization gives $P_w(H_3) = 1$ as expected: if Beauty wakes up on Wednesday without any memories from the two previous days, she can infer that the coin landed heads.

[14] Let Pr be such that $Pr(H_1) = Pr(H_2) = Pr(T_1) = Pr(T_2) = 1/4$. Then we have $EV(P_*) = \sum_w Pr(w)V(P_*, w)$. Since V is strictly proper, $\sum_w Pr(w)V(P_*, w)$ is maximized by $P_* = Pr$.

[15] See Schwarz (2025, p. 1076, footnote 7). Schwarz's "shifted conditionalization" is identical to classical conditionalization whenever $s(w) = w$ for every world w before the update, which is the case for the worlds in Beauty's sample space while dreaming.

[16] In the Elga variant with dreams, the successors of (heads, Sunday) are H_1 and H_2, and the

5.3. Counterintuitive consequences of average expected accuracy updating

It is an awkward consequence of AEA updating that it recommends different credences for the three variants of the Sleeping Beauty problem. The differences between these variants, from Beauty's perspective, are only the memories that she has on Wednesday. Intuitively, this should not affect her credences on Monday and Tuesday, since Beauty is not uncertain whether it is Wednesday on Monday and Tuesday in any of the variants. Moreover, the accuracy of Beauty's credences on Monday and Tuesday, on the one hand, and Wednesday, on the other, are unrelated: neither of the credences affects the accuracy of the other. They can be maximized separately.

The following case, inspired by Bostrom (2007), sharpens this counterintuitive result. Suppose that Beauty is awakened on a million subsequent days in case of tails. Suppose that her memories are erased after each awakening if the coin lands tails. However, she is immediately made aware of which day it is upon each awakening except the Monday and Tuesday awakening. Hence, Beauty's experience and evidence on Monday and Tuesday is just like that in the original Sleeping Beauty problem, whereas all subsequent days on tails are similar to Wednesday in the original. Intuitively, Beauty's optimally rational credence on Monday should thus not differ, or at least not differ much, from her credence in the original Sleeping Beauty problem.

Let us apply AEA updating. The only successor of (heads, Sunday) is H_1, while there are a million successors of (tails, Sunday), namely T_1 until $T_{1000000}$. Beauty's updating rule can pick different credence functions for each awakening day after Tuesday, called P_3 until $P_{1000000}$, while P is her credence function on both Monday and Tuesday. Clearly, Beauty chooses $P_i(T) = 1$ for $i \geq 3$, which maximizes expected accuracy on these days. To calculate the maximum for P, consider that the AEA of her updating rule is

$$EV(\mu) = \frac{1}{2}V(P,H_1) + \frac{1}{2} \cdot \frac{V(P,T_1) + V(P,T_2) + \sum_{i=3}^{1000000} V(P_i,T_i)}{1000000}.$$

The P-portion of $EV(\mu)$ is

$$\frac{1}{2}V(P,H_1) + \frac{V(P,T_1) + V(P,T_2)}{2000000}.$$

This is maximized for $P(H) = P(H_1) \approx 0.999998$. AEA updating gives the absurd result that Beauty should be almost sure that the coin has landed heads and that it is Monday.

---

successors of (tails, Sunday) are T_1, T_2, and T_3. On Wednesday Beauty knows that it is Wednesday, so the P-portion of the AEA is

$$\frac{1}{4}V(P_*,H_1) + \frac{1}{4}V(P_*,H_2) + \frac{1}{6}V(P_*,T_1) + \frac{1}{6}V(P_*,T_2).$$

Maximizing this gives $P_*(H_1) = P_*(H_2) = 3/10$ and $P_*(T_1) = P_*(T_2) = 1/5$. Bayesian conditionalization on $\neg H_2$ yields $P(H) = P(H_1) = 3/7$.

5.4. Revised expected accuracy updating and a disanalogy objection

The problem with Schwarz's principle appears to be that the agent averages over all successors, combining centred worlds in which she has different evidence and in which she can therefore have different credences. We can avoid the counterintuitive consequences by averaging only over the subjectively indistinguishable world-parts, that is, only those centred worlds in which the agent is forced to have the same credences.

This revision might be motivated as follows.[17] When Beauty wakes up, she should only care about the accuracy of her current credences. Her credences on Wednesday are irrelevant, since she knows it is not Wednesday. But since she does not know whether it is Monday or Tuesday given tails, optimizing the accuracy given tails requires her to give equal consideration to both.[18]

Let $s^*(w)$ be the set of sets of doxastic successors of w where each $W \in s^*(w)$ contains all centred worlds that are subjectively indistinguishable. The *revised average expected accuracy* (RAEA) of the updating rule μ is

$$EV^*(\mu) = \sum_{w \in \Omega} P_1(w) \sum_{W \in s^*(w)} \frac{\sum_{w' \in W} V(\mu(w'), w')}{|W|}.$$

Applying this to the Lewis variant gives

$$EV^*(\mu) = \frac{1}{2}\left(V(P, H_1) + V(P_w, H_3)\right) + \frac{1}{2}\left(\frac{V(P, T_1) + V(P, T_2)}{2}\right),$$

and this is maximized for $P(H) = 1/2$.

Like average and total expected accuracy updating, RAEA updating is the same as normal expected accuracy updating when doxastic successors are unique. Hence, it recommends Bayesian conditionalization when it is available.[19] Applying Bayesian conditionalization after learning that it is Monday then gives $P_+(H) = 2/3$. We get Lewisian halfing. The same result is obtained when applying RAEA updating to the Elga and Schwarz variants, as well as the variant with a million awakenings.

When Beauty has dreams, the RAEA is equal to the AEA from section 5.2, giving $P_*(H) = 1/2$. Applying Bayesian conditionalization then gives $P(H) = 1/2$. Hence, surprisingly, RAEA recommends different credences when Beauty does and does not dream.

[17] This motivation is similar to the motivation for a type of average expected accuracy maximization offered by Kierland and Monton (2005, p. 390).

[18] Note that this motivation sits uneasily with the premise that Beauty maximizes the accuracy of her future credences *on Sunday*. But given this premise, it is hard to see why Beauty should discount some of her future credences' accuracy at all, as done using the division in both AEA and RAEA. From her Sunday perspective, she is just as likely to end up in H_1 as *both* T_1 and T_2, and she spends an equal amount of time in these situations. Hence, it seems she should use *total* expected accuracy updating, rather than either AEA updating or RAEA updating.

[19] See footnote 15.

This makes available a disanalogy objection against my argument for the Lewisian halfer who adopts RAEA updating. If RAEA updating is a correct principle of rationality that is appropriately applied to both the dreaming and non-dreaming scenario, then Beauty ought to have different credences depending on whether she has dreams.

I argued in section 4.3 that someone who makes this argument has some explaining to do, especially if they, like Schwarz, maintain that RAEA updating is a principle that deals with situations of non-ideal rationality. Beauty's dreaming credences appear *ideally* rational. Since Beauty can derive them even if she doesn't dream, why shouldn't she also adopt them when she does?

To this it might be responded that proponents of expected accuracy updating don't care about the supposed similarity between such cases – they only care about optimizing accuracy. When new situations like the dreaming stage are added, the calculation becomes different. Apparently it is better in terms of expected accuracy for Beauty to be a thirder if she also dreams, but not if she doesn't – case closed.

But if we take these commitments of expected accuracy updating seriously, then the above calculations are much less certain. I have assumed throughout that Beauty considers only the accuracy of her immediate doxastic successors. But if expected accuracy is what matters above everything else, then Beauty should not ignore the expected accuracy of later successors. And if she doesn't, we don't get Lewisian halfing.

Consider again the no-dreaming scenario in which Beauty is told on Monday that it is Monday a little while after being woken up. Beauty is bound on Monday by Bayesian conditionalization after she is told that it is Monday, since this is the updating rule required by RAEA updating. Hence, she cannot decide on Sunday to overrule her Monday inclination to update by Bayesian conditionalization. She knows this on Sunday, so she should plausibly take into account the accuracy of her beliefs after being told it is Monday, when deciding on Sunday how to update immediately after waking up. Letting s_2^* be defined like s^* but for successors of successors, she should take into account the RAEA of the successors of her successors:

$$EV_2^*(\mu) = \sum_{w \in \Omega} P_1(w) \sum_{W \in s_2^*(w)} \frac{\sum_{w' \in W} V(\mu(w'), w')}{|W|}.$$

Let us suppose the expected accuracy of her successors of successors has the same weight as the expected accuracy her immediate successors, so Beauty maximizes $EV^*(\mu) + EV_2^*(\mu)$. The P-portion of this is

$$\frac{1}{2}\left(V(P, H_1) + V(P_w, H_3)\right) + \frac{1}{2}\left(\frac{V(P, T_1) + V(P, T_2)}{2}\right) + \frac{1}{2}V(P(\cdot \mid M), H_1) + \frac{1}{2}V(P(\cdot \mid M), T_1).$$

This quantity is different for different accuracy functions, so I will use the popular Brier score of inaccuracy. (Since this is an inaccuracy measure we need to minimize instead of maximize.) For a proposition X and a possible world w, let $X(w) = 1$ if X is true at w and $X(w) = 0$ otherwise. Then the Brier inaccuracy is $V(P,X) = \sum_{w\in\Omega}(P(w) - X(w))^2$. Writing $a = P(H_1)$, $b = P(T_1)$, $c = P(T_2)$, and $q = P(H \mid M) = a/(a+b)$, we get a P-portion of

$$\tfrac{1}{2}\left[(a-1)^2 + b^2 + c^2\right] + \tfrac{1}{4}\left[a^2 + (b-1)^2 + c^2\right] \\ + \tfrac{1}{4}\left[a^2 + b^2 + (c-1)^2\right] + (1-q)^2 + q^2.$$

This is minimized for $a \approx 0.424$, $b \approx 0.333$, $c \approx 0.244$, giving $P(H) \approx 0.42$.[20] So RAEA updating does not recommend halfing once Beauty takes the accuracy of her credences after being told it is Monday into account.[21]

It might be objected that this calculation inappropriately tilts the accuracy weights towards Monday. Although descriptions of the Sleeping Beauty problem do not clarify this, someone might claim that Beauty spends an equal amount of time awake on Monday and Tuesday. In that case, we should incorporate a symmetric accuracy for Beauty's Tuesday credences during the period after she would have been told it is Monday if it were Monday. That is, we should treat "tails and Tuesday after being awake for a while" as a doxastic successor of "tails and Tuesday immediately after being awake", for the purposes of EV_2^*. But since Beauty knows exactly how the experiment goes, she is able to deduce after this period of time that it is Tuesday as she has not been told it is Monday. Hence, all of her credences are trivial in that scenario: $P(\cdot \mid \neg M)$ puts probability 1 on T_2 and 0 on everything else. The Brier inaccuracy is $V(P(\cdot \mid \neg M), T_2) = 0$, so the additional term contributes nothing to the P-portion, and the calculation above is unchanged.

Where does this leave us? The RAEA defence of Lewisian halfing is its most plausible defence, since it allows for a disanalogy objection against my dreaming argument, applies to the original problem and all discussed variants, and does not have the same counterintuitive implications as AEA updating. However, the commitments of RAEA updating plausibly lead one to adopt a position that is *not* Lewisian halfing.

This puts the ball back in the Lewisian halfer's court. It is theoretically possible that there is another non-evidentialist defence of Lewisian halfing that allows for a disanalogy objection and that does not collapse to another position. It is up to Lewisian halfers to decide whether this is worth pursuing, or whether to abandon the position.

6. Conclusion

I gave a two-step argument based on simple diachronic Bayesian principles showing that Beauty should assign a credence less than 1/2 to the coin land-

[20] These values have been calculated using numerical approximation under the constraints $a + b + c = 1$ and $a, b, c \geq 0$.

[21] Note that maximizing total expected accuracy does not lead to an inconsistency between one-step and two-step maximization. Both lead to thirding.

ing heads. Apart from these principles – which Lewisian halfers traditionally accept – the argument relies only on highly plausible credence assignments while Beauty is dreaming. For Lewisian halfers, Sleeping Beauty's dream is a nightmare.

This does not refute double halfing, but double halfers have their own problems (Bradley, 2011; Conitzer, 2015; Pittard, 2015; Titelbaum, 2012). While it might have had some intuitive plausibility, the halfer position is becoming increasingly difficult to sustain.

References

Arntzenius, F. (2003). Some problems for conditionalization and reflection. *The Journal of Philosophy*, *100*(7), 356–370. https://doi.org/10.5840/jphil2003100729

Bostrom, N. (2007). Sleeping beauty and self-location: A hybrid model. *Synthese*, *157*(1), 59–78. https://doi.org/10.1007/s11229-006-9010-7

Bradley, D. J. (2003). Sleeping beauty: A note on dorr's argument for 1/3. *Analysis*, *63*(279), 266–268. https://doi.org/10.1111/1467-8284.00435

Bradley, D. J. (2011). Self-location is no problem for conditionalization. *Synthese*, *182*(3), 393–411. https://doi.org/10.1007/s11229-010-9748-9

Briggs, R. (2010). Putting a value on beauty. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (pp. 3–34, Vol. 3). Oxford University Press.

Conitzer, V. (2015). A devastating example for the halfer rule. *Philosophical Studies*, *172*(8), 1985–1992. https://doi.org/10.1007/s11098-014-0384-y

Dorr, C. (2002). Sleeping Beauty: in defence of Elga. *Analysis*, *62*(4), 292–296. https://doi.org/10.1093/analys/62.4.292

Draper, K., & Pust, J. (2008). Diachronic dutch books and sleeping beauty. *Synthese*, *164*(2), 281–287. https://doi.org/10.1007/s11229-007-9226-1

Elga, A. (2000). Self-locating belief and the Sleeping Beauty problem. *Analysis*, *60*(2), 143–147. https://doi.org/10.1093/analys/60.2.143

Elga, A. (2004). Defeating dr. evil with self-locating belief. *Philosophy and Phenomenological Research*, *69*(2), 383–396. https://doi.org/10.1111/j.1933-1592.2004.tb00400.x

Hedden, B. (2015). Time-slice rationality. *Mind*, *124*(494), 449–491. https://doi.org/10.1093/mind/fzu181

Horgan, T. (2004). Sleeping Beauty awakened: new odds at the dawn of the new day. *Analysis*, *64*(1), 10–21. https://doi.org/10.1093/analys/64.1.10

Horgan, T. (2007). Synchronic Bayesian updating and the generalized Sleeping Beauty problem. *Analysis*, *67*(1), 50–59. https://doi.org/10.1093/analys/67.1.50

Kierland, B., & Monton, B. (2005). Minimizing inaccuracy for self-locating beliefs. *Philosophy and Phenomenological Research*, *70*(2), 384–395. https://doi.org/10.1111/j.1933-1592.2005.tb00533.x

Kim, N. (2021). Limitation of analogy in the sleeping beauty debate: The case of singer's argument. *Synthese*, *198*(11), 10513–10528. https://doi.org/10.1007/s11229-020-02733-w

Kim, N. (2022). Sleeping beauty and the current chance evidential immodest dominance axiom. *Synthese*, *200*(6), 458. https://doi.org/10.1007/s11229-022-03926-1

Kim, N. (2024). Sleeping beauty and the evidential centered principle. *Erkenntnis*, *89*(5), 2073–2095. https://doi.org/10.1007/s10670-022-00619-6

Lewis, D. (2001). Sleeping Beauty: reply to Elga. *Analysis*, *61*(3), 171–176. https://doi.org/10.1093/analys/61.3.171

Meacham, C. J. (2008). Sleeping beauty and the dynamics of de se beliefs. *Philosophical Studies*, *138*(2), 245–269. https://doi.org/10.1007/s11098-006-9036-1

Milano, S. (2022). Bayesian beauty. *Erkenntnis*, *87*(2), 657–676. https://doi.org/10.1007/s10670-019-00212-4

Moss, S. (2015). Time–slice epistemology and action under indeterminacy. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology volume 5*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198722762.003.0006

Pittard, J. (2015). When beauties disagree: Why halfers should affirm robust perspectivalism. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (pp. 195–204, Vol. 5). Oxford University Press.

Pust, J. (2008). Horgan on sleeping beauty. *Synthese*, *160*(1), 97–101. https://doi.org/10.1007/s11229-006-9102-4

Pust, J. (2012). Conditionalization and essentially indexical credence. *The Journal of philosophy*, *109*(4), 295–315. https://doi.org/10.5840/jphil2012109411

Pust, J. (2013). Sleeping beauty, evidential support and indexical knowledge: Reply to horgan. *Synthese*, *190*(9), 1489–1501. https://doi.org/10.1007/s11229-011-9888-6

Pust, J. (2014). Beauty and generalized conditionalization: Reply to horgan and mahtani. *Erkenntnis*, *79*(3), 687–700. https://doi.org/10.1007/s10670-013-9532-6

Schoenfield, M. (2017). Conditionalization does not (in general) maximize expected accuracy. *Mind*, *126*(504), 1155–1187. https://doi.org/10.1093/mind/fzw027

Schwarz, W. (2025). Sleeping beauty and the demands of non-ideal rationality. *Noûs*, *59*(4), 1072–1092. https://doi.org/10.1111/nous.12545

Titelbaum, M. G. (2008). The Relevance of Self-Locating Beliefs. *The Philosophical Review*, *117*(4), 555–606. https://doi.org/10.1215/00318108-2008-016

Titelbaum, M. G. (2012). An embarrassment for double-halfers. *Thought: A Journal of Philosophy*, *1*(2), 146–151. https://doi.org/10.1002/tht3.21

Titelbaum, M. G. (2013a). *Quitting certainties: A bayesian framework modeling degrees of belief*. Oxford University Press.

Titelbaum, M. G. (2013b). Ten reasons to care about the sleeping beauty problem. *Philosophy Compass*, *8*(11), 1003–1017. https://doi.org/10.1111/phc3.12080

Weintraub, R. (2004). Sleeping beauty: A simple solution. *Analysis*, *64*(1), 8–10. https://doi.org/10.1093/analys/64.1.8