A Nightmare for Lewisian Halfers

Lennart B. Ackermans

17th January 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, it 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.

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 upon 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

were to land 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. Beauty was told how the experiment would go on Sunday. 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, is 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*, taking up the majority of published positions, argue that the solution is $1/3$ (among these are Dorr, 2002; Elga, 2000; Horgan, 2004, 2007; Kim, 2022a, 2022b; 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 update her credences based on 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 claim that this credence becomes greater than $1/2$ after being told it is Monday. Since this result is a bitter pill to swallow, Lewisian halfism is regarded as implausible by many – while still a coherent position. My aim in this article is to show that the situation is worse: Lewisian halfism may be inconsistent.

The central motivation for Lewisian halfism 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 sufficiently strong 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. If this argument succeeds, then Lewisian halfism is inconsistent. One must either accept thirdism, switch to another branch of halfism such as double halfism, or provide a fundamentally new defence of Lewisian halfism.

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 highly plausible premises about Beauty's dreaming credences. Hence, this result disproves Lewisian halfism.

If we add two plausible assumptions about Beauty's credences while sleeping, we can further derive the thirder position (section 3). Hence, this is also a novel argument for thirding, bearing similarities to existing arguments (Horgan, 2004, 2007; Milano, 2022). Unlike these existing 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 (2013). However, these arguments are open 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 in 1697.[1] Hence, the objector who claims that Beauty does not have dreams is offering a version that is seemingly inconsistent with the original.

In response, the Lewisian halfer might of course insist that she has always understood Beauty's sleep to be dreamless, and object that my version is disanalogous to that version. I turn to this objection in section 4, and argue that it is unlikely to succeed.

Finally, in section 5, I discuss a defence for Lewisian halfism by Schwarz (2025), which uses an updating principle based on an expected accuracy principle instead of the principles used by earlier Lewisian halfers. I show that Schwarz's version of Lewisian halfism is equally affected by a dreaming Beauty, since the same accuracy principle can be used to defend both thirding and halfing.

[1] As far as I'm 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 Lewisian halfism

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 (disregarding the possible incapacitation of the memory loss that results from the amnesia-inducing drug).[3]

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[4] 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,

[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] Two objections may be raised against this assumption (thanks to a reviewer for pointing this out). The first is that having dreams without incapacitation makes the scenario disanalogous to the original: see section 4 for further discussion. The second is that the norms of rationality might not require treating dreams the same as waking stages *even if* there is no incapacitation. I'm not sure how this can be motivated.

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

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 $P_*(H) = P_*(\neg H) = 1/2$.

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 to 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 "irrelevant" (according to Lewis) evidence of being woken up today. 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. 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 belief 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 use 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 is relevant for Beauty's belief in the outcome of the coin toss during her dream. But it is clearly 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 has 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 $P_*(H_1) < 1/2$.

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 has not changed 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 halfism, 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 use Bayesian conditionalization as normally. 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 halfism.

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.[5] Moreover, the alternative update that would be required to save Lewisian halfism has the counterintuitive consequence that all the dreaming credences assigned to H_2 are assigned to H_1 after waking up.[6]

[5] 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.

[6] 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 $H_2 \vee \cdots \vee H_{30}$. It seems you should now become quite sure of tails, as Bayesian conditionalization requires. The updating rule required to save Lewisian halfism instead leads you to assign an absurdly high credence to H_1 (heads and it's June 1) of 1/2.

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 thirdism

We can derive thirding by adding two additional plausible assumptions.

3.1. Beauty's dreaming credence that it is Monday

During Beauty's dream, she is uncertain which day it is. Symmetry considerations would 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 part of the same possible world (or sets of possible worlds) in which the coin lands heads. 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 independent events. 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.

Hence, Beauty's credences in M and H should be independent. Hence, using the dreaming credences argued for 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 argument itself and concede that *if* Sleeping Beauty has a dream, her credence in heads will become less than $1/2$ after waking up. But she might insist that if Sleeping Beauty does not dream, or is rationally incapacitated during her dreams, her credence in heads is $1/2$ after waking up. Hence, my version of the problem description would be disanalogous to a version in which it is added that Sleeping Beauty doesn't dream. And it's the latter version that the Lewisian halfer was always interested in!

4.1. Relation to other arguments by analogy

The disanalogy objection might have worked for previous arguments for thirdism 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 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 can't rule out and 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.[7] Second, Schwarz (2025, p. 1086) argues that the variants by Dorr and Arntzenius are disanalogous to the original because maximizing what he calls "average expected accuracy" with respect to Sunday recommends Lewisian halfism in the original scenario but thirding in the variants. This objection does not apply to the dreaming scenario: as I discuss in section 5 below, maximizing average expected accuracy with respect to Sunday also recommends

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

Lewisian halfism when Beauty has lucid dreams.

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, in two separate scenarios (real or hypothetical), an agent's relevant evidence is identical, then the set of rationally permissible credences in both scenarios is the same.[8] 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 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, as does Schwarz (2025). It then becomes possible that the way and order in which an agent comes to learn – and forget – evidence matters for norms of rationality.

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

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. It is unclear what such a reason could be. But even if it exists, a second problem immediately arises: in the dreaming scenario, Beauty could just as well ignore her dreams and directly update 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 version of Lewisian halfism defended by Schwarz (2025), discussed in the next section, is an example of exactly such a conflict.

The objector might be able to solve the problem by requiring that a rational agent always uses her last available degrees of belief for updating. Hence, 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 this raises a similar question: if Beauty did not dream, she would still be able to reason that thirding would be rationally required if she had dreamt. So if updating with respect to the dreaming credences is more rational than updating with respect to the Sunday credences if she dreams (as the objector is arguing), then shouldn't it still be more rational when she doesn't dream but is aware of the more rational credences if she did?

5. The expected accuracy argument for Lewisian halfism

Schwarz (2025) offers a new argument for Lewisian halfism that is entirely motivated by a (non-standard) diachronic accuracy principle. Since this argument is motivated by a new principle of rationality not used by earlier Lewisian halfers, it merits a separate treatment.

Schwarz maintains that the Sleeping Beauty problem is a problem of non-ideal rationality due to Beauty's amnesia and the unavailability of Bayesian conditionalization between Sunday and Monday. Schwarz's accuracy principle is proposed as a principle of non-ideal rationality to be used when more ideal principles are unavailable. It therefore seems that Schwarz is affected by my refutation of Lewisian halfism, which relies on the ideal (and therefore more rational) principle of Bayesian conditionalization (or

a closely related ideal principle that Schwarz calls "shifted conditionalization"[9]).

The Lewisian halfer might try to use this conflict between the ideal and non-ideal to make a disanalogy objection: when Beauty dreams, she has tools of ideal rationality available, and must therefore be a thirder; when she doesn't, Schwarz's non-ideal principle requires halfing. But this disanalogy argument clearly succumbs to the problem I raised at the end of the previous section. Credences required by ideal rationality are better than credences required by non-ideal rationality, and Beauty is capable of deriving the ideal credences even if she doesn't dream. Hence, she should be a thirder either way.

Alternatively, the Lewisian halfer might come up with some reason that Beauty is bound only by non-ideal principles of rationality even when she dreams (perhaps because of the possibility of amnesia, which affects her dreaming credences). I'm not sure if this response can be motivated. But even if it can, Schwarz's accuracy principle faces a serious problem: in the dreaming scenario, the principle can be used to defend both halfing and thirding. Hence, the principle seems to be in conflict with itself.[10]

In Schwarz's framework, propositions are centred 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, you are Beauty, and it is Tuesday.

An updating rule based on accuracy aims to maximize the expected accuracy of the agent's credences at the later time using their credences at the earlier time. After being put to sleep on Sunday Beauty undergoes what Schwarz calls "doxastic fission": there are two possible later times at which the update will take place (Monday and Tuesday). Doxastic fission raises the question of how expected accuracy is to be calculated. Schwarz shows that maximizing *total expected accuracy* between Sunday and Monday or Tuesday recommends thirding, while his preferred rule of maximizing *average expected accuracy* 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".[11] For each centred world w located at the

[9] Shifted conditionalization requires that an agent who receives evidence E should set their new credence in any proposition A to their old credence that A will be the case after the update conditional on that E will be the case after the update (Schwarz, 2025, p. 1076).

[10] Thanks to a reviewer for alerting me to this conflict.

[11] 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 maximum at

earlier time of the update, the *doxastic successors* of W are defined as a set $s(w)$ of centred worlds which follow w at which the update can be carried out. For example, when Beauty wants to update between Sunday and Monday after waking up, the earlier world (heads, Beauty, Sunday) has two successors: (heads, Beauty, Monday) and (heads, Beauty, Tuesday). Lastly, an updating rule $\mu(w)$ assigns a probability function to a centred world w.

Suppose the earlier probability function P_1 is defined on a sample space Ω of centred worlds. 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, average expected accuracy sums over the accuracy of updated probabilities in possible fission states divided by the number of successors of w. As Schwarz shows, maximizing expected accuracy of 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 $P_+(H) = 2/3$. Hence, we end up with the probabilities of Lewisian halfism.

In the simple scenario in which the agent receives the same evidence no matter what the state of the world, we have $\mu(w) = P_2$ for some probability function P_2, for all $w \in s(\Omega)$. In this situation the average expected accuracy of P_2 given P_1 is

$$EV(P_2) = \sum_{w \in \Omega} P_1(w) \sum_{w' \in s(w)} \frac{V(P_2, w')}{|s(w)|}.$$

Updating between Sunday and dreaming (on Monday and Tuesday) is such a simple scenario, as is updating between dreaming and being woken up.

When Beauty is dreaming, the average expected accuracy with respect to her Sunday credences is

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

This is maximized when $P_*(H_1) = P_*(H_2) = P_*(T_1) = P_*(T_2) = 1/4$.[12] Hence, we get the dreaming credences required by the argument for thirding in section 3.

After waking up, Beauty updates with respect to her dreaming credences P_*. There is no doxastic fission and $s(H_2) = \emptyset$, so the average expected accuracy with respect to P_* is

---

$P' = P$.

[12] 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 $P_* = Pr$.

$$EV(P) = \frac{1}{4}V(P, H_1) + \frac{1}{4}V(P, T_1) + \frac{1}{4}V(P, T_2).$$

This is maximized for $P(H_1) = P(T_1) = P(T_2) = 1/3$.[13] It follows that $P(H_2) = 0$, and so we have $P(H) = 1/3$.

The proponent of average expected accuracy might try to get around the problem by requiring that one should always update based on one's last available degrees of belief. This would require thirding when Beauty dreams and halfing when she doesn't. One again ends up with a version of the problematic disanalogy objection discussed above. It is up to Lewisian halfer to explain how such a seemingly irrelevant detail can change the demands of rationality so substantially.

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 landing heads. Apart from these principles – which Lewisian halfers 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 halfism, 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

[13] $EV(P)$ is maximized if $\varphi(P) = 1/3V(P, H_1) + 1/3V(P, T_1) + 1/3V(P, T_2)$ is maximized. Let Pr be a probability function on the sample space $\{H_1, T_1, T_2\}$ such that $Pr(H_1) = Pr(T_1) = Pr(T_2) = 1/3$. Then we have $\varphi(P) = \sum_w Pr(w)V(P, w)$. Since V is strictly proper, $\varphi(P)$ is maximized by $P = Pr$.

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

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

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. (2022a). Sleeping beauty and the evidential centered principle. *Erkenntnis*, 1–23. https://doi.org/10.1007/s10670-022-00619-6

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

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 *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. (2013). *Quitting certainties: A bayesian framework modeling degrees of belief*. Oxford University Press.