

NO DOUBLE-HALFER EMBARRASSMENT: A REPLY TO TITELBAUM

Joel Pust,
University of Delaware
jpust@udel.edu

1. *Introduction*

No doubt you have heard about Sleeping Beauty (Elga 2000). She is a subject in an unusual experiment. She will be put to sleep on Sunday evening and woken up on Monday morning. On Monday afternoon, after being told that it is Monday, she will be returned to sleep and the experimenters will toss a fair coin. If the coin lands heads, they will let her sleep until Wednesday. If the coin lands tails, she will have her memories of Monday morning erased and will be woken up on Tuesday morning, spend the morning in a state subjectively indiscernible from Monday morning, and then returned to sleep until Wednesday. On Wednesday, she will be immediately informed that the experiment is over. All of this is known to Beauty before she is first put to sleep.

The Sleeping Beauty Problem is the problem of determining the correct answer to the question: “How confident should Beauty be, when she is awakened on Monday morning (uncertain whether it is Monday or Tuesday), that the coin toss comes up heads?”¹ According to “thirders” (Elga 2000; Dorr 2002; Horgan 2004; Titelbaum 2013b), when she awakens, Beauty should have a credence of $1/3$ that the coin toss comes up heads. According to “halfers,” she should have a credence of $1/2$ in that proposition. Halfers differ amongst themselves about the

¹ For a bibliography of the large and growing literature on the problem see <https://philpapers.org/browse/sleeping-beauty/>. Why care about the problem? See Titelbaum (2013a) for ten reasons.

answer to a second question: “How confident should Beauty be, after she is told that it is Monday, that the coin toss comes up heads?” According to Lewis (2001), she should then have a credence of $2/3$ in the relevant claim. Other halfers (Meacham 2008; Briggs 2010; Bostrom 2007), so-called “double-halfers,” agree with thirder that she should then have a credence of $1/2$ in the relevant claim. So, the main options are as follows:

	SUN	MON	MON
	PM	AM	PM
Thirder	$1/2$	$1/3$	$1/2$
Lewisian-halfer	$1/2$	$1/2$	$2/3$
Double-halfer	$1/2$	$1/2$	$1/2$

2. Titelbaum’s Challenge to the Double-Halfer

Titelbaum (2012) provides an ingenious variation on the standard Sleeping Beauty scenario and claims that it demonstrates a serious problem with the double-halfer position. In order to see the force of his case, one must understand how Titelbaum sees the dialectic between the various positions. By Titelbaum’s lights, the main motivation for the halfer position is the intuition that, upon awakening, Beauty has no information that would justify her having a credence other than $1/2$ that the coin toss lands heads. However, this intuition is in tension with the conjunction of (a) the stronger intuition that, after she has been told that it is Monday (and so knows that the coin toss is in the future), her credence that the coin toss lands heads ought to be $1/2$, and (b) the view that Beauty’s credence on Monday afternoon should be

produced by conditionalizing on the fact that it is Monday.² Both (a) and (b) are typically endorsed by thirders.

Lewis accepted (b) and so was forced to deny (a), holding that, upon learning that it is Monday, Beauty acquires evidence (“inadmissible evidence” in his terms) which justifies her in having a credence of $2/3$ that the coin lands heads even though she then knows that the coin toss is in the future and that its outcome has no influence on her present situation. As Titelbaum sees things, many halfers do not want to join Lewis in this view because it violates their original intuition that nothing Beauty learns during the experiment is sufficient to justify a credence other than $1/2$ for the coin toss landing heads. As a result, double-halfers repudiate (b) and hold that Beauty is not obligated to update by conditionalization upon learning that it is Monday.

Titelbaum’s variant scenario is supposed to show that double-halfers must, *by their own lights*, allow that evidence of the sort Beauty gains during the scenario can sometimes justify a credence other than $1/2$ for the claim that a fair coin toss known to be in the future lands heads, just the sort of embarrassing commitment double-halfers aimed to avoid in parting ways with Lewis by denying that Beauty ought to conditionalize upon learning that today is Monday. Titelbaum’s variation merely adds to the original Sleeping Beauty scenario an additional toss of the fair coin by the experimenters on Tuesday. This second toss has no influence on anything to do with Beauty’s being awake or asleep. Titelbaum’s troublesome new question is this: “How

² This is because, so long as Beauty assigns a greater than zero credence to its being Tuesday, her credence that the coin lands heads, conditional on its being Monday, must be greater than $1/2$.

confident should Beauty be, when she is awakened on Monday (uncertain what day it is), that *today's* coin toss comes up heads?"

As Titelbaum shows, by the probability calculus, her credence in that claim must be greater than her credence that the Monday coin toss comes up heads. Since Beauty is certain that *today's* coin toss comes up heads just in case either it is Monday and the Monday toss comes up heads or it is Tuesday and the Tuesday toss comes up heads, her credence that *today's* coin toss comes up heads must be the sum of her credences in those two conjunctive claims. If, as halfers (of both the Lewisian and double variety) claim, Beauty's credence, upon awakening, that the Monday coin toss comes up heads, is $1/2$, and as the Monday coin toss coming up heads entails, given Beauty's background knowledge that she is now awake, that it is Monday, it follows, given that Beauty has a non-zero credence that it is Tuesday and the Tuesday coin toss comes up heads, that halfers must claim that Beauty's credence that *today's* coin toss comes up heads must be greater than $1/2$.³

This implication of the halfer position is, according to Titelbaum, unacceptable. As he puts it, incredulously, "Beauty is certain that *today's* coin is fair, is certain that it hasn't been flipped yet, and has no information about its outcome. Yet the halfer needs her to be greater than $1/2$ confident that *today's* flip will come out heads" (148). This result, he takes to be "ridiculous" and "embarrassing." By contrast, the thirder has no difficulty holding that Beauty's credence that *today's* coin toss lands heads is $1/2$. Indeed, standard thirder views yield exactly this verdict.

³ Peter Lewis (2010) is a halfer who might reject this reasoning as he maintains that even if P entails Q, P(Q) can be less than P(P).

Titelbaum is correct that the double-halfer is committed by her views to the claim that Beauty ought, during her awakenings in the Two-Toss scenario, to have a credence greater than $1/2$ in the proposition she would then express with “today’s coin toss lands heads.” However, I shall argue that this result is not any *distinctive* embarrassment for the double-halfer, as there appears to be an equally good reason of the same sort to think that Beauty ought *not to* have a credence of $1/2$ in the relevant claim. More precisely, I will show that direct inference from objective probabilities does not clearly justify a credence of $1/2$ (as implied by the thirder position) nor a greater than $1/2$ credence (as implied by the halfer position). Rather, some central principles of direct inference yield the conclusion that neither credence is ultimately justified by direct inference because the case for each one serves to undermine the case for the other. I leave open whether the credence required by either the thirder or halfer positions can be ultimately justified by a more complete theory of direct inference.⁴ I also leave open the relationship between direct inference principles and other non-coherence constraints on rational credences.

3. *Direct Inference and Admissibility*

Though many who have developed theories of direct inference are not Bayesians, from a Bayesian perspective, direct inference involves normatively constraining the value of some credences by objective probabilities. Theories of direct inference vary significantly in the way in which this core idea is developed. Many take the objective probabilities to attach to

⁴ For a recent attempt to directly justify the $1/3$ answer to the standard Sleeping Beauty Problem by appeal to direct inference, see Draper 2021.

propositional functions or open formulae (Kyburg 1974; Levi 1974; Pollock 1990) while others take them to attach to propositions (Lewis 1980). They differ as well on their account of the objective probabilities, variously appealing to objective chance (Lewis 1980; Levi 1974), frequency (Kyburg 1974), expected frequency (Thorn 2012), or nomic probability (Pollock 1990).

Given the wide variety of approaches, it is difficult to discuss direct inference in an ecumenical and reasonably concise fashion. As a result, I will here attempt to develop my argument utilizing a schematic account of direct inference loosely based on Lewis' well-known Principal Principle in appealing to objective chance, though one allowing that objective probabilities attach in the first instance to open formulae or propositional functions. My own view is that the picture of direct inference provided by Pollock (and, to some extent, Kyburg and Thorn) is often more perspicuous. However, in deference to the significantly wider contemporary familiarity with the broadly Lewisian framework, I opt for that framework here, trusting that the arguments I give would translate into alternative frameworks.

If we take credal probability or rational credence to be indicated using 'PROB' and objective probability to be indicated using 'prob' we might take the central simplified principle of direct inference to require that $\text{PROB}(Ta/\text{prob}(Tx/Rx) = n \ \& \ Ra \ \& \ E) = n$, when E is consistent with and *admissible* with respect to $\text{prob}(Tx/Rx) = n \ \& \ Ra$.⁵ If we take the sort of objective probability at issue to be chance, then this requires that one's credence in *Ta*, conditional on the claim that the chance that a system in initial state *R* will acquire property *T* is *n* and that system

⁵ This is much like Lewis' Principal Principle, though I here aim for a simpler version of Wallman and Hawthorne's "Generic Direct Inference Principle" (2020, 958). This example uses properties of individuals but the same principles hold for multi-place properties or relations.

a is in state R , should be n , provided that the remainder of one's total evidence, E , is admissible with respect to the aforementioned.

This core idea is at best schematic (and poised between vacuity and contradiction) without an account of admissibility. In particular, a theory of direct inference that can handle anything more than the simplest cases must take account of cases in which systems may be in multiple states with *different* chances of acquiring the attribute T . Working out the proper response to such cases (and hence, in Lewis' framework, partly explicating the notion of admissibility) requires an account which of those states is, in a given epistemic situation, the one (or ones), if any, which should ultimately be relied upon to determine the rational credence that the individual system acquires the target property.

Since Reichenbach's (1949) pioneering work, theorists of direct inference have typically agreed that two sorts of such cases are of particular interest. The first sort of case involves initial chance states which are such that one is logically stronger than the other and the second sort of case involves chance states which are not such that one is logically stronger than the other.

In the first sort of case, it is commonly accepted that one's credence ought to be based on the chances attaching to the logically stronger property. Hence, if, in addition to knowing that system a is in initial chance state R , one also knows that the system is in chance state S (where being in S entails being R), and that the chance that in chance state S has the target property is m ($\neq n$), then the following is a required conditional probability: $\text{PROB}(Ta/\text{prob}(Tx/Rx) = n \ \& \ Ra \ \& \ \text{prob}(Tx/Rx \ \& \ Sx) = m \ \& \ Sa \ \& \ E) = m$, where E is admissible. Coupled with a total evidence requirement, this principle requires that one's unconditional credences are, in the absence of

inadmissible evidence, constrained by the logically strongest state such that one knows the target system to occupy the state and one knows the relevant chances. In this way, one can capture Reichenbach's (1949, 374) suggestion that one ought, when engaging in direct inference, to base one's inference on the logically strongest reference property such that one knows that the individual has the property and one knows the objective probability that something with the known property has the property of interest.⁶

The second sort of case requiring qualification of the core principle of direct inference are cases in which a system is in two states which yield different chances but neither of those states is logically stronger than the other. If $\text{prob}(Tx/Sx) = m$, and $m \neq n$, and S is neither logically stronger than R nor logically weaker than R , then $\text{PROB}(Ta/\text{prob}(Tx/Rx) = n \ \& \ Ra \ \& \ \text{prob}(Tx/Sx) = m \ \& \ Sa \ \& \ E)$ need not be equal to m or n . (It might be equal to m or n , but not based on direct inference.) In such a situation, where $m \neq n$, if E is admissible for $\text{PROB}(Ta/\text{prob}(Tx/Rx) = n \ \& \ Ra \ \& \ \text{prob}(Tx/Sx) = m \ \& \ E) = n$, and for $\text{PROB}(Ta/\text{prob}(Tx/Sx) = m \ \& \ Sa \ \& \ \text{prob}(Tx/Rx) = n \ \& \ E) = m$, then Ra and Sa are jointly inadmissible relative to E .

Again, given a total evidence requirement, this yields the result that if one knows a system to occupy two chance states such that neither is logically stronger than the other and one knows the relevant chances to differ, then one's unconditional credence that some system in the chance states has the property of interest is unconstrained by those two chance claims. In this

⁶ This suggestion is not without difficulties. Some worries about it from a Bayesian perspective are explored by Wallman and Hawthorne (2020, Section 5.2). In addition, Pollock and Kyburg hold that a direct inference can also be defeated by knowledge of the objective probabilities attaching to logically stronger properties which are not known to be had by the individual but are known to be exclusive and exhaustive so that the individual must have one of them. Pollock calls such cases "joint subset defeaters."

way, one may capture Reichenbach's (1949, 374) claim that in such situations we ought not to adopt any credal probability on the basis of direct inference.

With this short sketch of some basic principles of direct inference in place, we can see that Titelbaum's argument involves a highly intuitive direct inference linking the chance that a toss of a fair coin known to be in the future lands heads to Beauty's credence that said coin toss lands heads. While he explicitly appeals to the Principal Principle, it also seems clear (see the quotations above) that Titelbaum regards the inference in question as intuitively extremely plausible, perhaps so plausible as to require accommodation by any acceptable theory of direct inference.⁷

In order to facilitate precision in what follows, let us adopt the following notation:

'Toss(x , s)' means ' x is a fair coin toss occurring in scenario s (of some sort) *on a day when Beauty is awake*' (an 'awakening day'),' $H(x, s)$ ' means ' x lands heads in scenario s ,' and ' Hx ' means ' x lands heads.' Where σ is the scenario in which Beauty is a subject, and τ is *today's* coin toss in σ , we can then see Titelbaum's argument as based on knowledge of the following:

$$(1) \quad \text{prob}(H(x, s)/\text{Toss}(x, s)) = 1/2,$$

which given knowledge of $H(\tau, \sigma)$ and only admissible further total evidence E yields,

⁷ Elsewhere, Titelbaum disavows commitment to the Principal Principle in its full generality (2013b, 43 & 215). For some general concerns about the usefulness of Lewis' Principal Principle for direct inference, see Kyburg (1981). For concerns about its usefulness specifically to settle the standard Sleeping Beauty Problem, see Dorr (2002, 293).

$$(2) \quad \text{PROB}(H\tau) = 1/2$$

by the aforementioned principle of direct inference.⁸

4. *Beauty's Inadmissible Evidence*

As noted above, whether (1) and accompanying knowledge that $H(\tau, \sigma)$ should determine Beauty's credence in heads by direct inference depends on whether Beauty has inadmissible evidence. I will argue that she does. In particular, she has, I maintain, evidence that would alone justify an alternative credence and so is inadmissible, relative to (1) and $H(\tau, \sigma)$. This evidence is as follows: Beauty knows upon awakening (just as she did on Sunday) that she is in a Two-Toss Sleeping Beauty scenario of the sort specified by Titelbaum. Moreover, she knows that the chance that a Two-Toss Sleeping Beauty system has a single toss occurring on an awakening day and that that single toss on an awakening day lands heads is $1/2$. She knows also that the chance that a Two-Toss Sleeping Beauty scenario has one awakening day toss landing heads and another awakening day toss landing tails is $1/4$. Finally, she knows that the

⁸ It might be claimed that Titelbaum's inference fails because the claim at issue is a "*de se*" or "self-locating" claim and that the Principal Principle doesn't clearly apply to such claims. This seems to me too quick as *some* self-locating claims seem sensibly accorded objective chances, especially some of those Lewis (1979) would regard as reductions of *de re* claims. While I would agree that it is dubiously coherent to hold that "today is Monday" has an objective chance, it seems quite sensible to suppose, with Titelbaum, that "today's coin toss lands heads" has one. Even if I have lost track of time, I might well know that a coin I am now holding is fair and, absent inadmissible evidence, adopt a credence of $1/2$ that its next toss will land heads. (I thank an anonymous referee for raising this issue.)

chance that a Two-Toss Sleeping Beauty scenario has no awakening day toss landing heads is 1/4.

Conditional on the first of these situations obtaining, the epistemic probability that today's toss lands heads is 1 because Beauty knows that she is awake today and this requires, given the first situation with its single observed toss, that the single awakening day toss land heads. Conditional on the second of these situations obtaining, the probability that today's toss lands heads is equal to the probability that today's toss is the single awakening day toss (of the two observed in the scenario) which lands heads. This is certainly greater than 0 and, given a minimal indifference principle, is plausibly 1/2. Considering, then, both the probabilities that her current Two-Toss Sleeping Beauty scenario is either of the two sorts mentioned and the conditional probabilities that today's toss lands heads conditional on those two possibilities, she should conclude that the probability that today's toss lands heads is certainly greater than 1/2 and is plausibly 5/8.

More formally, letting 'TSBs' means 's is a Two-Toss Sleeping Beauty scenario' (of the sort specified by Titelbaum), Beauty knows:

- (3) $\text{prob}(\exists!x(\text{Toss}(x, s) \ \& \ Hx \ \& \ \sim\exists y(y \neq x \ \& \ \text{Toss}(y, s)))/\text{TSBs}) = 1/2$
- (4) $\text{prob}(\exists!x(\text{Toss}(x, s) \ \& \ Hx) \ \& \ \exists!y(\text{Toss}(y, s) \ \& \ \sim Hy)/\text{TSBs}) = 1/4$
- (5) $\text{prob}(\sim\exists x(\text{Toss}(x, s) \ \& \ Hx)/\text{TSBs}) = 1/4.$

So, by the principle of direct inference earlier articulated, direct inference based on (3) – (5) and TSB_{σ} would, *if* the remainder of Beauty’s evidence is admissible, yield the following credal probabilities:

$$(6) \quad \text{PROB}(\exists!x(\text{Toss}(x, \sigma) \ \& \ Hx \ \& \ \sim\exists y(y \neq x \ \& \ \text{Toss}(y, \sigma)))) = 1/2$$

$$(7) \quad \text{PROB}(\exists!x(\text{Toss}(x, \sigma) \ \& \ Hx) \ \& \ \exists!y(\text{Toss}(y, \sigma) \ \& \ \sim Hy)) = 1/4$$

$$(8) \quad \text{PROB}(\sim\exists x(\text{Toss}(x, \sigma) \ \& \ Hx)) = 1/4.$$

Now, since $\text{Toss}(\tau, \sigma)$ is certain and (6) – (8) attach to an exclusive and exhaustive set of possible outcomes for coin tosses in σ , it follows (by the law of total probability) that

$$(9) \quad \text{PROB}(H\tau) = (1/2) \text{PROB}(\exists!x(\text{Toss}(x, \sigma) \ \& \ Hx \ \& \ x = \tau) / \exists!x(\text{Toss}(x, \sigma) \ \& \ Hx) \ \& \ \sim\exists y(y \neq x \ \& \ \text{Toss}(y, \sigma)))) + (1/4) \text{PROB}(\exists!x(\text{Toss}(x, \sigma) \ \& \ Hx \ \& \ x = \tau) / \exists!x(\text{Toss}(x, \sigma) \ \& \ Hx) \ \& \ \exists!y(\text{Toss}(y, \sigma) \ \& \ \sim Hy)).$$

As noted above, in (9), the first conditional probability has a value of 1. So, given that the second conditional probability is greater than 0, it follows that $\text{PROB}(H\tau) > 1/2$. Moreover, given a suitable centered indifference principle, the second conditional probability is equal to 1/2, and it follows that $\text{PROB}(H\tau) = 5/8$.⁹ Either result contradicts (2).

⁹ The envisaged centered indifference principle would permit equal credence that today’s coin toss is the single observed toss landing heads and that today’s coin toss is the single observed coin toss landing tails each conditional on the scenario being one with an observed heads toss and an observed tails toss.

The result is that direct inferences supported by (1) and (3) – (5) separately yield contradictory probabilities for $H\tau$. Moreover, and this is the crucial point, *none of these inferences is based on a logically stronger state (of one and the same chance system) than another inference*. The direct inference from (1) is based on states of scenario-coin toss pairs and the direct inferences from (3) – (5) on a states of scenarios.¹⁰ A property or state of pairs cannot be logically stronger or weaker than a property of individuals and so the notion of one state being logically stronger than another cannot be invoked to justify a direct inference appealing to one over a direct inference appealing to another.

The result is that this is a case in which $TSB\sigma$ is inadmissible with respect to $Toss(\tau, \sigma) \& (1) \& (3) - (5) \& E$ while $Toss(\tau, \sigma)$ is inadmissible with respect to $TSB\sigma \& (1) \& (3) - (5) \& E$. Since Beauty's total evidence is $TSB\sigma \& Toss(\tau, \sigma) \& (1) \& (3) - (5) \& E$, no credence in $H\tau$ is justified by direct inference in Titelbaum's variant scenario.¹¹ To put the point another way, her knowledge regarding the chances that the Two-Toss Sleeping Beauty scenario would yield various outcomes is inadmissible relative to the direct inference invoked by Titelbaum and her knowledge regarding the chances that a fair coin toss lands heads is inadmissible relative to the remainder of her total knowledge. The result, given the basic principles of direct inference, is

¹⁰ The same result would follow if we were to treat Titelbaum's direct inference as based solely on a property of fair coins and so were comparing chances of individual coin tosses landing heads with chances of two-toss scenarios having various numbers of observed tosses landing heads.

¹¹ Interestingly, it seems that, given Beauty's background knowledge, E, upon awakening, she couldn't have lacked the total knowledge she has and so her situation is one in which she couldn't have lacked the inadmissible evidence.

that her rational credence that today's toss lands heads cannot be determined by direct inference.

The upshot is that the credence required by the double-halfer in Titelbaum's new scenario is *not* one that contravenes some generally accepted principles governing how chances should constrain credences because, surprisingly, such principles seem to imply that Titelbaum's scenario is one in which they do not require a specific credence that today's coin lands heads. However, these general principles do not permit the Monday afternoon credence required by the Lewisian on the original Sleeping Beauty problem as there is no defensible chance assignment on which the chance that the single Monday toss lands heads is $2/3$ which can be said to defeat that derived from the fact that the Monday toss is a toss of a fair coin. Since, according to standard principles of direct difference, there is a relevant difference between "the (single) coin toss lands heads" in the original scenario and "today's toss lands heads" in Titelbaum's variant, the double-halfer need not worry that her commitments on Titelbaum's variant case are at odds with her rejection of the Lewisian halfer position.¹²

5. Conclusion

I have argued that Titelbaum is mistaken in thinking that direct inference yields $\text{PROB}(H\tau) = 1/2$ in his variant scenario, at least if standard general principles of direct inference

¹² See the Oscar Seminar's "An Objectivist Argument for Thirdism" (2008) for an attempt to defend directly the thirder position on the original Sleeping Beauty problem by appeal to direct inference. Pust (2011) argues against their case by suggesting, like the present paper, that they neglect inadmissible evidence. See Thorn (2011) for a rejoinder to Pust, Draper (2017) for a rejoinder to Thorn, and Horgan (forthcoming) for a rejoinder to Draper.

are supposed to warrant such a conclusion. While Beauty knows today's coin toss to be fair, Titelbaum fails to account for the evidential relevance of the fact that Beauty also knows herself and today's coin toss to be part of a larger chance system which supports a conflicting value for $\text{PROB}(H\tau)$. Lacking a complete theory of direct inference, I have not argued that direct inference ultimately supports the credences mandated by the double-halfer. Rather, I have argued that, given some widely accepted principles of direct inference, it appears that *neither* the halfer *nor* the thirder can provide a straightforward undefeated direct inference for the credence they require of Beauty in Titelbaum's variant case. Still, in the absence of some yet-to-be-produced reason to think that an acceptable direct inference to $\text{PROB}(H\tau) = 1/2$ is available in the variant case or justification for the view that $\text{PROB}(H\tau) = 1/2$ is more plausible than the theoretical principles of direct inference on which I have relied, my argument suffices to justify concluding that, contrary to Titelbaum, the double-halfer need not be more embarrassed that her view implies $\text{PROB}(H\tau) \neq 1/2$ than Titelbaum's thirder ought to be embarrassed that her view implies that $\text{PROB}(H\tau) = 1/2$.¹³

¹³ For helpful comments on previous drafts of this paper, I thank [names and locations omitted for blind review] and three anonymous referees.

References

- Bostrom, N. (2007). "Sleeping Beauty and Self Location: A Hybrid Model." *Synthese* 157: 59-78.
- Briggs, R. (2010). "Putting a Value on Beauty." In T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology* 3: 3-34, Oxford: Oxford University Press.
- Dorr, C. (2002). "Sleeping Beauty: In Defense of Elga." *Analysis* 62, 53-62.
- Draper, K. (2017). "Even for Objectivists, Sleeping Beauty Isn't So Simple." *Analysis* 77, 29-37.
- Draper, K. (2021). "Direct Inference and the Sleeping Beauty Problem." *Synthese* 198: 2253-2271.
- Elga, A. (2000). "Self-Locating Belief and the Sleeping Beauty Problem." *Analysis* 60, 143-47.
- Horgan, T. (2004). "Sleeping Beauty awakened: new odds at the dawn of the new day." *Analysis* 64: 10-21.
- Horgan, T. (forthcoming). "Direct Inference, Reichenbach's Principle, and the Sleeping Beauty Problem." *Episteme*.
- Kyburg, H. (1974). *The Logical Foundations of Statistical Inference*. Boston: Reidel.
- Kyburg, H. (1981). "Principle Investigation." *Journal of Philosophy* 78, 772-78.
- Levi, I. (1977). "Direct Inference." *Journal of Philosophy* 74, 5-29.
- Lewis, D. (1979). "Attitudes *de dicto* and *de se*." *Philosophical Review* 88: 513-543.
- Lewis, D. (1980). "A Subjectivists Guide to Objective Chance." In Jeffrey (ed.), *Studies in Inductive Logic and Probability*, Vol. 2, 263-94. Berkeley: University of California Press.
- Lewis, D. (2001). "Sleeping Beauty: Reply to Elga." *Analysis* 61, 171-76.
- Lewis, P. (2010). "Credence and Self-Location." *Synthese* 175, 369-82.

- Meacham, C. J. G. (2008). "Sleeping Beauty and the Dynamics of De Se Beliefs." *Philosophical Studies* 138: 245-270.
- Pollock, J. (1990). *Nomic Probability and the Foundations of Induction*. New York: Oxford University Press.
- Pust, J. (2011). "Sleeping Beauty and Direct Inference." *Analysis* 71, 290-93.
- Reichenbach, H. (1949). *A Theory of Probability*. Berkeley: University of California Press.
- Seminar, O. (2008). "An Objectivist Argument for Thirdism." *Analysis* 68, 149-55.
- Titelbaum, M. (2012). "An Embarrassment for Double-Halfers." *Thought* 1, 146-51.
- Titelbaum, M. (2013a). "Ten Reasons to Care About The Sleeping Beauty Problem." *Philosophy Compass* 8, 1003-11.
- Titelbaum, M. (2013b). *Quitting Certainties*. New York: Oxford University Press.
- Thorn, P. (2011). "Undercutting Defeat via Reference Properties of Differing Arity: A Reply to Pust." *Analysis* 71, 662-67.
- Thorn, P. (2012). "Two Problems of Direct Inference." *Erkenntnis* 76, 299-318.
- Wallmann, C. & James Hawthorne. (2020). "Admissibility Troubles for Bayesian Direct Inference." *Erkenntnis* 85: 957-993.