Philosophy of Science (2026), **00**, 1–10
doi:10.1017/psa.2026.10257

PSA PHILOSOPHY OF SCIENCE ASSOCIATION
PHILOSOPHY OF SCIENCE

ARTICLE

Applying the Copernican Principle to the Problem of Induction

Jared Warren

Stanford University, USA
Email: jaredwar@stanford.edu

(Received 08 June 2025 UTC; revised 01 July 2026 UTC; accepted 05 July 2026 UTC)

Abstract

This article applies the Copernican principle and Gott's delta-*t* method to the infamous Humean problem of induction. Copernican reasoning uses *only* the observed duration of a phenomenon to predict its future duration, so this reasoning is noncircular. It also does not involve positing necessary connections, as do many other attempts to answer Hume's challenge. Copernican reasoning is controversial, as are many of its applications, but the application to Hume's problem is one that even most critics of the Copernican principle should accept. If Copernican reasoning is *ever* applicable, it is applicable here. Copernicus might have an answer for Hume.

1 Induction and the uniformity of nature

Nature is uniform. Not perfectly uniform, of course. But there are relatively time-stable objects with relatively time-stable properties, and there are detectable regularities between neighboring events.[1] Nature has been uniform for all of human history, and—from what we can tell—for long before as well.

Does the observed past uniformity of nature provide us with any reason for thinking that uniformity will continue into the future? The most obvious argument that it does is viciously circular—it simply assumes that the past serves as a rule for the future, which is the very point at issue. We've known this since 1739, when Hume published the first volume of his *Treatise of Human Nature*.

Hume infamously used this point to challenge the good standing of inductive reasoning, as such. In effect, he argued that *every* inductive argument presupposes

[1] This sounds like I'm endorsing three-dimensionalism over four-dimensionalism and anti-Humeanism over Humeanism about laws of nature, but I am not. The uniformity that I'm gesturing at can be stated in four-dimensional and Humean terms, although a bit less naturally. Here, I won't try to be more specific than this about the content of this general uniformity claim, but see Mackie (1979) for some relevant discussion.

© The Author(s), 2026. Published by Cambridge University Press on behalf of Philosophy of Science Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press

that nature is uniform across the observed and unobserved portions of reality. But because this uniformity claim is contingent, any successful argument for it must itself be inductive, and thus, it will presuppose the very uniformity claim at issue, making it viciously circular. If Hume was right, a successful justification of induction requires a noncircular argument that the observed uniformity of nature extends to the hitherto unobserved portions of reality.[2]

There is one popular contemporary strategy for noncircularly establishing general uniformity on the basis of observed uniformity. This is the *necessitarian* strategy, which first uses inference to the best explanation (abduction) to infer the existence of nomic necessary connections in order to explain observed uniformity, then deduces the uniformity of the unobserved from these inferred necessary connections.[3] This strategy is controversial. Many philosophers doubt that strong nomic necessary connections can reasonably be inferred using abduction on observed uniformities.[4]

Here I will explain an alternative noncircular strategy for vindicating belief in the uniformity of nature, one that does not appeal to either abduction or necessary connections. The strategy instead applies the widely discussed but controversial *Copernican principle*.

2 The Copernican principle and the delta-*t* method

With apologies to Aristarchus of Samos, Copernicus taught us that we are not observing the universe from a special spatial position. This is the general Copernican principle: *ceteris paribus*, your position as an observer is probably not special. This means that in certain circumstances and in a certain sense, you could just as well have been observing from a different position, and this is relevant to our reasoning. Copernicus applied this principle spatially, but it can also be applied temporally.

This basic idea leads to a completely general and widely applicable method—the delta-*t* method—for estimating the future duration of some event or process based solely on its past duration at the time of observation. This method was pioneered by, and has been widely applied by, the physicist J. Richard Gott.[5] It has subsequently been discussed very widely, even in popular books and articles.[6]

Gott originally came up with the delta-*t* method when he visited the Berlin Wall in 1969. At the time, the wall was 8 years old. Gott used the past age of the wall, along with the Copernican principle, to estimate its future duration, reasoning as follows.

Suppose that the total time that the Berlin Wall will stand is T. The time between Gott's observation, t_{now}, and the beginning of the wall, t_{start}, is 8 years. Gott reasoned that because t_{now} was equally likely to be at any point in T, there was a .5 probability

[2] Of course, Hume's claim that inductive reasoning presupposes a contingent uniformity assumption has itself been challenged. For example, see Sober (1988) for an argument that Hume commits a quantifier shift fallacy here, and see Stove (1986) for an argument that induction does not and could not require the truth of any contingent claim.

[3] Versions of this strategy are pursued in, for example, Armstrong (1983), Ellis (2002), and Foster (2004).

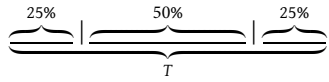
[4] See Beebee (2011) for this sort of criticism, and see Hildebrand (2018) for a reply.

[5] See Gott (1993, 1996, 1997a, 1997b) and Tyson et al. (2016).

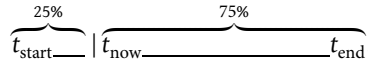
[6] See Ferris (1999), Gott (2002), and Poundstone (2019).

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press

that t_{now} occurred in the *middle* 50% of T:



If t_{now} is at the start of the middle 50% of T, as in this illustration:



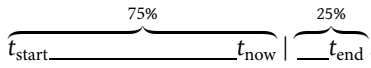
then:

$$T = 4 \times (t_{\text{now}} - t_{\text{start}}),$$

and:

$$(t_{\text{end}} - (t_{\text{now}} - t_{\text{start}})) = 3 \times 8 \text{ years} = 24 \text{ years}.$$

If t_{now} is instead at the end of the middle 50% of T, as in this illustration:



then:

$$T = \frac{4}{3} \times (t_{\text{now}} - t_{\text{start}}),$$

and:

$$(t_{\text{end}} - (t_{\text{now}} - t_{\text{start}})) = \frac{1}{3} \times 8 \text{ years} = 2\frac{2}{3} \text{ years}.$$

Putting these points together, we have that at t_{now} (in 1969), the probability that the Berlin Wall will last between $2\frac{2}{3}$ years and 24 years beyond 1969 equals .5:

$$\text{Prob}\left(2\frac{2}{3} \text{ years} < (t_{\text{end}} - t_{\text{now}}) < 24 \text{ years}\right) = .5.$$

Indeed, 20 years later, in 1989, the Berlin Wall fell. This successful prediction led Gott to publish his delta-*t* method.

The philosophical justification behind the delta-*t* method is rooted in the Copernican principle, applied temporally. Unless you have reasons for thinking that your time of observation is somehow special, you should reason *as if* you were randomly placed within the duration of some observed process or event.[7] This is the philosophical justification for the general delta-*t* method, which tells us, for a real

[7] This seems equivalent to reasoning as if you were randomly sampled from the class of relevantly similar (possible?) observers, but this framing raises *metaphysical* questions that I won't be discussing (although see Bostrom [2002] for discussion of these and other relevant background issues). It also raises epistemic issues, including, as a referee for this journal has noted to me, an assumption of randomness analogous to the randomness assumption made in the combinatorial approach to the problem of induction (Williams 1947; Stove 1986). For arguments that no special and problematic assumption of randomness is needed in this context, see McGrew (2001) and Campbell and Franklin (2004). It is also worth noting that in some such contexts, a fairly weak principle of symmetry or exchangeability allows us to derive the needed probabilistic randomness claim.

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press

number $0 \leq r \leq 1$, there is probability r that

$$\left[\left(\frac{1-r}{1+r}\right)(t_{\text{now}} - t_{\text{start}})\right] < [t_{\text{end}} - (t_{\text{now}} - t_{\text{start}})] < \left[\left(\frac{1+r}{1-r}\right)(t_{\text{now}} - t_{\text{start}})\right].$$

This general formula allows prediction, at credence level r, of how much longer a process will last after the time of observation based solely on the duration of the process at the time of observation.

The delta-t method is applicable *only* when you have no additional information about how long the observed process is likely to last. This means, for example, that we cannot apply the method to conclude that an 87-year-old man is more likely to live for a further 15 years than a 30-year-old woman. This is because we have empirical information about (for instance) the average human life span, relative life spans of human sexes, and the limits of human longevity. This kind of information would make the delta-t method inapplicable. So would information about a correlation between the probability of an observation and the duration of a process—when attending a wedding, you can't apply delta-t to predict a quick divorce because in this case, it is certainly not as if your observation were randomly placed somewhere in the duration of the marriage.

Despite these limitations, Gott and others have applied the delta-t method very widely. Critics have usually argued that, at least for many of these applications, we do have additional information that makes the delta-t method inapplicable. Yet even these critics typically agree that when we have absolutely no prior information about the expected duration of a process, and there is no observation selection effect involved in our point of observation, the delta-t method is rationally applicable. I will return to these crucial background assumptions, but first let me explain how the delta-t method can be used to argue that the so-far-observed uniformity of nature will almost certainly continue into the future.

3 The Copernican principle and the uniformity of nature

Here is the argument. Let's start with the provisional factual claim that humans have directly observed the uniformity of nature for the last $200,000$ years. We did not observe the beginning of uniformity, but we know that it began earlier than $200,000$ years ago. Later, I will discuss rolling back this assumption, but for now, let us proceed with it. Suppose that the total duration of the uniformity is U. We are at u_{now}, and we know that u_{start} was at least $200,000$ years ago. That is, we know that:

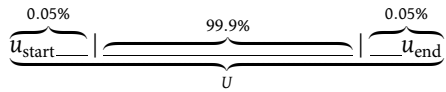
$$(u_{\text{now}} - u_{\text{start}}) \geq 200,000 \text{ years}.$$

To simplify things, let us assume:

$$(u_{\text{now}} - u_{\text{start}}) = 200,000 \text{ years}.$$

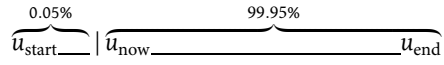
Calculations based on this assumption will give us estimates based on the greatest lower bound for $(u_{\text{now}} - u_{\text{start}})$. Now we apply the Copernican principle and the delta-t method at a 99.9% confidence level:

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press



Obviously, this illustration is not drawn to scale!

Assume first that our observation takes place at the earliest point in the 99.9% confidence interval:



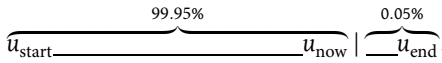
Then we have:

$$U = 2,000 \times (u_{\text{now}} - u_{\text{start}}),$$

and:

$$(u_{\text{end}} - (u_{\text{now}} - u_{\text{start}})) = 1,999 \times 200,000 \text{ years} = 399.8 \text{ million years}.$$

If u_{now} is instead at the other end of the middle 99.9% of U:



then we have:

$$U = \frac{2,000}{1,999} \times (u_{\text{now}} - u_{\text{start}}),$$

and:

$$(u_{\text{end}} - (u_{\text{now}} - u_{\text{start}})) = \frac{1}{1,999} \times 200,000 > \text{years} \approx 100\frac{1}{20} \text{ years}.$$

Putting these points together and assuming that the uniformity will end, the probability that the observed uniformity of nature will continue for between $100\frac{1}{20}$ years and 399.8 million years is .999:

$$\text{Prob}\left(100\frac{1}{20} \text{ years} < (u_{\text{end}} - (u_{\text{now}} - u_{\text{start}})) < 399.8 \text{ million years}\right) = .999.$$

Obviously, this entails that it is rational to be virtually certain that the uniformity of nature will continue through the end of all of our current experiments and endeavors:

$$\text{Prob}((u_{\text{end}} - (u_{\text{now}} - u_{\text{start}})) > 1 \text{ year}) \approx 1.$$

For smaller and smaller temporal intervals, we get closer and closer to rational certainty. Our expectation that the uniformity of nature will continue for the next few seconds is as close to rational certainty as empirical beliefs ever get, even if we use *only* the delta-*t* method applied to the observed duration of uniformity.

In full generality, suppose that some particular uniformity U has lasted for duration $t = t_{\text{now}} - t_{\text{start}}$. What is the probability that U will last for another $\frac{1}{n}$th of t? To answer this, suppose you predict that U will last at least another $\frac{t}{n}$. This prediction

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press

will only fail if the total duration of U, T_{total} is less than $t + \frac{t}{n} = t(1 + \frac{1}{n}) = t(\frac{n+1}{n})$. So, we have the following:

$$\text{Prob}\left(U \text{ will last another } \frac{t}{n}\right) = \text{Prob}\left(T_{\text{total}} \geq t\left(\frac{n+1}{n}\right)\right).$$

To calculate this probability, we can use the Copernican principle to determine the probability that our observation point is in the last $(n+1)$th of U. The odds against this are n to 1, so we have the following:

$$\text{Prob}\left(U \text{ will last another } \frac{t}{n}\right) = \frac{n}{n+1}$$

And this holds for any n. So, the probability that U will last for another $\frac{1}{100}$th of t is $\frac{100}{101}$, and for divisions of t into smaller intervals, the relevant probability gets closer and closer to 1.[8]

Even if we knew nothing save that the uniformity of nature has lasted for over $200,000$ years, it would be rational for us to be virtually certain that said uniformity will continue into the near future. On its face, this reasoning does not circularly use inductive reasoning, nor does it use abductive reasoning, nor does it infer the existence of controversial necessary connections.[9] Assuming this is right, we have a justification for induction that depends only on the Copernican principle applied temporally by way of the delta-t method. Given this, we can rationally infer that the observed uniformity of nature will almost certainly continue into the near future, and will quite probably continue into the distant future as well.

4 Trouble with the Copernican principle?

As far as I'm aware, nobody has yet attempted to apply the Copernican principle to the problem of induction. At least, not in anything like this form.[10] Perhaps this is because the general Copernican principle, understood as vindicating the delta-t method, is controversial.

For my purposes here, most of these controversies can be set aside. My overall argument is conditional: *If* the delta-t method can be applied to observed uniformity, *then* we have a noncircular reason for thinking that uniformity will continue into the future. Those skeptical of the method itself will resist discharging this conditional.

Yet, as noted earlier, even many severe critics of the delta-t method have admitted that it is applicable in certain special circumstances. Not every critic will agree, as

[8] This reasoning essentially uses the Copernican principle to derive the central formula of Harrod (1956). Harrod's "principle of experience" is, I believe, implicitly Copernican.

[9] We could also imagine using inference to the best explanation to conclude that some longer-lasting events (e.g., a structure continuing to stand) will likely continue because (continuing with the example) it having stood for such a long period is best explained by it being built in a particularly sturdy, long-lasting way. This reasoning, though, is different from the pure Copernican reasoning—see Monton and Kierland (2006) for some discussion.

[10] There is a forgotten strand in the literature on induction that pursues ideas that I believe to be related to the Copernican principle, starting with Harrod (1956), with both Blackburn (1973) and Mackie (1979) partially following up on Harrod's approach. A full discussion of these and other partial historical antecedents and their connection to the delta-t method must wait for another occasion.

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press

some will reject the entire framework or argue that certain Bayesian assumptions preclude Gott-style reasoning.[11] But I believe it is true even of Gott's most hostile critics, such as Carlton M. Caves. Initially, Caves suggests that a proper Bayesian analysis shows that Copernican reasoning using Gott's rule would be applicable only if we knew that a phenomenon was in progress but had no information about its past duration. But according to Caves, this blocks Gott-style applications because obtaining information about the past duration immediately makes *the moment of observation* special, thus undermining the Copernican approach and blocking any projections.[12]

However, later in his discussion, Caves discusses a possible rejoinder to this objection, writing:

> We have now uncovered the prior information that underlies the use of Gott's rule as a predictor of future duration; namely, knowing that a phenomenon is in progress, you cannot identify any time scales associated with the phenomenon either into the past or into the future. One way of thinking about this is that for a phenomenon that has no time scales, discovering the present age does *not* put you at a special time in the phenomenon' s history, so some consequences of the temporal Copernican principle survive. Whether the scale-free prior information is appropriate must be judged case by case; it is not a universal rule. (Caves 2000, 151)[13]

There remains disagreement about exactly what an "uninformative" prior is, in Bayesian terms, but even most of Gott's critics admit that *if* our prior was uninformative, in the relevant sense, then Copernican reasoning would be applicable.[14] Yet this is—arguably, at least—exactly the situation we are in when faced with the uniformity of nature.

At least this seems to be the case when we face Hume's problem in the internalist sense in which he originally posed it. Hume allows us our direct observations and past information about said observations, but little else. The allowance for past observations includes observations made by other humans, along with records of the past. If we were to instead allow only information derived from our own personal memories, then obviously the numbers would change, but the basic epistemic situation would not.

[11] For instance, an associate editor for this journal has noted to me a comparison between the delta-*t* reasoning and certain kinds of statistical reasoning that Bayesians are generally skeptical about. From a different direction and more generally, as noted later, some philosophers are skeptical of every attempt to assign prior probabilities *a priori*.

[12] As Gott (2000) observed in reply, this seems to make virtually all observations happen at a special time, in relation to some ongoing process.

[13] See also Bostrom (2002).

[14] Buch (1994) argued that a Bayesian account of Gott's (1993) delta-*t* method requires a non-normalizable prior probability and is problematic for that reason, but Gott (1994) replied that this was not a problem because (i) the posterior probability is still normalizable, and (ii) by putting appropriate bounds in place, a normalizable prior can be produced. This exchange connected the delta-*t* method to a so-called Jeffreys prior (see Jeffreys 1946). Caves (2000, 2008) criticizes this connection. Millican's (1982) reply to an argument from Mackie (1979) seems to me to be similar to parts of Caves's criticism.

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press

For this reason, I have not focused on worries that the base claim I started with—that humans have observed uniformity for 200,000 years—can only be arrived at with inductive reasoning and is thus, in this context, circular and therefore illegitimate. I think that there are some potentially promising responses to this sort of worry.[15,16] But even if these responses don't ultimately succeed, all hope is not lost for applying the Copernican principle to the problem of induction.

Even limiting our observations of uniformity to the uniformity that we have personally observed allows a reasonably long duration of "observed" uniformity—for adult humans, corresponding to the majority of our lives. Each of us can then use the Copernican strategy to get a reasonably high rational probability that uniformity will continue for the next small fraction of our lives as an observer. This pure version of the Copernican strategy is still available even if the just-discussed worries about the impure version prove decisive.

I won't reproduce the details of the full calculation, but for even an average 20-year-old with 15 years of memories of observed uniformity, we have, using the same kinds of assumptions from earlier, the following:

$$\text{Prob}\left(2\frac{3}{4}\ \text{days} < (u_{\text{end}} - (u_{\text{now}} - u_{\text{start}})) < 30,000\ \text{years}\right) = .999.$$

Treating the left-hand side as a lower bound, we can see that in any given moment, the delta-*t* method sanctions high confidence approaching certainty that the observed uniformity will continue into the next day, let alone the next moment. So, even the "pure" version of the strategy makes significant progress.

If we were to allow *only* our observations in a given moment, things are obviously much different, but this is not a problem with the Copernican strategy—*everyone* will have difficulty justifying a story of the world starting from this extremely limited basis.[17]

There are questions that can be raised about the Copernican argument. Yet many philosophers and scientists believe that delta-*t* reasoning is at least approximately applicable in cases where we observe some phenomenon without having any prior information about its expected duration or about positive or negative correlations of

[15] For example, one response would be to argue that it is *abduction*, not *induction*, that justifies our starting claim. This approach would make the delta-*t* reply a very close structural parallel to the standard necessitarian strategy mentioned at the outset (see the references in footnote 3). Another response would be to appeal to internalist principles that enable testimonial evidence to support claims quite directly, perhaps even *a priori*—for example, see Burge (1993) and Huemer (2001).

[16] In addition to pressing the circularity worry, a referee for this journal also raised a related arbitrariness worry: why limit our data to only human ancestors, thus excluding the observations of early hominids? I think that there are two strategies for replying to this. First, we could appeal to the details of one of the responses to the circularity worry suggested in the previous footnote. It may be that, for example, an abductive case leaves us a secure basis only for a more limited collection of observations. Second, we could appeal to a kind of self-sampling principle as mentioned in footnote 7, according to which observers should reason as if they were randomly sampled from the class of relevantly similar observers. One could then argue that only *Homo sapiens* observers are relevantly similar to us, in the needed sense.

[17] It is perhaps not completely hopeless, largely because human perception does not take place in frozen moments but instead over an interval known as the *specious present*. This was noted by Kelly (1882) and popularized by James (1893). This allows for some temporal structure even in the minimal units of awareness. This makes room for applying the delta-*t* method even in this case, but the initial gains would be small.

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press

our observations with the duration of the phenomenon.[18] Even this is not beyond all dispute. Some critics are skeptical of *all* predictions made in the absence of specific empirical information.[19] I think that this would cripple our ability to reason, but I won't try to refute the position here.[20]

Most of us think that Copernican reasoning is at least occasionally applicable. And that, if it is ever applicable, it is likely applicable to cases where we don't have prior information about the duration of a process. That is exactly the position we find ourselves in when facing down the infamous Humean problem of induction.

To be clear: I am not simply making the claim that there is a method that, if applicable, solves the problem of induction. That kind of conditional is not always interesting. Magic would solve the problem of induction, if magic were applicable to the problem. The point here is that there is a plausible, non–*ad hoc* method for making predictions of duration absent any empirical information, making the method apparently tailor-made for Hume's problem. Even critics of the method generally agree that the method is sometimes applicable, and if sometimes applicable, applicable in the situation in which it has been applied here.

Yet despite this, the method has never been applied to Hume's problem. My modest goal here has simply been to bring it to the attention of philosophers that Copernicus just might be able to vanquish Hume.

Acknowledgments. Thanks to J. Richard Gott, Bradley Monton, and two referees and an editor of this journal for comments.

Funding statement. None to declare.

Declarations. None to declare.

References

Armstrong, David M. 1983. *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Beebee, Helen. 2011. "Necessary Connections and the Problem of Induction." *Noûs* 45 (3):504–27. https://doi.org/10.1111/j.1468-0068.2010.00821.x.

Blackburn, Simon. 1973. *Reason and Prediction*. Cambridge: Cambridge University Press.

Bostrom, Nick. 2002. *Anthropic Bias: Observation Selection Effects in Science and Philosophy*. New York: Routledge.

Buch, Peter. 1994. "Future Prospects Discussed. "*Nature* 368:107–8. https://doi.org/10.1038/368107b0.

Burge, Tyler. 1993. "Content Preservation." *Philosophical Review* 102 (4):457–88. https://doi.org/10.2307/1523046.

Campbell, Scott, and James Franklin. 2004. " Randomness and the Justification of Induction." *Synthese* 138 (1):79–99. https://doi.org/10.1023/b:synt.0000012206.01154.c7.

[18] One might argue that uniformity is required for observers like us to exist. This is almost certainly at least approximately true, but it doesn't undercut the reasoning used here (see also Bostrom's [2002] discussion of necessary conditions for delta-*t* to be applicable). In more general terms, we are considering foundational epistemological reasoning, so no assumptions can be made up front about our observations depending on uniformity. That is something we only reach at a much later epistemic stage. One might also go on to argue that uniformity must exist for a long period before observers like us can evolve, but the same reply also works here, with the additional point that the greater past duration of uniformity simply gives us more grist for the Copernican mill.

[19] Sober (2003) is an example.

[20] See Monton and Kierland (2006) and Nelson (2009) for replies on behalf of delta-*t*. And see section VII of Putnam (1971) for a more general but related argument that all empirical reasoning requires *a priori* plausibility orderings.

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press

Caves, Carlton M. 2000. "Predicting Future Duration from Present Age: A Critical Assessment." *Contemporary Physics* 41 (3):143–53. https://doi.org/10.48550/arXiv.astro-ph/0001414.
Caves, Carlton M. 2008. "Predicting Future Duration from Present Age: Revisiting a Critical Assessment of Gott's Rule." *arXiv*. https://doi.org/10.48550/arXiv.0806.3538.
Ellis, Brian. 2002. *The Philosophy of Nature: A Guide to the New Essentialism*. New York: Routledge. https://doi.org/10.1515/9780773581715.
Ferris, Timothy. 1999. "How to Predict Everything." *New Yorker* 75 (18):35–36.
Foster, John. 2004. *The Divine Lawmaker: Lectures on Induction, Laws of Nature, and the Existence of God*. Oxford: Oxford University Press. https://doi.org/10.1093/0199250596.001.0001.
Gott, Richard J. 1993. "Implications of the Copernican Principle for Our Future Prospects." *Nature* 363: 315–19. https://doi.org/10.1038/363315a0.
Gott, Richard J. 1994. "Future Prospects Discussed." *Nature* 368:108. https://doi.org/10.1038/368108a0.
Gott III, Richard J. 1996. "Our future in the Universe." In *Clusters, Lensing, and the Future of the Universe*, edited by V. Trimble and A. Reisenegger, pp. 140–151. San Francisco: Astronomical Society of the Pacific.
Gott, Richard J. 1997a. "Tories out on Cue." *Nature* 387:842. https://doi.org/10.1038/43045.
Gott, Richard J. 1997b. "A Grim Reckoning." *New Scientist* 156 (2108): 36–39.
Gott, Richard J. 1998. "The Copernican Principle and Human Survivability." In *Human Survivability in the 21st Century. Transactions of the Royal Society of Canada*, series VI, vol. IX, edited by David M. Hayne, 131–47. Toronto: University of Toronto Press.
Gott, Richard J. 2002. *Time Travel in Einstein's Universe: The Physical Possibilities of Travel through Time*. Boston, MA: Houghton Mifflin.
Harrod, Roy. 1956. *Foundations of Inductive Logic*. London: Macmillan.
Hildebrand, Tyler. 2018. "Natural Properties, Necessary Connections, and the Problem of Induction." *Philosophy and Phenomenological Research* 96 (3):668–69. https://doi.org/10.1111/phpr.12351.
Huemer, Michael. 2001. *Skepticism and the Veil of Perception*. Lanham, MD: Rowan and Littlefield.
James, William. 1893. *The Principles of Psychology*. New York: Henry Holt and Company.
Jeffreys, Harold. 1946. "An Invariant Form for the Prior Probability in Estimation Problems." *Proceedings of the Royal Society of London* 186 (1007):453–61.
Kelly, E. Robert. 1882. *The Alternative: A Study in Psychology*. London: Macmillan.
Mackie, John L. 1979. "A Defence of Induction." In *Perception and Identity: Essays Presented to A.J. Ayer with His Replies to Them*, edited by Graham E. Macdonald, 113–30. Ithaca, NY: Cornell University Press.
McGrew, Timothy. 2001. "Direct Inference and the Problem of Induction." *The Monist* 84 (2):153–78. https://doi.org/10.5840/monist200184211.
Millican, Peter J. R. 1982. "Mackie's Defence of Induction." *Analysis* 42 (1):19–24. https://doi.org/10.1093/analys/42.1.19.
Monton, Bradley, and Brian Kierland. 2006. "How to Predict Future Duration from Present Age." *Philosophical Quarterly* 56 (222):16–38. https://doi.org/10.1111/j.1467-9213.2005.00426.x.
Nelson, Kevin. 2009. "How and How Not to Make Predictions with Temporal Copernicanism." *Synthese* 166 (1):91–111. https://doi.org/10.1007/s11229-007-9259-5.
Poundstone, William. 2019. *The Doomsday Calculation: How an Equation That Predicts the Future Is Transforming Everything We Know about Life and the Universe*. New York: Little, Brown, and Spark.
Putnam, Hilary. 1971. *Philosophy of Logic*. New York: Harper & Row.
Sober, Elliot. 1988. *Reconstructing the Past*. Cambridge, MA: MIT Press.
Sober, Elliot. 2003. "An Empirical Critique of Two Versions of the Doomsday Argument—Gott' s Line and Leslie' s Wedge." *Synthese* 135 (3):415–30. https://doi.org/10.1023/a:1023545820214.
Stove, David C. 1986. *The Rationality of Induction*. Oxford: Clarendon Press.
Tyson, Neil deGrasse, J. Richard Gott, and Michael Strauss. 2016. *Welcome to the Universe*. Princeton, NJ: Princeton University Press.
Williams, Donald C. 1947. *The Ground of Induction*. Cambridge, MA: Harvard University Press.

Cite this article: Warren, Jared. 2026. "Applying the Copernican Principle to the Problem of Induction." *Philosophy of Science*. https://doi.org/10.1017/psa.2026.10257

https://doi.org/10.1017/psa.2026.10257 Published online by Cambridge University Press