

Immortal Beauty: Does Existence Confirm Reincarnation?

Jens Jäger

To cite this article: Jens Jäger (2022) Immortal Beauty: Does Existence Confirm Reincarnation?, Australasian Journal of Philosophy, 100:4, 789-807, DOI: [10.1080/00048402.2021.1938150](https://doi.org/10.1080/00048402.2021.1938150)

To link to this article: <https://doi.org/10.1080/00048402.2021.1938150>



Published online: 09 Aug 2021.



Submit your article to this journal [↗](#)



Article views: 517



View related articles [↗](#)



View Crossmark data [↗](#)



Immortal Beauty: Does Existence Confirm Reincarnation?

Jens Jäger 

New York University

ABSTRACT

I argue that a popular view about self-locating evidence implies that there are cases in which agents have surprisingly strong evidence for their own reincarnation. The central case is an ‘Immortal Beauty’ scenario, modelled after the well-known Sleeping Beauty puzzle. I argue that if the popular ‘thirder’ solution to the puzzle is correct, then Immortal Beauty should be confident that she’s going to be reincarnated. The essay also examines another pro-reincarnation argument due to Michael Huemer (2021). I argue that his argument fails, and that my argument establishes an alternative way in which mere existence can be evidence for reincarnation. I then examine whether my result generalizes.

ARTICLE HISTORY Received 11 April 2020; Revised 4 May 2021

KEYWORDS Sleeping Beauty; immortality; *de se*; probability; Bayesianism

1. Introduction

In *Reasons and Persons*, Derek Parfit contemplates what sort of evidence might justify belief in reincarnation [1984: 227]:

There might, for example, have been evidence supporting the belief in reincarnation. One such piece of evidence might be this. A Japanese woman might claim to remember living a life as a Celtic hunter and warrior in the Bronze Age. On the basis of her apparent memories she might make many predictions which could be checked by archaeologists. Thus she might claim to remember having a bronze bracelet, shaped like two fighting dragons. And she might claim that she remembers burying this bracelet beside some particular megalith, just before the battle in which she was killed. Archaeologists might now find just such a bracelet buried in this spot, and their instruments might show that the earth had not here been disturbed for at least 2,000 years. This Japanese woman might make many other such predictions, all of which are verified.

Parfit is quick to add that ‘[w]e do not in fact have the kind of evidence described above’ [ibid.: 228].¹ Further context reveals that he has in mind a particular kind of reincarnation scenario—that of an immaterial Cartesian Ego, surviving past a body’s death, and occupying a different body at a later time. This explains Parfit’s focus on memories as potential evidence for reincarnation: one might think that the surviving Ego retains the memories, such that future incarnations have access to them.²

¹ See Almeder [1992] for an attempt to argue that we do have the requisite evidence. See Dilley [1995] and Hales [2001] for rebuttals.

² See also Ayer [1956: 220].

This essay presents a number of cases—modelled after the well-known Sleeping Beauty puzzle—in which an agent does have surprisingly strong evidence for their own reincarnation. Unlike in Parfit’s scenario, however, the evidence does not involve memories of past lives. Instead, it consists merely in coming to know that one is currently alive.

Reincarnation, as I understand it, is a matter of personal identity. More precisely, I take the following to be a sufficient condition for reincarnation: a subject occupying a body F is a reincarnation of a subject occupying a body G if (1) F exists wholly after the biological death of G and (2) the subject occupying F is *the same person as* the subject occupying G. This condition is also necessary if we disregard (as I will) the possibility of a single subject existing in multiple separate bodies at the same time (as might be the case in some time travel scenarios).

Stipulating the existence of Cartesian Egos is one way to satisfy (1) and (2). But it’s not the only way. For example, reincarnation is in principle consistent with psychological theories of personal identity and with theories that take personal identity as primitive.³

Here’s the plan for this essay. Initially, I review an attempt by Michael Huemer [2021] to establish a similar, albeit stronger, conclusion (section 2). Huemer argues that ‘existence is evidence of immortality’ in a very broad range of cases. I argue that his argument is unsuccessful (section 3).

Section 4 presents my new argument. I develop an ‘Immortal Beauty’ case involving reincarnation that is analogous to Sleeping Beauty. I argue that *if* the popular ‘thirder’ solution to Sleeping Beauty is correct, then Immortal Beauty has strong evidence that she is going to be reincarnated. I’ll show how this establishes an alternative sense in which ‘existence is evidence for immortality.’ Afterwards, in section 5, I’ll examine whether these results generalize, including to our own situation.

2. Huemer’s Argument

In ‘Existence Is Evidence of Immortality’ [2021], Huemer argues that, in a temporally infinite universe (infinite toward both the past and the future), our present existence is a probability-zero event unless we are infinitely reincarnated. Given a Bayesian conception of evidence, we should thus become confident of infinite reincarnation once we learn of our present existence.

At the centre of Huemer’s argument is the following ‘lemma about unrepeatabe events’ [ibid.: 135]:

Let E be any event that may occur at a particular time. Suppose we are given the following:

- (i) E can occur at at most one time in the history of the universe.
- (ii) The history of the universe is infinite.
- (iii) E is initially no more likely to occur at any given time than at any earlier time.

Conditional on just these assumptions, the probability of E occurring now or in the recent past is zero.

The lemma’s proof is as follows. Partition the universe’s time axis into finitely long non-zero time intervals—centuries, say. Then, by (iii), the probability of E’s occurrence

³ To take an example from Huemer [2021], in eternally recurrent universes, theories that identify subjects based on certain qualitative similarities—e.g. personality profile and upbringing—imply reincarnation.

is (weakly) monotonically decreasing over the centuries. By (i), these probabilities are additive (E's occurrence in any given century excludes its occurrence in any other century). But, by (ii), there are countably infinitely many centuries in the past. Thus, the only probability function satisfying (i)–(iii) is everywhere zero.⁴ From this, we conclude that the probability of E's occurring in the present century is zero.

Huemer now asks us to 'substitute *your being incarnated* (that is, beginning a life as a conscious being) for E' [ibid.: 136]. Clause (i) is then satisfied by 'restrictive' views of personal identity, according to which one is incarnated at most once. Clause (ii) is satisfied if we assume that the universe is temporally infinite. And clause (iii) is supposed to hold *a priori*: Huemer claims that 'a priori, if you were to live in some century, it is no more likely that it would be the present century than that it would be any given previous century' [ibid.]. With (i), (ii), and (iii) satisfied, the lemma's conclusion follows: the probability of my being incarnated in the present century is zero. Suppose, further, that my present existence conditional on *infinite* reincarnation has non-zero probability. Then infinite reincarnation is infinitely confirmed over singular incarnation.

How sound is this argument? Huemer's reasoning for the instance of clause (iii) doesn't seem good if we interpret 'I am alive in the present century' as expressing a purely self-locating (or '*de se*') proposition. This is because there's *a priori* good reason to privilege my existence *now* over my existence in the past (say, five centuries ago). Indeed, 'I am alive now' is a strong candidate for an *a priori* knowable truth, while 'I am alive five centuries ago' is not.^{5,6} Huemer's indifference reasoning therefore seems unsound: there *is a priori* reason to 'expect to have been born in the present century [read: *now*] rather than the previous century' [ibid.: 143].⁷

⁴ Given countable additivity, this would entail probability zero that you are ever born. Let us assume here, for Huemer's sake, merely finite additivity. This permits non-zero probability that you are ever born, despite zero probability of being born at any given century. We could also choose to model the situation with a hyper-real-valued probability function. We would still forsake countable additivity, but we may recover regularity. The approach wouldn't change any of the paper's essential results. If our present existence has infinitesimal probability given H, positive real-valued probability given $\neg H$, and the prior probabilities of H and $\neg H$ are positive and real-valued, then one should still become confident in H upon learning that one is presently alive (one's posterior of H should be infinitesimally close to 1). For convenience, I'll stay with real-valued probability functions.

⁵ This is echoed by popular formalizations of *de se* content. Consider a centred-worlds account, on which belief contents are sets of agent-time-world triples. A triple $\langle x, t, w \rangle$ represents that [I am x and t is present and w is actual]. Hence only triples $\langle x, t, w \rangle$ such that x is located at t in w represent epistemic possibilities. (For certainly I am not both x and currently at t in w if x isn't at t in w !) But 'I exist now' plausibly expresses the set of triples $\langle x, t, w \rangle$ such that x is located at t in w . Hence, it exhausts all epistemic possibilities and is thus *a priori* knowable. Meanwhile, 'I exist five centuries ago' expresses the set of triples $\langle x, t, w \rangle$ such that x is located five centuries prior to t in w . This set does not exhaust all epistemic possibilities: it leaves out centres where the agent hasn't existed five centuries prior. (It also includes epistemically impossible centres.)

⁶ Perhaps someone is tempted to object that 'I exist now' isn't *a priori* knowable, on the basis that 'I exist (sometime)' isn't *a priori* knowable. But I have two comments on that. (1) Even if 'I exist (sometime)' is only knowable *a posteriori*, the conditional 'I exist (sometime) \rightarrow I exist now' might still be knowable *a priori*. In that case, Huemer's argument is still blocked. For Huemer provides no reason for thinking that my existence at some time should make me confident of reincarnation. Learning 'I exist (sometime) \rightarrow I exist now' is instead doing all of the work in his argument. (2) My main concern isn't *a priori* knowability *per se*, but only that 'I exist now' (or the conditional) has more *a priori* plausibility than 'I exist five centuries ago' (or its corresponding conditional). And that much strikes me as hard to deny. (In further support of this, note that my learning a *de se* proposition that resists Huemer's indifference reasoning is compatible with my *also* learning something further, to which Huemer's reasoning does apply—namely, a *de re* proposition (see the main text). So, the view that I learn something that is on equal footing, epistemically, to similar propositions about the past can be accommodated.)

But we may take Huemer's argument in a different way. That there is some reading on which 'I am alive in the present century' is *a priori* more likely than 'I am alive five centuries ago' doesn't rule out that there are other readings on which indifference is justified. Indeed, that Huemer repeatedly emphasizes the importance of non-qualitative evidence suggests the following alternative:⁸ read the content of 'I am alive in the present century' as involving its time directly (*de re*).⁹ This makes Huemer's clause (iii) more plausible. 'I am alive in the present century' now expresses a different content at different times (for each content involves its time directly), and it is plausible that my prior credence should distribute uniformly over these contents. For, as Huemer notes, 'different times are intrinsically, qualitatively identical' [Huemer 2021: 144]; so, intuitively, nothing privileges my existence at one particular time (*de re*) over my existence at another time. So, Huemer's argument gets off the ground.¹⁰

To examine whether the argument ultimately succeeds, let's represent it more formally. Let SI (singular incarnation) be the thesis that the universe is temporally infinite in both directions and that you are incarnated (born) exactly once. Thus, SI entails clauses (i) and (ii),¹¹ and it excludes the hypothesis of infinite reincarnation, II. That's the thesis that the universe is temporally infinite in both directions and that you're infinitely reincarnated, with no first or last incarnation.

A *de re* reading of clause (iii), with 'one's being incarnated' substituted for E, gives us this:

⁷ I've been equivocating somewhat, running together 'I exist now' and 'I am alive now.' This potentially invites the following further complaint: 'I exist now' might be *a priori* knowable, but 'I am alive now' is not; for it is not *a priori* knowable that I am a living thing rather than an inanimate object. For the purposes of our discussion, we may grant this complaint. For it won't save Huemer. 'I exist now' and 'I am alive now' are epistemically equivalent, conditional on 'I am alive at any time at which I exist' (something that's plausibly true of agents). And it's hard to see why learning this latter proposition should shift my credences in infinite reincarnation or singular incarnation. (The proposition doesn't play any role in Huemer's own argument.) So, for the purpose of examining our posterior credences in reincarnation, the equivocation is harmless: we may conditionalize our probabilities on 'I am alive at any time at which I exist' without distorting the result. (Note that Huemer also freely switches between 'alive'-talk and 'existence'-talk throughout his paper.)

⁸ One might think, however, that Huemer's disregard of qualitative evidence is premature. Qualitative evidence would still allow for a strong update in favour of immortality, *if* the evidence excluded all but a finite number of centuries. (For 'I am alive in such-and-such a century', where 'such-and-such a century's extension is finite, is infinitely more likely on infinite reincarnation than on singular incarnation.) Huemer doesn't consider this case, probably due to the kinds of cosmological assumptions that he adopts. For example, if the universe is infinitely recurrent, no purely qualitative evidence about the present time picks out a merely finite set of centuries. Note, though, that my response to Huemer in section 3 also responds to a qualitative version of his argument.

⁹ I understand *de re* as follows: on a possible worlds account of content, a set S of worlds 'involves a time t *de re*' iff each world in S has t (or a counterpart of t) as a part. On a Russellian account of content, a content 'involves a time t *de re*' iff t is one of the content's constituents. (Not everyone uses '*de re*' in this way. For example, Lewis [1979: 538–43] suggests a weaker notion, where belief *de re* about an object O is compatible with some worlds' in one's belief content not having O, or any of O's counterparts, as a part.)

¹⁰ One might think that Huemer explicitly rejects a *de re* reading of his argument, given that he also stresses the importance of the 'indexicality of evidence' [2021: 133]. I have two comments on this. (1) Huemer draws the relevant contrast between 'indexical expressions such as "I" and "now" and "purely qualitative terms". This isn't an exhaustive distinction, since it ignores proper names, like 't'. (2) Huemer's examples don't discriminate between indexicals and proper names. For example, 'Flippy The Coin comes up heads ten times in a row' is just as little evidence for many coins' being flipped as '*this specific coin* comes up heads ten times in a row' [ibid.].

¹¹ SI is a slight strengthening of [(i) and (ii)], since clause (i) is compatible with not existing at all. But this is a harmless convenience: given that one is alive at t, SI is epistemically equivalent to [(i) and (ii)]. So, one's posterior credence in SI, after learning that one is alive at t, should equal one's posterior credence in [(i) and (ii)].

- (A) **Indifference.** For any century t and any century t^- before t , the prior probability of one's being incarnated at t is no greater than the prior probability of being incarnated at t^- .

From (A) and Huemer's lemma, it follows that, conditional on SI, the prior probability of being incarnated (and hence alive) at t is uniformly zero:

$$\text{For all } t, P(\text{I am alive at } t | \text{SI}) = 0. \quad (1)$$

Huemer requires a crucial second premise for his argument (which he does not explicate):

- (B) **Probable Existence Given Reincarnation.** For any century t , the probability of one's being incarnated at t conditional on II is non-zero:

$$\text{For all } t, P(\text{I am alive at } t | \text{II}) > 0. \quad (2)$$

Further, Huemer assumes:

- (C) **Non-Zero Priors.** The prior probabilities of SI and II are each non-zero:

$$P(\text{SI}) > 0 \text{ and } P(\text{II}) > 0. \quad (3)$$

Finally, we have (D):

- (D) **Bayesian Conception of Evidence.** If H and H' are hypotheses with $P(H) > 0$ and $P(H') > 0$, then evidence E confirms H over H' iff $P(E|H) > P(E|H')$.¹²

From (A), (B), (C), and (D), we conclude thus:

- (E) **'Existence Is Evidence of Immortality'.** My evidence (for some t) that I am alive at t confirms II over SI.

The argument is valid. Is it also sound? As I mentioned, Huemer holds that Indifference is an *a priori* truth. By his own lights, his reasoning here is 'reminiscent of the controversial Principle of Indifference' [ibid.: 143].¹³ While Huemer doesn't defend the Principle in generality (not here, at least; see his [2018]), he does defend this particular application [2021: 143–4]. Let me suppose that Huemer's defence succeeds. Further, I even grant the remaining premises, (B)–(D).¹⁴

¹² For any H with $P(H) > 0$, the conditional probability is given, as usual, by the ratio formula: $P(E|H) = P(E \& H) / P(H)$.

¹³ One possible formulation of the Principle of Indifference is this: if there is no reason to favour (epistemically) one possibility over another, then the two possibilities have the same probability.

¹⁴ Note that the truth of Premise (B) isn't entirely obvious. Conditional on II, there is an uncountable set U of infinite candidate sets of centuries specifying when I'm alive. Premise (B) requires that the subset of U of sets whose members include t has non-zero measure. This needs to be supported with argument. One way might be to argue that, conditional on infinite reincarnation, there's a non-zero probability that I'm alive in every century. This could be the implication of a multiverse theory, together with an account of personal identity according to which, upon my local death, I immediately continue existence in a counterpart of mine in another part of the multiverse. Other accounts might be feasible, too.

The problem is that the argument's conclusion is too weak for Huemer's purposes. It fails to establish that, upon learning that one is presently alive, one should become confident of immortality. This is because 'I am alive at t' does not exhaust one's total evidence. We have supposed that, in addition to the purely self-locating proposition that I am alive now, I also learn the *de re* proposition that I am alive at t. But, if this is so, I also learn a third self-locating proposition—that t is *now*. As we'll see, this proposition neutralizes the update in favour of reincarnation.

3. Update Neutralized

3.1. An Analogy

The epistemic situation of an agent who learns that she is presently alive is analogous, I claim, to the following situation.

Lottery. You are presented with a 100-ticket lottery. You are told that the lottery contains either only a single winning ticket or ten winning tickets. The game master, who knows which tickets win, is about to reveal to you the number of some winning ticket. You are told that he does this by selecting a ticket at random from among the winning tickets. The game master reveals number 34.

There is an analogue here of Huemer's argument, to the conclusion that, after seeing that ticket 34 is revealed, you should become confident that there are many winning tickets. After all, ticket 34's winning is more likely if there are many winning tickets. But the argument fails, for it ignores how the ticket is selected. Call the hypothesis that exactly one ticket is a winner [that exactly ten tickets are winners] OneWinner [TenWinners]. We have these:

$$P(\#34 \text{ wins} | \text{TenWinners}) = 1/10,$$

$$P(\#34 \text{ wins} | \text{OneWinner}) = 1/100.$$

However, our evidence also includes that the game master reveals 34. By stipulation, he does this by random selection; so we have this:

$$P(\#34 \text{ revealed} | \#34 \text{ wins} \ \& \ \text{TenWinners}) = 1/10,$$

$$P(\#34 \text{ revealed} | \#34 \text{ wins} \ \& \ \text{OneWinner}) = 1.$$

Hence, our total evidence is neutral about the total number of winning tickets (note that '#34 revealed' entails '#34 wins'):

$$P(\#34 \text{ revealed} | \text{OneWinner}) = P(\#34 \text{ revealed} | \text{TenWinners}) = 1/100.$$

The analogy with the incarnation setup arises since, prior to learning when we exist, we have large self-locating uncertainty. Given Huemer-style indifference assumptions, this self-locating uncertainty is analogous to the game master's random selection over the subset of winning tickets.

More concretely, consider a finite incarnation setup:

Finite Universe. Upon being born, you entertain two rival hypotheses: first, the universe has a total lifespan of 100 centuries and you are alive in precisely one of those centuries (SI_t); second, the universe has a total lifespan of 100 centuries and you are alive in precisely ten (II_t).

Huemer's reasoning shows—correctly—that Π_f makes 'I am alive at t ' more likely than SI_f does: *a priori*, I have no reason to favour any collection of ten centuries over any other such collection, as being locations where I am alive. Given Π_f , exactly 1/10 of these collections have me existing at t . Hence,

$$P(\text{I am alive at } t | \Pi_f) = 1/10. \quad (4)$$

Analogously,

$$P(\text{I am alive at } t | SI_f) = 1/100. \quad (5)$$

But now I also learn two *de se* propositions—that I am alive *now*, and that t is *now*. For the purpose of expositing Huemer's argument, let's assume that 'I am alive *now*' doesn't influence my confidence in SI_f or Π_f .¹⁵

Inert De Se Evidence. $P(\text{I am alive now} | SI_f) = P(\text{I am alive now} | \Pi_f)$.

In this case, we can treat P as already conditionalized on 'I am alive *now*'. (Inert *De Se* Evidence does not beg the question against Huemer. As I've previously argued, his indifference reasoning does not get off the ground when applied to 'I am alive *now*', read *de se*.)

Now, if I'm singularly incarnated, then my being alive at century t implies that I'm alive *only* at t . Given, further, that I am alive *now*, the fact that I am alive only at t entails with certainty that t is *now*. Hence,

$$P(t \text{ is now} | \text{I am alive at } t \ \& \ SI_f) = 1. \quad (6)$$

Meanwhile, given Π_f , the information that I am alive at t implies merely that t is one of ten centuries where I am alive. We have (7):

$$P(t \text{ is now} | \text{I am alive at } t \ \& \ \Pi_f) = 1/10. \quad (7)$$

Putting together Eqs. (4), (5), (6), and (7), we derive (8):

$$P(E | SI_f) = P(E | \Pi_f) = 1/100, \quad (8)$$

where E is my total evidence, 't is now & I am alive at t .' Both hypotheses support our evidence equally. An update in favour of multiple reincarnation is unwarranted.¹⁶

This is not to say that there aren't, in principle, situations in which learning when you are alive *does* increase your confidence in reincarnation. This happens when the information is learned without a neutralizing *de se* proposition. Suppose that a time-travelling alien grants you access to a Cosmic Demoscope, a large database listing, for every century, all of the persons alive in that century. Suppose that you pick some century c at random, look up the corresponding entry in the Demoscope, and find that you are alive at c . In this case, you are warranted to update in favour of multiple reincarnation, since you haven't obtained any relevant self-locating

¹⁵ Given that 'I am alive *now*' is *a priori* knowable, there's a question whether a framework of hypothetical priors would even in principle allow non-trivial updates on the proposition. I'll return to this question in note 25.

¹⁶ Here I'm relying on the 'total evidence requirement': rational credences should reflect one's total evidence, and not merely a proper part of it. Huemer himself accepts the requirement, for familiar reasons (cf. his [2021: 7]).

information. Similarly, if in the lottery case you simply pointed to ticket number 34 and demanded that it be revealed to you—irrespective of whether it is a winner—then, upon learning that it is a winner, you should become more confident that there are many winning tickets.

It remains to consider the case of infinite reincarnation. The calculations are essentially the same. The probability that I'm alive at t , conditional on SI, is zero. Since the probability that t is now, conditional on my being alive at t and SI, is well-defined (namely, 1), we then have $P(E|SI) = 0$. The probability that I'm alive at t , conditional on II, is also well-defined (although its value isn't as obvious; see note 14). Further, the probability that t is now, conditional on my being alive at t and II, is 0, because infinite reincarnation induces infinite self-locating uncertainty. Hence:

$$P(E|II) = P(E|SI) = 0. \quad (9)$$

Once the total evidence is in, you shouldn't update in favour of immortality.

3.2. Evidence of Times, De Se and De Re

We've supposed that, upon learning that one is currently alive, one makes epistemic contact to the present time under two distinct guises, 'I am alive now' and 'I am alive at t .' Is that really what our evidence is like?

Note that there's no trouble in formally representing this situation. On a centred-worlds account, 'I am alive now' expresses the set of triples $\langle x, t, w \rangle$ such that x is alive at t in w . 'I am alive at t ' plausibly expresses the set of triples $\langle x, s, w \rangle$ (where ' s ' is a variable ranging over all (actual and non-actual) times) such that x is alive at t in w . These sets are clearly distinct (indeed, the former is a proper subset of the latter). (With a bit of effort, we can also represent these two contents on a Guise Russellian account. On the version defended by Salmon [1989] and Braun [2016], agents admit, besides the binary relation *believes*, a ternary relation, *BEL*. This relation attaches to an agent, a proposition, and a propositional guise. We can represent the *de re* reading of 'I am alive in the present century', for each century, if we help ourselves to a distinct guise for each century. To represent the *de se* reading, we would stipulate, in addition, a single 'now' guise.)¹⁷

Of course, formally distinguishing between the contents isn't enough to show that we can obtain each of them. Consider the *de re* evidence 'I am alive at t .' Is gaining it simply a matter of giving the current time a particular name? This would fit with the idea, implicit throughout Huemer's paper, that gaining evidence for immortality is easy. But it is hard to believe that our evidence changes so drastically by a simple

¹⁷ It's also straightforward to rehash the previous calculations in the two frameworks. Take the centred-worlds account: in (Finite Universe), conditional on 'I am alive at t ' and any specific ten-times reincarnation thesis (specifying who I am, when I live, and what world is actual) consistent with it, my probability distributes uniformly over ten centred worlds, exactly one of which has t for its time index. Hence the conditional probability of ' t is now' is $1/10$. This yields Eq. 7; other equations are obtained similarly. The calculations for the Guise Russellian are exactly analogous, granted that there's a distinct guise for each century, and a single 'now' guise. For any time s , let $g(s)$ denote s 's '*de re*' guise. Conditional on 'I am alive at t ' and any specific ten-times reincarnation thesis consistent with it, my probability distributes uniformly over ten proposition-guise pairs, one of which is (t , t is identical to), $\ulcorner g(t)$ is now \urcorner). This yields Eq. 7; similarly for other equations.

act of christening. Alternatively, perhaps it's enough to look at a clock. But it's hard to see how this affords us the right sort of non-qualitative evidence about the present time, which Huemer claims to require.¹⁸ A third strategy is to define one's evidence in relation to a particular event—for example, the birth of Jesus Christ (as in 'I am alive in 2021 AD').¹⁹ A potential challenge here would be to defend a uniform prior of incarnation—in particular, to deny that a bias toward centuries that are closer to the event in question is rationally permissible.²⁰ However, worries about *de re* evidence are Huemer's problem. If they cannot be answered, *his* argument is in danger: Without a viable *de re* reading of the evidence, it doesn't get off the ground.

But perhaps one thinks that the trouble lies with our *de se* evidence. Perhaps, more concretely, when we learn of our present existence, our evidence of the present time is purely *de re*. The self-locating bit either doesn't arise at all or is reduced to the *de re*. (Our evidence is perhaps already satisfactorily represented by a set of uncentred worlds or by a guiseless Russellian proposition.) This would allow Huemer's *de re* argument to succeed without the threat of neutralizing *de se* evidence.

Against this line, I note two points. First, I take well-known cases from Perry [1979] and Lewis [1979] to show, rather conclusively, that irreducibly self-locating evidence exists. Consider Lewis's [ibid.] 'two gods' case: a god who has exhaustive knowledge of how the world is like—where that might include knowing all *de re* facts (that is, what individuals inhabit the world and how they are like)²¹—can still be uncertain about some things. In particular, he might not know who *he himself* is. This is irreducibly *de se* uncertainty, and evidence that resolves it is irreducibly *de se* evidence. An analogous case arises for evidence about times: surely a god who knew all of the *de re* facts of a world—including what instants of time the world consists of—could still be uncertain about what time it is *now*.

Second, if the Perry-Lewis cases were somehow overcome, there would be substantial pressure to admit that one's existence at times *de re* is *a priori* knowable. For it's plausible that we get at *something* right with our common inclination that there's a relation to times that we instantiate, *a priori*, whenever we exist. One straightforward way to explain this inclination would be to admit (somewhat surprisingly) that, where *t* is the present time, my existence at *t* is *a priori* knowable. But then we've come full circle: given that we only live once, a bias in favour of the present time seems justified. So, Huemer's indifference argument doesn't get off the ground.

Finally, the fact that Huemer's conclusion is highly surprising only increases his burden of proof for a reductive or eliminative account of *de se* evidence. In other words, anyone who wishes to resist Huemer's argument currently has two easy ways out: either accept that, when learning that we're presently alive, we only gain *de se* evidence of times, or accept that we gain both *de re* and *de se* evidence, while noting that the latter neutralizes the former.

¹⁸ See note 8, though. The issue is arguably not about qualitiveness *per se*, but about whether one's evidence is specific enough to pick out a finite number of centuries.

¹⁹ This wouldn't, properly speaking, be *de re* evidence of *time*, but it would fulfil the same role in our calculations.

²⁰ Additionally, some might question whether the birth of Jesus Christ even features *de re* in my thoughts, or whether it is instead picked out by description. See Lewis [1979: 538–43] for a defence of the latter view.

²¹ Cf. Lewis [1979: 522–4].

4. Immortal Beauty

As it turns out, an analysis of *de se* evidence enables a different, more limited, argument for immortality. Where it applies, the argument calls into question—rather strikingly—Inert *De Se* Evidence (section 3). We start by noting a parallel between the immortality case and the well-known Sleeping Beauty puzzle.

Sleeping Beauty. Beauty goes to sleep on Sunday, planning on sleeping for two full days. On Sunday night she is told that, in the night from Sunday to Monday, a team of scientists will flip a coin. If the coin comes up Heads, Beauty is woken on Monday and, shortly after, put back to sleep. If the coin comes up Tails, Beauty is woken on Monday *and* on Tuesday. Again, shortly after waking on Monday, she is put back to sleep; this time, while asleep, her memory of the Monday awakening is erased.

Now consider the following variation:

Immortal Beauty. God creates a universe temporally infinite toward the past and the future. God then flips a coin. If the coin comes up Heads, She incarnates Beauty in precisely one century (singular incarnation). If Tails, She incarnates Beauty in countably infinitely many centuries, with no earliest or latest century (infinite reincarnation). Prior to flipping the coin, God fixes a plan for when to incarnate Beauty, conditional on either outcome of the coin flip. Upon incarnating Beauty, God informs her of the setup, although not about the outcome of the coin flip or about what her plans say.

Thirdering is a popular view about Sleeping Beauty: upon waking on Monday, Beauty should assign credence $1/3$ to the hypothesis that the coin came up Heads (and thus $2/3$ to Tails). Roughly, thirdering think that, since there are more awakenings conditional on Tails, Beauty should become more confident of Tails upon waking—starting with a $1/2$ credence in Tails on Sunday, she should update to $2/3$ on Monday.²²

The new argument is this. If thirdering is correct in Sleeping Beauty, then ‘zeroing’ is correct in Immortal Beauty. After being informed that she is in an Immortal Beauty experiment, Beauty should assign credence 0 to Heads (singular incarnation) and credence 1 to Tails (infinite reincarnation).

I’ll make this argument by going through four intermediate cases, successively closing the gap between the two setups. Consider, first, this case:

Modified Sleeping Beauty. The protocol is as in Sleeping Beauty, except that if the coin comes up Heads then Beauty is woken *either* on Monday or on Tuesday (never on both days). The scientists have received prior instructions specifying on what day Beauty should be woken on Heads (Beauty doesn’t know these instructions).

Next consider this:

Created Sleeping Beauty. After several medical breakthroughs, the scientists are able to create adult humans *in vitro*. Beauty does not yet exist on Sunday. The scientists flip a coin on Sunday night. If Heads, they create Beauty *in vitro* on either Monday or Tuesday. Which day she is created is specified by previously received instructions. Upon creation, Beauty is initially asleep; she is then woken and, some time after, put back to sleep until Wednesday. If Tails, the scientists create Beauty on Monday, again initially asleep. She is woken and, some time after, put back to sleep, and any memory that she has made on Monday is erased. The scientists wake her again on Tuesday, before putting her back to sleep until Wednesday.

Next up:

²² Prominent defenders of thirdering include Elga [2000], Dorr [2002], Horgan [2004], and Titelbaum [2008].

Reincarnated Beauty. Pleased with the team's scientific progress, God decides to grant the scientists some of Her divine powers: the scientists gain the ability to reincarnate the deceased. They rush to test out the newfound powers by setting up a case like Created Sleeping Beauty, with a few (morbid) changes. The Heads protocol remains unchanged. On Tails, however, rather than putting Beauty back to sleep on Monday, they give her a deadly potion. The scientists then create a perfect duplicate of Beauty's body, as it was first created, and, with their newfound powers, reincarnate Beauty into this body on Tuesday.

Next:

Multiply Reincarnated Beauty. The scientists are no longer satisfied with mere two-day experiments. They decide to set up a Reincarnated-Beauty-like case, except that it is run over n days. More specifically, if the coin comes up Heads, Beauty is incarnated and woken precisely once during the next n days (according to previously received instructions). If the coin comes up Tails, Beauty is incarnated, before being given the deadly potion, on each of m of the next n days (again, according to previously received instructions).

The argument from thirding in Sleeping Beauty to zeroing in Immortal Beauty goes as follows.

Immortal Thirders

- (1) If thirding is correct in Sleeping Beauty, thirding is correct in Modified Sleeping Beauty.
- (2) If thirding is correct in Modified Sleeping Beauty, thirding is correct in Created Sleeping Beauty.
- (3) If thirding is correct in Created Sleeping Beauty, thirding is correct in Reincarnated Beauty.
- (4) If thirding is correct in Reincarnated Beauty, $1/(m+1)$ -ing is correct in Multiply Reincarnated Beauty.
- (5) If $1/(m+1)$ -ing is correct in Multiply Reincarnated Beauty, zeroing is correct in Immortal Beauty.

(C) If thirding is correct in Sleeping Beauty, zeroing is correct in Immortal Beauty.

'Zeroing' in Immortal Beauty means assigning posterior probability 0 to Heads, and 1 to Tails. Thus, by Immortal Thirders, a thirder, upon being informed that they are in an Immortal Beauty setup, should become confident of Tails—and hence of her own immortality. I'll briefly defend each of the five premises.

Premise (1). In Modified Sleeping Beauty, Beauty doesn't know on which day she awakes given Heads; otherwise, the situation is identical to the original Sleeping Beauty case. But, intuitively, whether the Heads-awakening occurs on Monday or on Tuesday doesn't seem significant for Beauty's credences upon waking. More specifically, let P be Beauty's credence function upon waking. In Modified Sleeping Beauty, she distributes her credence between two hypotheses:

(H-Mon) The scientists' instructions say to wake her on Monday if the coin comes up Heads.

(H-Tue) The scientists' instructions say to wake her on Tuesday if the coin comes up Heads.

H-Mon and H-Tue partition Beauty's epistemic possibility space. Conditional on H-Mon, the experiment is just the original Sleeping Beauty experiment; so, if thirding is correct, $P(\text{Heads}|\text{H-Mon}) = 1/3$. Similarly, if H-Tue is true, the current experiment is just like the original case but with the Heads-awakening occurring on Tuesday rather than Monday. But it's hard to see what should motivate a different credence of Heads in this case. The time of the single Heads-awakening

relative to the Tails-awakenings plausibly shouldn't influence Beauty's credences upon waking.²³

Premise (2). In Created Sleeping Beauty, Beauty doesn't yet exist on Sunday, and so can't yet form an initial credence about Heads. Some thirder accounts derive Beauty's intra-experiment credence diachronically from her Sunday credences (see, for example, Titelbaum [2008] and Schulz [2010]). Without a Sunday period, these diachronic schemes are silent on what it is rational for Beauty to believe upon waking/creation. But this seems like a defect of these schemes, since surely there are some such rationality constraints. Synchronic schemes, by contrast, don't have this issue. They obtain Beauty's credences at a given time from a constant prior, by conditionalizing the prior on her total evidence at that time. Synchronic schemes aren't silent in the absence of temporally prior probability. So, it's reasonable to think that diachronic schemes need fixing here, and the fact that some are sensitive to the lack of a Sunday period isn't good evidence that Beauty's credence function should mirror this sensitivity. Finally, synchronic updating schemes plausibly predict the same intra-experiment credences in Created Sleeping Beauty as in Modified Sleeping Beauty. The relevant total evidence post-creation seems to be the same as the relevant total evidence post-waking. This supports Premise (2).

Premise (3). In Reincarnated Beauty, Beauty is killed and reincarnated on Tails, rather than put to sleep and woken. This fact might make an emotional or ethical difference, but it shouldn't affect Beauty's credence in Heads.

Premise (4). The argument here is similar to the argument for Premise (1). In Multiply Reincarnated Beauty, also the scientists' Tails-instructions (non-trivially) partition Beauty's epistemic possibility space. Conditionalizing on a cell of this partition, Beauty knows which m days are her Tails-awakenings. The resulting situation is like Reincarnated Beauty, except with m rather than two Tails-awakenings. Again, the relative position of the Heads-awakening plausibly doesn't matter for Beauty's credence in Heads/Tails. This suggests that if thirdering is true in Reincarnated Beauty, then Beauty's credence in Heads should be $1/(m+1)$.

Premise (5). Immortal Beauty is essentially an infinite limit of Multiply Reincarnated Beauty, with days swapped for centuries. The substitution of days with centuries plausibly shouldn't affect Beauty's probabilities. What about the infinite limit? Any one of God's possible plans for (re)incarnation can be viewed as the limit of an infinite sequence of scientists' instructions (where each set of instructions specifies an earlier first, and a later final Tails-incarnation than its predecessor). For any finite m , a set of instructions specifying that Beauty has m Tails-incarnations yields $1/(m+1)$ -ing. And there is no particular reason to expect that Beauty's credences should be discontinuous in the infinite limit. So, Immortal Beauty requires zeroing. This supports Premise (5), and thus conclusion (C).

Here's one way to formally capture the update in favour of Tails in the foregoing cases. Consider Multiply Reincarnated Beauty. If $1/(m+1)$ -ing in Multiply Reincarnated Beauty is true, then the posteriors P' (Heads) and P' (Tails) (Beauty's probabilities

²³ Several concrete accounts of thirdering concur. Among them is the Self-Indication Assumption (SIA) (cp. Bostrom [2010] and Manley [ms.]); Elga's [2000, 143] argument for thirdering based on long-run frequencies; and Horgan's [2004; 2008] synchronic updating scheme for 'preliminary' probabilities.

after being informed of the setup), are related to their respective priors as follows:

$$\frac{P'(\text{Heads})}{P(\text{Heads})} = \frac{1}{m} \cdot \frac{P'(\text{Tails})}{P(\text{Tails})}. \quad (10)$$

Suppose that, in the original Sleeping Beauty case, Beauty gets her posteriors by updating on new evidence.²⁴ The same will then be plausible in Multiply Reincarnated Beauty. Let I^* be the new evidence. By Eq. 10, we have (11):

$$P(I^*|\text{Heads}) = 1/m \cdot P(I^*|\text{Tails}). \quad (11)$$

If, furthermore, $P(I^*) > 0$, we can express the conclusion that Tails is infinitely confirmed over Heads as the limit of Eq. 11 for infinitely many centuries:

$$P(I^*|\text{Heads}) = 0 \cdot P(I^*|\text{Tails}). \quad (12)$$

This makes precise the sense in which Beauty's credence in Immortal Beauty is the 'infinite limit' of her credence in Multiply Reincarnated Beauty.

Eq. 11 and Eq. 12 also show how 'existence is evidence of immortality.' Let H_1 be 'Heads and it's Monday', T_1 be 'Tails and it's Monday', and T_2 be 'Tails and it's Tuesday.' In the original Sleeping Beauty case, Beauty's new evidence is $H_1 \vee T_1 \vee T_2$. The analogue of this evidence for Multiply Reincarnated Beauty, I^* , is a disjunction of disjunctions: each subordinate disjunction is associated with a specific set of scientists' instructions; it is of the form $H_i \vee T_{j_1} \vee T_{j_2} \vee \dots \vee T_{j_m}$, where the i and j_k specify the days of waking conditional on Heads and Tails, respectively. Crucially, once Beauty is informed of the setup, she can infer I^* from the simple fact that she is alive now. Indeed, given the background knowledge of the setup, the two pieces of information are equivalent. Hence, from Eq. 11,

$$P(\text{I am alive now}|\text{Heads}) = 1/m \cdot P(\text{I am alive now}|\text{Tails}). \quad (13)$$

On this view, then, thirders deny Inert *De Se* Evidence (section 3).²⁵ Thus, Immortal Thriders captures a sense in which 'existence is evidence of immortality.' The evidence is a purely *de se* proposition: I am alive now. (Note that we've achieved this result without assuming that my-existence-now is somehow on equal footing, *a priori*, with my-existence-five-centuries-ago. The argument, and thirthing in particular, are compatible with assuming that the former is *a priori* knowable, and the latter isn't.)²⁶

²⁴ Many thirders agree with this: see, e.g., Horgan [2008] and Manley [ms.]. One account that disagrees is Elga's [2000].

²⁵ We've noted that 'I am alive now' is plausibly *a priori* knowable. Hence, Eq. 13 requires that some hypothetical priors don't assign probability 1 to all *a priori* knowable truths. This might seem surprising, since rational credences arguably don't have this feature (cf. Gallow [ms.: 24]). But hypothetical priors needn't perfectly align with rational credence functions. The purpose of hypothetical priors is to encode epistemic norms in the absence of any evidence at all, and this includes *a priori* knowable evidence. Prior to conditionalizing on such evidence, hypothetical priors may thus differ from rational credences. An alternative way to read Eq. 13 would be as a direct restriction on rational initial credences. The restriction would say, roughly, that, as m grows large, any rational initial credence function should favour Tails over Heads. This would permit equating hypothetical priors with credences. If we take this route, I'll need to adjust my conclusion slightly: existence isn't evidence for immortality; rather, prior belief in immortality is rationally required by an *additional* epistemic norm.

²⁶ Since the argument is a conditional argument from thirthing, people who are certain of halving won't be moved much. But I take it that most people, if not thirders, at least lend thirthing considerable credence.

5. Generalizing the Result

The results from Immortal Beauty generalize, to an extent. The presence of a coin-flipping deity simplifies the analysis but isn't required for strong pro-immortality conclusions. For example, suppose that you distribute all of your credence between the following two hypotheses:

Lone Mortal. The universe is temporally infinite in both directions. You are incarnated exactly once, and live in exactly one century. No other person exists.

Lone Immortal. The universe is temporally infinite in both directions. You are reincarnated infinitely often, and each century contains at most one of your incarnations, and each incarnation lives in exactly one century. No incarnation has any recollection of previous lives or leaves any trace of their existence after death. No other person exists.

The remaining details may be filled in such that, as in Immortal Beauty, the qualitative evidence you receive upon being born doesn't tell in favour of either hypothesis. Your epistemic situation is then relevantly like Beauty's in Immortal Beauty.

A salient difference is that there's no coin flip to fix the priors. But a realistic choice of priors assigns both Lone Mortal and Lone Immortal some non-zero probability, for reasons given in section 2 (recall, specifically, note 14). The situation is then analogous to an Immortal Beauty setup with a biased coin, with the bias exactly equalling the prior probabilities of Lone Mortal and Lone Immortal. Tails would still be infinitely confirmed over Heads upon registering your present existence. Thus, an agent who decides between Lone Mortal and Lone Immortal should become confident of her own immortality, upon registering her present existence.

We can further complicate this situation by allowing the agent to entertain additional reincarnation hypotheses, besides Lone Immortal and Lone Mortal. Indeed, a temporally infinite universe allows for infinitely many incarnation hypotheses: for every natural number n , there is the hypothesis that you are n times incarnated. Lone Immortal will be infinitely confirmed over any such finite reincarnation hypothesis. Given merely finite additivity, it's still possible that Lone Immortal's posterior probability isn't close to 1. But we would still get the result that, for any natural number n , you should have posterior 1 that you are incarnated at least n times. So, you should become confident that you are reincarnated *arbitrarily* often. And if the probabilities *are* countably additive in this case, we recover the stronger result that Lone Immortal is infinitely confirmed over the infinite disjunction of all finite reincarnation hypotheses. In this case, you should become confident of immortality proper.

Next, we may consider cases where the universe is more populated. Suppose that you are uncertain between two theories: either the universe is temporally infinite, but only the current (say) trillion years are populated, containing (say) one quadrillion people; or the universe is temporally infinite, and infinitely many disjoint trillion-year periods are populated, with the same exact quadrillion people in each. Let the rest of the details be as before. Contemplating these two hypotheses is plausibly analogous to contemplating whether the coin came up Heads or Tails in a 'many-subjects' Sleeping Beauty experiment. This is a Sleeping Beauty experiment where, given Heads, a group of people is woken, together, only once, while on Tails they are woken multiple times,

Upon finding themselves in an Immortal Beauty setup, they should still update in favour of immortality, in proportion to their credence in thirring.

with memory erasures in between. (The group of people in this experiment is analogous to the quadrillion people.) It's plausible that, given thirdering, each subject should 'third' also in this many-subjects variant: that is, the subject should assign credence $1/3$ to Heads upon waking. (One way to argue for this is to note that, for any person in the group, *conditional on oneself being that person*, one's epistemic situation is analogous to Beauty's in a conventional Sleeping Beauty case. Moreover, since Heads and Tails both involve the exact same persons, which person you are isn't evidence in favour of Heads or Tails.) Therefore, by an analogue of Immortal Thirderers, each of the quadrillion people should update in favour of reincarnation upon registering their present existence.²⁷

5.1 Limits

There are also limits to how far Immortal Beauty generalizes.²⁸ Suppose that you divide your credence between Lone Immortal and this:

Many Mortals. The universe is temporally infinite in both directions and contains at most one person per century. There are infinitely many persons, each living in exactly one century, and nobody is reincarnated after death.

Assume, further, that the distribution of bodies over the centuries is probabilistically independent of each hypothesis. Intuitively, the two hypotheses then merely differ on the distribution of personal identities over those bodies. Many Mortals specifies that all of the past and future bodies belong to distinct people; Lone Immortal specifies that all past and future bodies belong to *you*.

This epistemic situation is arguably disanalogous to Sleeping Beauty. This is because thirdering typically (although often tacitly) presupposes that there are no additional observer moments anywhere in the universe with the same evidential state as Beauty intra-experiment. For Beauty's intra-experiment credences are typically sensitive to the ratio of the number of observer moments equivalent to Beauty's given Heads to the number of such observer moments given Tails. In the original Sleeping Beauty case, that ratio is 1:2, inducing thirdering. The existence, on Heads, of additional subjects with the same evidential state as Beauty's, threatens this rationale.²⁹

²⁷ More concretely, the case for the analogy would go like this: start with a 'Modified Many-Subjects SB' (leaving unspecified the dates of the group's Heads awakening); then 'Created Many-Subjects SB' (the group is created *de novo* intra-experiment, rather than already being alive on Sunday); 'Reincarnated Many-Subjects SB' (the group dies and is reincarnated, on Tails); and 'Multiply Reincarnated Many-Subjects SB' (the group is multiply reincarnated given Tails); and 'Immortal Many-Subjects SB' (the universe is temporally infinite in both directions, and the group is infinitely reincarnated given Tails).

²⁸ Here is another consideration. What if we replace the infinite universe in Immortal Beauty with a *finite* universe? Beauty doesn't outlive the universe, and so literal immortality is off the table. But she should still update in favour of *finite* multiple reincarnation. The update is weaker than in the infinite case: from Eq. 10 it follows that, if you start with a prior probability of reincarnation lower than $1/m$, then your posterior probability still favours singular incarnation (Heads). For example, suppose that your priors are $P(\text{Tails}) = 1/n$ and $P(\text{Heads}) = 1 - 1/n$ for some natural number $n > 1$. Then your posterior probabilities are related as $P'(\text{Heads}) = \frac{n}{n-1} \cdot \frac{n}{m} \cdot P'(\text{Tails})$. For large n , this is approximately $P'(\text{Heads}) \approx \frac{n}{m} \cdot P'(\text{Tails})$. So, for $n > m$, your posterior still favours Heads. In the finite case, the update isn't strong enough to overcome priors heavily biased against reincarnation. (Thanks to an anonymous referee for pressing me to clarify this.)

²⁹ For example, in the SB-Twin case below, synchronic accounts of thirdering, like Horgan's [2008] or SIA, and arguably also diachronic accounts like Titelbaum's [2008], will treat Heads and Tails symmetrically, and as a result recommend halving.

But a satisfactory analogy to someone who decides between Lone Immortal and Many Mortals requires additional subjects on the Heads outcome in Sleeping Beauty. Consider:

SB-Twin. Everything is as in Sleeping Beauty, except that, conditional on Heads, an additional subject, Twin, is woken on Tuesday, in a separate but qualitatively indistinguishable room. Twin's memories have been manipulated to be indistinguishable from Beauty's memories when she wakes.

SB-Twin can be analogized, via familiar steps, to an 'Immortal SB-Twin' case: this is a case where Heads and Tails agree on the (infinite) distribution of bodies across the centuries, but where Heads says that each body belongs to a different person, and Tails says that all bodies belong to the same person.³⁰ The analogy invites an 'Immortal Thirder's' argument, to the effect that, if thirdering in SB-Twin is correct, then zeroing is correct in Immortal SB-Twin. Beauty's epistemic situation in Immortal SB-Twin, in turn, is analogous to someone who decides between Lone Immortal and Many Mortals. Overall, then, if thirdering in SB-Twin is right, that person should become confident of Lone Immortal.

What, if anything, can be said in favour of thirdering in SB-Twin? It's clear that the position assigns a central role to personal identity in determining what Beauty should believe. After all, if it was Beauty rather than Twin who awoke on Tuesday given Heads (with Beauty's Monday-morning mental state), the coin flips would have symmetrical outcomes, and halving would be irresistible.

One path to leveraging personal identity in favour of thirdering in SB-Twin is via certain (externalist) 'conservative' diachronic norms. Being informed of the setup on Sunday, Beauty knows that *she* is the subject who undergoes a standard Sleeping Beauty protocol. Upon waking on Monday, Beauty initially maintains this belief. Conservative norms may prescribe that Beauty keep this belief even after she updates on being awake.³¹ If any such view *also* prescribes thirdering in the traditional Sleeping Beauty case, it will support thirdering in SB-Twin.

However, conservative norms are unlikely to yield thirdering in Sleeping Beauty. For they would probably want to enjoin Beauty to *also* maintain her Sunday credence in Heads—namely, 1/2.³² A significant further challenge for the conservative on the way to thirdering in 'Immortal SB-Twin' will be to defend thirdering in 'Created SB-Twin'. This scenario eliminates the Sunday period, whose existence may be crucial for the conservative rationale for thirdering in this case.

Are there other ways for an externalist to achieve thirdering in SB-Twin? I, at least, do not see any. So, I'm inclined to think that whether existence is evidence for immortality depends on what thesis is being contrasted with reincarnation. When choosing between Lone Mortal and Lone Immortal, thirderers should indeed become confident

³⁰ More concretely, the analogy from SB-Twin would proceed like this: start with a 'Modified SB-Twin' (leaving the dates of Beauty's and Twin's Heads awakenings unspecified); 'Created SB-Twin' (Beauty and Twin are both created intra-experiment); 'Reincarnated SB-Twin' (Beauty is killed and reincarnated); 'Multiply Reincarnated SB-Twin' (on Heads, many distinct twins of Beauty's are created, each on separate days; the total number of twins equals the number of Beauty's Tails-incarnations minus 1); 'Immortal SB-Twin' (as described).

³¹ There are obvious parallels here to Elga's [2004] Dr. Evil case, where conservative views correspondingly say that Dr. Evil should maintain his belief that *he* is Dr. Evil even after the creation of his clone.

³² See Schwarz [ms.] for more on the distinction between 'conservative' and 'evidentialist' credence dynamics. Schwarz defends a particular conservative diachronic norm, which recommends halving in Sleeping Beauty.

of their own immortality. But when Lone Immortal is contrasted with Many Mortals, thirders should resist.

5.2 Do We Have Evidence for Reincarnation?

How does this bear on our own epistemic situation? (I'll focus here, as I've done throughout, on the portion of our epistemic space where the universe is temporally infinite in both directions.) Our beliefs about the overall number and distribution of conscious lives are, plausibly, approximately independent of whether reincarnation is true. For one, we draw our beliefs about how many *physical bodies* there are largely from our theories of physics (for instance, in the case of a temporally infinite universe, whether the universe is cyclic somehow and, if so, how a cycle evolves). Plausibly, our credences in such theories are largely independent of our theories of personal identity. For another, whether these bodies are *conscious* is largely independent of these theories, too. Plausible contenders for theories of consciousness are sensitive only to whether the bodies are functionally and/or biologically similar to our bodies. This is, plausibly, largely probabilistically independent of *who* occupies the bodies.

As a result, our epistemic situation is more closely analogous to someone who chooses between Lone Immortal and Many Mortals. Recall that, given this choice, the distribution of conscious bodies is independent of reincarnation. More specifically, our situation is analogous to choosing between 'many subjects' versions of these hypotheses. These versions accommodate finitely many agents in each century, who are either each infinitely reincarnated or live only once. In turn, this choice situation can be analogized to a 'many-subjects' version of SB-Twin (where, each day, many people are woken—the same people on Tails, distinct people each day on Heads). It's plausible that the thirder's resistance to thirring in SB-Twin carries over to this many-subjects version.³³ As a result, even if thirring in conventional Sleeping Beauty is true, we shouldn't update in favour of immortality.

Now, it's implausible that our beliefs about the number of conscious lives are *perfectly* independent of whether reincarnation is true. For example, one might think (admittedly idiosyncratically) that there are only finitely many possible Cartesian Egos. In such a situation, singular incarnation places an upper limit on the number of conscious agents that could exist: there are not enough Egos to go around! (Maybe the other lives are zombies, or maybe they're never born.) In this case, it would be rational to update in favour of immortality, even for someone in our situation. But the scenario is arguably far-fetched, and so won't receive much epistemic weight. I suspect that the same is true for other scenarios on which singular incarnation limits the total number (or duration) of observers.

So, I'm inclined to think that, given our actual epistemic situation, it's unjustified to infer reincarnation from the fact that we are alive now.³⁴ For some possible

³³ For, given any agent *M* living on Monday and any agent *T* who lives on Tuesday iff Heads comes up, conditional on [being *M* or being *T*], the epistemic situation is analogous to Beauty's in SB-Twin. Further, the disjunction [being *M* or being *T*] doesn't favour Heads over Tails.

³⁴ To put what's been said more formally, let *Perm* be the proposition that a 'permissive' theory of personal identity is true (i.e. a theory that allows for reincarnation), and let *Restr* (short for 'restrictive theory') be the negation of *Perm*. Suppose that there are *n* mutually exclusive and jointly exhaustive candidate theories—*C*₁, ..., *C*_{*n*}—

agents, no doubt, existence is evidence for immortality. For us, though, it probably shouldn't be.^{35,36}

Disclosure Statement

No potential conflict of interest was reported by the author.

ORCID

Jens Jäger  <http://orcid.org/0000-0003-4615-4593>

References

- Almeder, Robert F. 1992. *Death and Personal Survival: The Evidence for Life After Death*, Lanham, MD: Rowman & Littlefield.
- Ayer, A.J. 1956. *The Problem of Knowledge*, London: Macmillan.
- Bostrom, Nick 2010. *Anthropic Bias: Observation Selection Effects in Science and Philosophy*, New York: Routledge.
- Braun, David 2016. The Objects of Belief and Credence, *Mind* 125/498: 469–97.
- Dilley, Frank B 1995. Review of Almeder [1992], *International Journal of Philosophy of Religion* 37/3: 185–8.
- Dorr, Cian 2002. Sleeping Beauty: In Defence of Elga, *Analysis* 62/4: 292–6.
- Elga, Adam 2000. Self-Locating Belief and the Sleeping Beauty Problem, *Analysis* 60/266: 143–7.
- Elga, Adam 2004. Defeating Dr. Evil with Self-Locating Belief, *Philosophy and Phenomenological Research* 69/2: 383–96.
- Gallow, J. Dmitri manuscript. Two-Dimensional Deference.
- Hales, Steven D. 2001. Evidence and the Afterlife, *Philosophia* 28/1–4: 335–46.
- Horgan, Terry 2004. Sleeping Beauty Awakened: New Odds at the Dawn of the New Day, *Analysis* 64/1: 10–21.
- Horgan, Terry 2008. Synchronic Bayesian Updating and the Sleeping Beauty Problem: Reply to Pust, *Synthese* 160/2: 155–9.
- Huemer, Michael 2018. *Paradox Lost: Logical Solutions to Ten Puzzles of Philosophy*, Cham: Palgrave Macmillan.
- Huemer, Michael 2021. Existence Is Evidence of Immortality, *Noûs* 55/1: 128–51.

about the number and distribution of conscious bodies in the universe. Then $P(C_1 \vee C_2 \vee \dots \vee C_n) = 1$. If we assume that these theories are entirely independent of issues of personal identity, we get $P(C_i | Perm) = P(C_i | Restr)$, for all $i = 1, \dots, n$. Further, as we've learned, once the distribution of conscious bodies is fixed, thirders needn't think that existence is evidence for immortality. That is, $P(I \text{ am alive now} | C_i \vee Perm) = P(I \text{ am alive now} | C_i \vee Restr)$, for all $i = 1, \dots, n$. Putting this together, we have $P(I \text{ am alive now} | Perm) = P(I \text{ am alive now} | Restr)$. Given the independence of each C_i from issues of personal identity, learning that you're currently alive (*de se*) doesn't favour permissive theories over restrictive theories.

³⁵ A referee asks whether there could be a spatial analogue of the Immortal Beauty argument, perhaps concluding that there would be evidence that we'll be simultaneously reincarnated in multiple bodies. Such an argument might start with a variant of Wolfgang Schwarz's [ms.] 'Broken Duplication Machine': if the coin lands heads, Beauty is woken on Monday; if the coin lands tails, Beauty fissions into two copies, each of whom is woken on Monday (in separate rooms). The argument would conclude with 'Universal Beauty': Beauty is either incarnated once or incarnated in many bodies simultaneously. 'Thirding' in the variant of Schwarz's case leads to assigning probability 1 to simultaneous incarnation in Universal Beauty. Perhaps there is a case here, but I note that (i) it's arguably impossible for Beauty to be identical to its fission products (who are presumably not identical to each other); (ii) without personal identity, the case for thirdering in the variant Schwarz's case is arguably weaker than in Sleeping Beauty; and (iii) as with Immortal Beauty, our own situation is not analogous to that of Universal Beauty.

³⁶ Thanks most of all to David Chalmers, for helpful comments and advice at every stage of the paper's development. Many thanks also to Cian Dorr, for extensive written comments and discussion. I am also indebted to Mike Huemer, Marko Malink, Soren Schlassa, and Patrick Wu for helpful discussion. Finally, thanks to two anonymous referees of this journal.

- Lewis, David 1979. Attitudes *De Dicto* and *De Se*, *The Philosophical Review* 88/4: 513–43.
- Manley, David [manuscript](#). On Being a Random Sample.
- Parfit, Derek 1984. *Reasons and Persons*, Oxford: Clarendon Press.
- Perry, John 1979. The Problem of the Essential Indexical, *Noûs* 13/1: 3–21.
- Salmon, Nathan 1989. Illogical Belief, *Philosophical Perspectives* 3: 243–85.
- Schulz, Moritz 2010. The Dynamics of Indexical Belief, *Erkenntnis* 72/3: 337–51.
- Schwarz, Wolfgang [manuscript](#). Evidentialism and Conservatism in Bayesian Epistemology.
- Titelbaum, Michael 2008. The Relevance of Self-Locating Beliefs, *The Philosophical Review* 117/4: 555–606.