

## EXISTENCE IS NOT EVIDENCE FOR IMMORTALITY

ABSTRACT. Michael Huemer argues, on statistical grounds, that “existence is evidence for immortality”. On reasoning derived from the anthropic principle, however, mere existence cannot be evidence against any non-indexical, “eternal” hypothesis that predicts observers. This note attempts to advertise the much-flouted anthropic principle’s virtues and workings in a new way, namely by calling attention to the fact that it is the primary intension of one’s indexically-described evidence that best characterizes one’s epistemic position.

### Theories of persons and the Bayesian conception of evidence

Michael Huemer (2019) notes that if the universe is infinite, then under certain natural hypotheses every configuration of “stuff” (physical matter, energy, psychons, ectoplasm, whatever) that occurs in a bounded region of it will recur, up to an arbitrarily fine degree of approximation, infinitely many times.<sup>1</sup> Since you occur, presumably, in a finite region of the universe, the configuration of stuff that constitutes you would on these hypotheses recur, to arbitrarily fine approximation, infinitely many times. Call such recurrences your “Poincaré clones”. If your Poincaré clones are you, you are immortal. So: *are your Poincaré clones you?*

Huemer thinks that this is a question between two substantive theories (and not, I take it, just a distinction between two ways of talking). Call a theory of persons on which your Poincaré clones are you *Permissive*. Call a theory of persons on which your Poincaré clones are not you *Restrictive*. Huemer thinks that statistical reasoning can decide which theory is correct. I shall argue, however, that Huemer has neglected something important in his calculations—the anthropic principle.

Huemer begins his discussion by introducing the following seeming truism:

*Bayesian Conception of Evidence:*  $E$  supports  $H$  if  $P(E|H) > P(E|\sim H)$ .

He illustrates with an example running (very loosely) as follows. Imagine two competing hypotheses about an experiment in which participants toss fair coins.  $H$  is the hypothesis that there are 1000 participants and 1000 coins, with each participant tossing their coin ten times.  $\sim H$  is the hypothesis that (for some suitable numbering) only participants 1-10 toss their coins. After the experiment you know that you will be handed a coin, chosen uniformly at random from those coins that were tossed, and told how that coin landed during the experiment. Suppose that you are told that the coin landed heads ten consecutive times.

---

<sup>1</sup>Huemer alludes to the Poincaré recurrence theorem here, although he does neglect to mention the crucial hypothesis that one is dealing with a volume (measure) preserving transformation. To tweak an example he discusses where the theorem fails due to unbounded phase space, consider a system where two particles move apart monotonically, with distance approaching a finite limit  $L$ . This system is bounded in the relevant sense, but states don’t recur.

Huemer considers two candidate expressions that might be taken to capture your total relevant evidence proposition in an utterance. The first is indexical:  $E =$  “this specific coin lands heads ten consecutive times”. We believe that Huemer would say here that  $P(E|H) = P(E|\sim H) = 2^{-10}$ , so that there is no confirmation of  $H$  by  $E$ . The second is  $E =$  “some coin in the experiment lands heads ten consecutive times”. (I’ll mostly ignore the second, less tempting expression.)

Now we return to the main argument. Let  $H$  be the expression “a Permissive theory of persons is true”, and let  $\sim H$  be the expression “a Restrictive theory of persons is true”. By analogy with the coins case your relevant evidence is captured by the expression  $E =$  “this specific person exists in this century”.<sup>2</sup> (By “this specific person” you ostend yourself.) Since according to a Restrictive theory you get to exist at most once, then assuming the universe is infinite, Huemer claims that  $P(E|\sim H) = 0$ . But, he goes on to say, Poincaré clones of you exist with positive probability in any given century. So on a Permissive theory *you* exist with positive probability in any given century. This implies that  $P(E|H) > 0$ , so by the Bayesian conception of evidence  $E$  confirms  $H$ . In fact, it refutes  $\sim H$  a.s.

Proponents of anthropic reasoning will see a problem with this argument. For example, it seems that Huemer translates indexical expressions such as “this specific person exists in this century” into propositions by their so-called *secondary intension*. As is well known, however, it is the expression’s primary intension that most closely tracks the epistemic position of the speaker.<sup>3</sup> To wit, if I compose a specific enough indexical expression, e.g. something like (but far more specific than) “I am sitting at my desk, there are two computer screens, that I own, in front of me, several markers that I also own, a broken pencil, a deck of cards to my left, my cell phone lies to my right (it just buzzed at me), I can see my garage out my window, as well as my car, a black 2005 Sequoia, my dogs (Ramsey and Hobbes) are on the couch...” etc. then the primary intension of this expression should pick out (to some degree...employing the primary intension does not come to a “description” theory of reference) something like my “epistemic alternatives”...the set of contexts that I might occupy, for all I know. In particular, it should pick out a set of contexts containing those of my Poincaré clones.<sup>4</sup>

Apart from apparently clear-cut intuitions that it is the primary intension that captures epistemic position (one doesn’t, after all, always even know that the

<sup>2</sup>Or *now*, which is somewhat stronger, but makes little difference to what we have to say.

<sup>3</sup>The terminology is from Chalmers (1996, p. 57): “The primary intension of a concept is a function from worlds to extensions reflecting the way that actual world reference is fixed. In a given world, it picks out what the referent of the concept would be if that world turned out to be actual.” And (p. 64), “Composing the primary intensions of the terms involved yields a primary proposition, which holds in precisely those contexts of utterance in which the statement would turn out to express a truth.” The secondary intension, meanwhile, treats referring terms as rigid: “reference is fixed in the actual world, depending on how the world turns out...”.

<sup>4</sup>Note that “a context of utterance...is an ordered pair consisting of a world and a center representing the viewpoint within that world of an agent” (Chalmers 1996) using the expression in question. This implies, in particular, that the primary intension of an expression picks out not a classical proposition, but a so-called “centered proposition”, i.e. a set of such pairs.

propositions picked out by the secondary intensions of our asserted indexicals are true), trying to adapt the secondary intension to this end breeds monsters. For example, if it's right to identify the secondary intension of "this specific person exists in this century" with one's evidence in the current case, then it isn't clear what could prevent one from using it to identify the secondary intension of "this specific coin lands heads ten consecutive times" with one's evidence in the coins case. But that leads to disaster, because if  $H$  is true and the coin is numbered above 10, one would thereby get  $P(E|\sim H) = 0$ , since *that specific coin*, on its secondary intension, wouldn't be tossed at all conditional on  $\sim H$ .

No such trouble arises for the primary intension, according to which  $P(E|H) = P(E|\sim H) = 2^{-10}$ , consistent with the intuition that neither  $H$  nor  $\sim H$  should be confirmed. Note that on this interpretation the probability of *this specific coin is tossed by that specific participant* is 1; you knew prior to being shown the coin/participant pair that you would affirm that indexical.<sup>5</sup>

Back in the main argument, we can assume that your relevant evidence is captured by the indexical expression "this specific person exists in this specific century". The secondary intension of this expression has positive probability (again by analogy with the Poincaré recurrence theorem, which says that in a finite measure preserving system almost every orbit originating in a positive measure set returns to that set with positive asymptotic density) conditional on  $H$  (the Permissive theory) but zero (or infinitesimal) probability conditional on  $\neg H$  (the Restrictive theory). So Huemer's argument would succeed if it were the secondary intension that tracked epistemic position. Again, though, it is the primary intension that does this...and if  $E$  is the primary intension of the expression "this specific person exists in this specific century" then  $E$  cannot confirm anything, as  $P(E|H) = P(E|\sim H) = 1$ . For the primary intension of an utterance is true in those contexts where it would express a truth...all contexts, in this case.

### The firing squad thought experiment

But consider John Leslie's firing squad (cf. Leslie 1996, p. 141).<sup>6</sup> If fifty marksmen have just taken aim at you and missed then "I exist now" seems to be, in your context, evidence for the proposition that someone has conspired to keep you alive. So is it just wrong to say that "I exist now" cannot be evidence *for anything*?

The primary intension of "I exist now" does not in fact support the conspiracy theory—though something like "I was recently convicted, set before a firing squad and fired upon by fifty marksmen, and yet I am still alive" does. Most observers at contexts where the latter is true are at contexts where "someone conspired to keep me alive" is also true. It is not however the case that most observers at contexts where "I exist now" is true are at contexts where "someone conspired

<sup>5</sup>One could decompose the primary intension formally as  $E_p = \bigvee_j (j \text{ is observed and lands all heads})$ , where the disjunction runs over all tossed coins  $j$ . Note in particular that  $E_p \neq \bigvee_j (j \text{ lands all heads}) = \text{some coin lands all heads}$ .

<sup>6</sup>Thanks to Michael Huemer for calling my attention to this example, and for other helpful remarks on an earlier draft of this paper.

to keep me alive” is also true. So “I exist now” is not evidence for “someone conspired to keep me alive”. If you still think that it is, imagine that you fainted just before the bullets were fired and woke up an hour later with amnesia. You could then truthfully assert “I exist now”, but you wouldn’t have any evidence for “someone conspired to keep me alive”. Such evidence requires greater specificity.

Huemer responds to the objection that expressions such as “I exist now” cannot constitute evidence in Section 5.7 of his paper along similar lines:

Some philosophers hold that one’s own existence cannot serve as Bayesian evidence for any hypothesis. The reason is that the epistemic probability of one’s own existence, in any epistemic position one might be in, is 100%. By a general theorem of probability, for any  $e, h$ , if  $P(e) = 1$ , then  $P(h|e) = P(h)$ . But on a Bayesian conception of evidential support,  $e$  supports  $h$  only if  $P(h|e) > P(h)$ . Thus, it appears that when  $P(e) = 1$ ,  $e$  cannot support  $h$ . (...) This claim has perfectly mundane counterexamples. For instance, the fact that you exist is obviously evidence, for you, that your parents were not infertile. (...) Perhaps the critic of Bayesianism would argue that the preceding example is not genuinely one of Bayesian reasoning, that there is some account of the reasoning that does not involve application of Bayes’ Theorem or anything like it. Such a move, however, would obviate the objection to the argument for Reincarnation. For whatever account is to be given of how your existence supports that your parents (were not infertile), that account may be deployed to explain how your existence supports Reincarnation. (It does not matter whether the account qualifies as truly Bayesian.) Both pieces of reasoning involve the claim that if a certain hypothesis were true, one would be less likely to exist than otherwise; hence, one’s existence supports the rejection of that hypothesis. The objection raised to the reincarnation inference, that you were never uncertain of your existence, applies just as well to the infertility inference. There is no reason to think that whatever would block the objection in the infertility case would somehow allow the objection to succeed in the reincarnation case.

Huemer’s “critic” might be correct to think that the infertility example is not genuinely one of “Bayesian” reasoning.<sup>7</sup> Applications of Bayes’ theorem require prior and posterior probability functions. But “my parents were infertile” is indexical, and the evaluation of an indexical’s primary intension requires “situatedness”. I may be able, *ex post facto*, to construct some prior probability in the expression’s secondary intension (based on facts about the frequency of infertility in the population of adult persons<sup>8</sup>), but in order to have *ex post facto* access to a plausible prior in the expression’s primary intension, I would need to entertain, as if from

<sup>7</sup>Though I wouldn’t characterize the critic as a “critic of Bayesianism” for thinking this.

<sup>8</sup>Since Huemer views the current objection as a species of the “confirmation by old evidence” problem, I suspect that this may be precisely what he takes himself to be doing.

a vantage of non-existence, some distribution or other over possible contexts of utterance. That is, I would have to engage in self-selection (anthropic) reasoning.

### Carter's formulation of the anthropic principle

The *locus classicus* on modern anthropic reasoning is Carter (1983). His formulation of what Huemer calls the Bayesian Conception of Evidence is as follows:

$$\frac{p_E(A)}{p_E(B)} = \frac{p_S(E|A) p_S(A)}{p_S(E|B) p_S(B)}.$$

We will interpret this equation as follows.  $E$  is an indexical expression taken to capture your current evidence, and  $p_E(\cdot) = p_S(\cdot|E)$ . Here  $p_S$  is a “Selected” or “Subjective” probability function that Carter is careful to distinguish from “Original” or “Objective” probability  $p_O$ . (It may be helpful to think of  $p_O$  as objective chance.) Carter writes: “ $S$  denotes the totality of all the selection conditions that are implied by the hypothesis of application of the theory to a concrete experimental or observational situation, but which are not necessarily included in the abstract theory”. Here a “theory” (or “abstract theory”) is a collection of objective chances, population means or expected statistical frequencies. One may think of these as probabilities for classical propositions.  $p_S$  in contrast is defined over centered propositions, and is sensitive to one’s prior distribution over contexts.

Suppose, for example, that you are asleep in a room and that a fair coin is being tossed in an adjoining room. The coin is tossed two minutes after each *tails* toss and one minute after each *heads* toss. Since the coin is fair, the “abstract theory” has it that the probability of each toss landing *tails* is  $\frac{1}{2}$ , independent of the other tosses. Suppose now that you wake up and contemplate the indexical expression “the previous toss landed tails”. You should assign this expression’s secondary intension probability  $\frac{1}{2}$  (in agreement with theory), but the probability that you assign to its primary intension should, in the absence of further evidence, be  $\frac{2}{3}$ . (In the absence of further evidence, you should take yourself to have awakened at a uniformly random time.) So here we would say that  $p_O$ (“previous toss lands tails”) =  $\frac{1}{2}$ , but  $p_S$ (“previous toss lands tails”) =  $\frac{2}{3}$ .

Similarly, in the case where  $B$  is “my parents were fertile” Carter would say that  $p_O(B)$  might be .95 (some figure derived from a fertility baserate), but  $p_S(B) \approx 1$ . So it seems plausible to say that, when  $A$  is an indexical expression, Carter intends by  $p_O(A)$  the probability of the expression’s secondary intension, and intends by  $p_S(A)$  the probability of the expression’s primary intension. Since it is only after accounting for selection conditions (i.e. employing  $p_S$ ) that Bayes’ theorem is applied, that the posterior  $p_E(B)$  comes out to be greater than .95 isn’t due to distinctively “Bayesian” reasoning. It’s due rather to anthropic reasoning.

The critic’s work is not done, though, for as Huemer notes, it doesn’t matter whether we view the infertility example as distinctively “Bayesian” in flavor or not. Some feature of it blocks the critic’s objection, so in order for the critic’s objection to the reincarnation argument to be successful, the critic should say why such blocking does not occur in that particular application. Now the

feature blocking the critic's objection in the infertility example is plainly that  $p_S(B) \neq p_O(B)$ . So what the critic must show is that in the reincarnation case, one has  $p_S(H) = p_O(H)$ . This is right, I'll argue, for the following reason. In the reincarnation case the competing expressions  $H =$  "some Permissive theory of persons is true" and  $\sim H =$  "some Restrictive theory of persons is true" are non-indexical, and the primary intension of a non-indexical expression agrees with its secondary intension. So, since  $p_S(H)$  is the probability of  $H$ 's primary intension and  $p_O(H)$  is the probability of  $H$ 's secondary intension, we have  $p_S(H) = p_O(H)$ .

### The event/theory distinction

My proposal, then, is that (to some degree of approximation) it is the primary intension of the expressions to which one would assent that constitute one's evidence, and that, accordingly, hypotheses communicated by non-indexical expressions (call these *theories*) are not subject to the same selection pressures as hypotheses communicated by indexical expressions (call these *events*).<sup>9</sup>

Such an event/theory distinction is implicit in Carter (1983), in that he only dares apply his anthropic principle to the testing of rival hypotheses  $A$  and  $B$  of the "theory" type, and for these pairs always assumes that  $\frac{p_S(A)}{p_S(B)} = \frac{p_O(A)}{p_O(B)}$ . It is relatively more explicit in Lewis (2001), where one finds the following thought experiment involving three equally likely theories:

- A: God tosses coin, creates 10 persons if *heads* and 1000 persons if *tails*.
- B: God tosses coin, creates 10 persons if *heads* and  $10^{10^{100}}$  persons if *tails*.
- C: God tosses coin, creates 1000 persons if *heads* and  $10^{10^{100}}$  persons if *tails*.

His presentation is cryptic, but it's clear (it helps to be antecedently aware of an event/theory distinction) that Lewis subjects the toss of the coin (an indexical matter) to selection effects, but does not subject the choice of theory (a non-indexical matter) to such effects. In particular, he advocates that observers put the probability of  $M =$  *there are*  $10^{10^{100}}$  *persons* at approximately  $\frac{2}{3}$ .<sup>10</sup>

There are traces of an event/theory distinction in the writings of philosophers as well. John Leslie (1996, p. 248-250) suggested that selection pressures apply to Doomsday alternatives only if the process choosing which is actual is "radically indeterministic", and Nick Bostrom (2007) noted (the so-called "Presumptuous Philosopher gedanken") that it is less tempting to subject a "theory of everything"

<sup>9</sup>There are exceptions. Theories that predict no observers with certainty are obviously refuted by any observations whatsoever, and contingent propositions picked out by non-indexical expressions might (if there are such things) be exceptional as well. Lewis (2001) discusses theories that predict no observers with non-trivial probability, but offers only speculations as to how anthropic reasoning should treat such cases. These issues run orthogonal to our task here.

<sup>10</sup>Lewis would put  $p_S(B) = p_S(C) = p_O(B) = p_O(C) = 1/3$  (no selection effect), and  $p_S(M|B) = 10^{10^{100}} / (10^{10^{100}} + 10)$ ,  $p_S(M|C) = 10^{10^{100}} / (10^{10^{100}} + 1000)$  (selection effect). Note: Lewis is explicit that his anthropic reasoning treats finite cases as if they were iterated (under the auspices of the same theory) *ad infinitum*. So this case is relevant for us, despite its finiticity.

to the selection pressures that rational agents (Sleeping Beauty thirder, more generally) subject indexicals such as “the coin in question landed heads” to.<sup>11</sup>

### A symbolic dynamics example

Here’s a parting shot for those still on the fence. Suppose you know that you are the only conscious being in the universe, are immortal, and that all of your conscious observations are completely determined by a process that spits out characters  $a_1, a_2, \dots$  from the alphabet. Initially there are two equally likely candidate stationary theories describing the local statistical features of the output:

Theory A:  $P(a_n = \text{“C”}) = \frac{1}{10^3}, P(a_{n+1} = \text{“E”} | a_n = \text{“C”}) = \frac{1}{3}$ .

Theory B:  $P(a_n = \text{“C”}) = \frac{1}{10^6}, P(a_{n+1} = \text{“E”} | a_n = \text{“C”}) = \frac{2}{3}$ .

You know that you spend most of your existence asleep, but that whenever a “C” is spat out, you are aroused and get to observe the next letter spat, after which your memory of the episode is erased and you are put back to sleep.

Now...if you find yourself awake and having an observation of an “E”, which theory should that confirm? On the method of anthropic reasoning I have been advocating for in this paper, it confirms Theory B. (Your posterior probability that Theory B is true should be  $\frac{2}{3}$ .) That the first process would give you more (per unit time) “E”-observations is irrelevant, because your unconscious moments don’t count for anything. (Why would they?) In other words, it’s the frequency of “E” observations in the stream of observations that matters, not the frequency of “E” observations in the stream of all moments (conscious or not). To put it another way, no selection effect favors the first process.

### Conclusion

If the arguments I’ve given are correct then Huemer left out something crucial in assuming, with no justification, that there was “no reason to think that whatever would block the objection in the infertility case would somehow allow the objection to succeed in the reincarnation case”. There is a reason, because what blocks the objection in the infertility case, namely selection pressure induced by the context sensitivity or indexicality of the rival hypotheses considered there, is not present in the reincarnation case, since the rival hypotheses there are non-indexical. If that’s right, it’ll take some new argument to resurrect reincarnation. Until then, I’d advise against crossing anything off your bucket list you don’t remember doing.

---

<sup>11</sup>Probably neither Leslie nor Bostrom would affirm the event/theory distinction entirely, however, especially conditional on the universe being finite. In a passage reminiscent of Huemer (2019), for example, Bostrom (2001) wrote: “...it is hard to see what the relevant difference is between” a case in which the number of observers is a matter of theory and a case in which the number of observers is determined by coin tossing. However, given that Bostrom held that members of an observer’s “reference class” ought to be actual, in an infinite universe his anthropic reasoning agrees with that of Carter and Lewis. (In an infinite universe, the asymptotic statistics converge to that of the underlying theory almost surely.) Huemer’s position may be closest to the one that Bostrom attributes to his “Presumptuous Philosopher”, in which observers belonging to nomologically impossible worlds (Permissive theory of persons worlds, perhaps) can figure into the calculations of probabilities in self-selection experiments.

## References

- 
- Bostrom, Nick. 2001. The Doomsday Argument, Adam & Eve, UN<sup>++</sup>, and Quantum Joe. *Synthese* 127(3):359-387.
- Bostrom, Nick. 2007. Sleeping Beauty and Self-Location: a Hybrid Model. *Synthese* 157(1):59-78.
- Carter, Brandon. 1983. The Anthropic Principle and its Implications for Biological Evolution. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*. 310:347-363.
- Chalmers, David. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. OUP, Oxford.
- Huemer, Michael. 2019. Existence is Evidence for Immortality. *Nous*. Forthcoming.
- Leslie, J. 1996. *The End of the World: The Science and Ethics of Human Extinction*. New York: Routledge.
- Lewis, Antony. 2001. Comparing cosmological theories. Online. Available at <http://cosmologist.info/anthropic.html>.
- rmcctchn@memphis.edu*