Time Machines, Oracles, and Easy Arguments

David Friedell and Patrick Grafton-Cardwell

Forthcoming in *Disputatio*. Please cite final version.

Abstract

Katrina Elliott's argument that time travel is unlikely can be extended to show that infallible predictors are unlikely. These are interesting metaphysical results, but the epistemology is just as interesting. Both arguments rely on trivial empirical observations about human psychology and about how easy it is to do certain things that would count as changing the past or defying an oracle. No physics or engineering required. In this way, the arguments might seem too easy. There are other easy arguments that rely on trivial empirical observations: Hempel's Raven Paradox, Bostrom's Simulation Argument, Leslie's Doomsday Argument, and Builes and Harris's Whale Argument. This paper defends the argument for infallible predictors being unlikely and suggests that those who endorse any of these other easy arguments should keep an open mind about Elliott's time travel argument and its extension to infallible predictors.

1. Introduction

Katrina Elliott (2018), building upon work by Paul Horwich (1987), argues that time travel, though metaphysically possible, is unlikely. Here we argue that Elliott's argument can be extended to show that oracles and other infallible predictors are metaphysically possible but unlikely. These are interesting metaphysical results, but the epistemology is just as interesting.

Both arguments rely on trivial facts about human psychology and about how easy it is to do certain things that would count as changing the past or defying an oracle. No sophisticated physics or engineering required. In this way, these arguments seem too easy and are epistemically strange.

These arguments, however, are not alone in their strangeness. This may be revealed by considering Hempel's (1945) Raven Paradox, Bostrom's (2003) Simulation Argument, Leslie's (1990) Doomsday Argument, and Builes and Harris's (2023) Whale Argument. In each case, there's an argument that certain trivial observations are evidence for the likelihood of a particular conclusion in a surprising way that feels too easy. We call these *easy arguments*.

Our central goals are three-fold. First, we will show how Elliott's time-travel argument can be extended to the case of infallible predictors. Second, we will defend the argument for infallible predictors being possible but unlikely by responding to some objections. Third, we will suggest that that if you're comfortable with any of the other easy arguments you should be open-minded about Elliott's time-travel argument and its extension to infallible predictors.

2. Time Travel

Imagine a time traveler, Fitz. Fitz forgets about his sister's birthday, and she doesn't hear from him on her birthday. The next day he realizes his error but remembers something important: he has a time machine! So, he uses the time machine to travel to yesterday, his sister's birthday. Upon arrival he tries to use his cell phone to call her to wish her Happy Birthday.

What happens next? Well, we'll let you finish the story. But here's something that won't happen, provided this mundane time-travel story stays logically consistent: Fitz won't wish his

sister Happy Birthday on her birthday. For, as stipulated above, she doesn't hear from Fitz on that day. Something will happen to thwart his attempt. Maybe his phone won't work. Maybe he'll go to the library to use a computer to email his sister but on the way slip on a banana peel. Maybe he'll type the email, but it will fail to send because of a technological issue he doesn't notice. No matter how hard Fitz tries, and no matter how many attempts he makes— via phone, email, text, or telegram—he will not wish his sister Happy Birthday.

This story might seem contradictory. On the one hand, here Fitz is on his sister's birthday with a desire to wish her Happy Birthday. There are various ways for him to contact her, and they all are easy. Given how easy it is to wish her Happy Birthday it seems Fitz can wish his sister Happy Birthday. On the other hand, given that he will not wish his sister Happy Birthday, every way he tries to wish her Happy Birthday will fail. This consideration makes it seem that, alas, he cannot wish her Happy Birthday. So, it seems that Fitz both can and cannot wish her Happy Birthday on the right day. And this might seem like a contradiction. One might infer that time travel to the past is impossible.

That would be a mistake. David Lewis (1976: 149) shows that there is no contradiction. To say someone can do something is relative; it means that their doing it is consistent with certain facts. Relative to the fact (let us suppose) that you ski frequently and are in good shape, you can ski. Relative to the fact that you are currently reading a philosophy paper and are nowhere near a snowy mountain, you cannot ski. The claim that you can ski and the claim that you cannot ski are both true, because each claim is made in relation to different facts. Likewise, relative to the fact that Fitz has access to the marvels of modern technology he can wish his sister Happy Birthday. Relative to the fact that his sister never hears from him on this day, he cannot. No contradiction. There might be some completely independent reason why time travel to the

past is metaphysically impossible, or at least inconsistent with our physical laws, but one should not infer from the story of Fitz alone that time travel to the past leads to any contradiction or is metaphysically impossible.

Instead, one should conclude, as Katrina Elliott (2018) does, that frequent time travel to the past is unlikely. Elliott's argument is a variant of Paul Horwich's (1987). Here is, roughly, her variant. She starts with this proposition: if people were to frequently travel to the past, many time travelers would try to change the past by doing easy things. This proposition is supported by trivial empirical facts, including facts about human psychology and facts about how easy it is to accomplish certain tasks. Some time travelers might try to wish their siblings Happy Birthday on days their siblings never hear from them. Wishing a sibling Happy Birthday is easy. Other time travelers might covet real estate they have never purchased and then go in the past to try to buy it at cheap prices. Buying cheap real estate is easy. Others might try to save a pedestrian who dies in a car accident by going back in time and pushing the victim out of harm's way. Pushing someone out of harm's way (especially when you know a car is coming) is easy.

Any attempt to change the past will fail. Doing something in the past that didn't happen is contradictory. It doesn't matter how easy it is to wish someone Happy Birthday, buy cheap real estate, or push someone out of harm's way. If something didn't happen, then it didn't happen. If you try to wish your sister Happy Birthday when she doesn't hear from you, you will fail. If you try to buy cheap real estate when you didn't buy it you will fail. If you try to prevent someone from dying when they died, you will fail. Time travelers who try to change the past by doing easy tasks will fail at easy tasks.[1]

And those with a time machine would try to change the past. Even we would try, though

[1] For discussion of models of time travel that allow changing the past, see Bernstein (2017).

we know it's impossible. It's just human nature. Thus, given this obvious fact about human psychology, there are two main possibilities: (a) there will be frequent time travel to the past, and many time travelers will try to change the past by trying to do easy tasks; all of those who do so will fail, or (b) there will not be frequent time travel to the past. Elliott recognizes that the first possibility is comparatively unlikely. It is unlikely that all the relevant time travelers will fail at accomplishing easy tasks. It is comparatively likely that, instead, there will not be frequent time travel to the past. We therefore have a reason—albeit a defeasible one—to think that there will not be frequent time travel to the past. Elliott concludes that it is unlikely that there will be frequent time travel to the past.

Elliott's argument tells us nothing about what prevents there from being frequent time travel. The explanation might be that time travel is too expensive, or that nobody learns how to build a time machine, or that it's physically impossible to time travel, or that some institution bans time travel, etc. Elliott's reasoning tells us only that frequent time travel to the past is unlikely without telling us why.

3. Infallible Predictors

Here we will extend Elliott's argument to argue that infallible predictors are possible but unlikely. In the next section we will defend the argument about infallible predictors against objections.

Suppose determinism is true. Now suppose you learn that someone is an oracle of a certain kind. She has complete knowledge of the initial conditions of the universe and the laws of nature, and the ability to accurately predict facts about what will happen by deducing her

predictions from this information. Now, suppose the oracle tells you that she knows with absolute certainty that you will order coffee in the next five minutes. The oracle's prediction doesn't initially strike you as very impressive. After all, you are with the oracle in a coffee shop, and you are visibly tired and in need of coffee. Eager to refute the oracle, however, you plan not to order coffee in the next five minutes. "I'll show you, oracle!" you think to yourself. All you have to do is not order coffee in the next five minutes. What could be easier than that?

This case generates the so-called "paradox of predictability." On the one hand, the oracle's prediction carries a lot of weight. Given that determinism is true and this is what allows the oracle to make accurate predictions, it seems impossible for you to falsify what she says (assuming she is honestly reporting her prediction). On the other hand, you will try not to order coffee in the next five minutes, and it is extremely easy for you not to order coffee in the next five minutes. Although this case is not genuinely paradoxical, there is something puzzling afoot. It both seems impossible for you to falsify the oracle's predictions and it seems easy for you to falsify them. How can that be?

Michael Scriven (1965) presents a seminal version of this sort of case. Some theorists, including Stefaan Cuypers and Stefan Rummens (2009) and Jenann Ismael (2016: 170-181), conclude that it is impossible for there to be an infallible predictor who (a) makes predictions about what you will do, and (b) shares those predictions with you. Ismael reaches this conclusion because she thinks you would be able to "thumb your nose" at any purportedly infallible predictor. That is, you would be able to falsify the predictor's prediction. If anyone claims to know with absolute certainty that you will order coffee in the next five minutes, you can easily defy them. Ismael infers two things from this. First, she infers that the laws are a certain way to resolve the "paradox"—viz. the laws should be thought of in a Lewisian way, as global patterns

in local matters of fact. Second, she infers that people have a certain kind of freedom that is compatible with determinism.[2] These inferences are related. The view of laws is posited as a general solution to the paradox of predictability that allows us to recognize that classical determinism doesn't pose a threat to freedom.

Ismael may be right that adopting a Lewisian understanding of the laws resolves the paradox of predictability in some cases where determinism plays an important role. This, however, does not provide a general solution to the paradox of predictability. One reason why is that equally puzzling situations can arise even if determinism is false. Let us describe two cases to show this.

Suppose indeterminism and theism are both true. Now suppose God tells you that God knows with absolute certainty that you will order coffee in the next five minutes. On the one hand, God's predictions carry a lot of weight. It's God. God knows everything. God knows you'll order coffee in the next five minutes. It seems like you will. On the other hand, you really want to take the opportunity to spoil God's prediction, and it is extremely easy for you not to order coffee in the next five minutes.

Take another scenario in which you, in an indeterministic universe, learn that someone is a time traveler from the future. The time traveler tells you that she knows with absolute certainty that you will order coffee in the next five minutes. She saw you do it already. On the one hand, the time traveler's predictions carry a lot of weight. She was there. She saw you order coffee. On the other hand, you really want to take the opportunity to spoil the time traveler's prediction, and

[2] Note that Ismael does not conclude that it is impossible for there to be an infallible predictor tout court. She remains neutral about whether it is possible for there to exist outside of our universe a being who makes infallible predictions about us and doesn't share any of those predictions with us. She thinks merely that the paradox of predictability shows that it is impossible for there to be an infallible predictor who exists inside of our universe and shares with us their predictions about us.

it is extremely easy for you not to order coffee in the next five minute. These scenarios show the paradox of predictability doesn't require the truth of determinism. Determinism being true is just one conceivable way one could come to meet an infallible predictor. Since the underlying problem could conceivably arise in an indeterministic world, we should want a more general solution than the adoption of a certain view about the laws.

We will now highlight two ways time machines and infallible predictors are analogous. First, we should not infer from the story of Fitz that time travel to the past is metaphysically impossible. There is nothing contradictory about Fitz traveling to the past, trying to wish his sister Happy Birthday by trying to do easy tasks, and then failing. Moreover, there is nothing contradictory about many time travelers trying to change the past by doing easy tasks and all of them failing. Analogously, we should not infer from our story of the oracle that infallible predictors who share predictions about many people's behavior are metaphysically impossible. There is nothing contradictory about an infallible predictor sharing her prediction you will order coffee in the next five minutes and then you trying but failing to prove the predictor wrong. Moreover, there is nothing contradictory about many people trying to refute an infallible predictor by trying to do easy tasks and all of them failing.

Second, just as the case of Fitz reveals that frequent time travel to the past is unlikely, the case of the oracle in the coffee shop reveals that an infallible predictor sharing predictions with many people about their behavior is unlikely. Here is why. Like Elliott's argument concerning time travel, we start with this proposition: if there were an infallible predictor who shared predictions with many people about their behavior, then many people would try to falsify some of the predictor's predictions by trying to do things that are easy. This proposition is supported by trivial empirical facts, including facts about human psychology and facts about how easy it is

to accomplish certain tasks. Some people might try not to order coffee when an infallible predictor predicts they will. Not ordering coffee is easy. Some people might try to say "hi" when an infallible predictor predicts they will say "hello" instead. Saying "hi" instead of "hello" is easy. Some people might try to jump when the predictor says they will sit still. Jumping is easy. It is human nature to try to falsify predictions of any purportedly infallible predictor. Many people would want to show that they have free will—or that they are unpredictable creatures, more nuanced than any predictor or algorithm can imagine. If nothing else, many people would try to falsify the predictions of an infallible predictor out of mere curiosity: a desire to see what will happen when they try to refute the predictor.

Any attempt to falsify the prediction of an infallible predictor will fail. Infallible predictors make only correct predictions. If an infallible predictor predicts you will do something, then you will do it. Not doing something you do is contradictory. It doesn't matter how easy not ordering coffee is or how easy saying "hi" is. If you try not to order a cup of coffee when you do order a cup of coffee you will fail. If you try to say "hi" when you instead say "hello" you will fail. People who try to falsify predictions made by an infallible predictor about their behavior will fail at easy tasks.

Given these facts, then, there are two main possibilities: (a) an infallible predictor will share predictions with many people about their behavior, many of those people will try to falsify the predictor's prediction by trying to do easy tasks; all of them will fail, or (b) there will not be an infallible predictor who shares predictions with many people about their behavior. The first possibility is comparatively unlikely. It is unlikely that all the relevant wannabe falsifiers will fail at accomplishing easy tasks. It is comparatively likely that, instead, that there will not be an infallible predictor who shares predictions with many people about their behavior. We thus have

a reason—albeit a defeasible one—to think that there will not be an infallible predictor who shares predictions with many people about their behavior. Note that our reasoning does not rely on claims about whether determinism is true or false. Nor does it rely on claims about the nature of the laws. Relying only on basic claims about human psychology and about how easy it is do certain tasks, we conclude that it is unlikely that infallible predictors will share many predictions with people about their behavior.

We should make explicit an assumption of our reasoning and Elliott's. Like Elliott (2018: 6-8) we assume that there are objective "macro-chances" about events that are independent from relative frequencies, micro-chances, and subjective degrees of confidence. We assume, for example, that the likelihood or macro-chance of a coin landing on heads may be 50% even if (a) coins (randomly) land slightly more on heads than on tails (which might very well be true if there are an odd number of coin flips in the universe), or (b) there's an underlying microscopic deterministic process that guarantees the coin will land on heads in this instance, or (c) a relevant observer thinks the chance is higher or lower than 50%. In this sense, it's unlikely for time travelers to fail to wish their siblings Happy Birthday. Likewise, it's unlikely for someone to order coffee when they try really hard not to. Crucially, this second failure is unlikely, even if an infallible predictor has predicted it. The infallible predictor has merely predicted something unlikely will happen.[3]

Just as Elliott does not tell us why time travel is unlikely—only that it is unlikely—our reasoning is silent about why it is unlikely infallible predictors will share many predictions with people about their behavior. It might be that it is too expensive to build a machine that infallibly

[3] We won't argue for this metaphysics of probability here. Our argument and Elliiott's argument are still epistemologically interesting, even if one treats them as a reductio of macro-chances.

predicts human behavior. It might be that there are infallible predictors, but they would rather keep their predictions to themselves. It might be that a government will ban the existence of infallible predictors. Our reasoning tells us only that it is unlikely that infallible predictors will share many predictions with people about their behavior, without telling us why. We don't know why. Alas, we don't have a crystal ball.

4. Is it Ever Unlikely to do what an Oracle Predicts?

One might object to our argument about infallible predictors by claiming that it is actually *likely* that someone will order coffee when an infallible predictor predicts they will order one, even if they try really hard not to order a coffee.[4] More generally, one might claim that if someone ever tries to disprove an infallible predictor, then failure is likely. In fact, it's guaranteed! We claim that some of these failures are unlikely, especially when someone is trying to defy a predictor by trying something easy (such as not ordering a coffee). If it turns out, as the objector claims, that such failures are likely, then our argument doesn't work.

In response, we insist that in the relevant sense of "unlikely" infallible predictors can predict unlikely outcomes. That's part of what makes them so impressive. Even if it's likely (and guaranteed) in the micro-sense of "likely" that someone will order coffee when they try not to—because their failure is entailed by a complete description of the initial conditions of the universe and the physical laws—their ordering coffee may still be unlikely in the macro-sense of "unlikely". This point might be easier to appreciate if we suppose that the oracle is making frequent predictions about all sorts of people. That is, the predictor doesn't pick on incompetent people, or people who

[4] Elliott (2018: 18-22) discusses analogous concerns regarding time travel.

are otherwise prone to accidentally ordering things at cafes.

Here's an analogy one could offer in response. The statistician Persi Diaconis and a group of engineers have built a machine that, when working properly, flips a coin heads 100% of the time (or thereabouts). Upon learning this, how surprised should you be? Perhaps you should be somewhat surprised, since it's an impressive feat of engineering. But it would be wrong to think it's unlikely for such a machine to exist on the grounds that it's unlikely for a coin to land on heads many times in a row. The machine by its very design makes such an outcome likely. Similarly, one might object that it's wrong to think it's unlikely for someone to accidentally order coffee when an infallible predictor predicts they will order coffee. The prediction is like the Diaconis machine. It makes it likely the predictee will order coffee. At least, that's what one could argue.[5]

We respond with an analogy of our own. Imagine you watch a highlight reel of the ten biggest upsets in sports history. One of the upsets is "The Miracle on Ice", when the U.S. men's hockey team defeated the much better Soviet Union team in the 1980 Olympics. This upset is likely in a subjective sense of "likely"; even if you know little about sports history, you are confident the U.S will win, given that you are watching a highlight reel of *upsets*. Plausibly, it's even likely *now* in the macro-sense of "likely" that the U.S. won, given that they did. None of this changes the fact that in 1980 before the game occurred the outcome was unlikely in the macro-sense. It was an upset, after all. We think an oracle's prediction that you will drink coffee is more like the highlight reel than the Diaconis machine. That machine *makes* coins come up heads. The oracle's prediction doesn't make you drink coffee. Nor did the future existence of the highlight reel make the U.S team win in 1980. Moreover, the fact that an oracle has predicted you will order coffee does not make it any harder for you to refrain from doing so. It's still easy to refrain, which

[5] Thanks to Amy Seymour for helpful discussion.

makes your failure to refrain unlikely.

One might object that it's crucial that the highlight reel exists *after* the victory, whereas the oracle's prediction occurs *before* you drink coffee. The oracle's prediction, on this line, makes it likely that you will order coffee, in a way that future highlight reels can't affect how likely outcomes were in the past. We note in response that the mere fact someone says something will happen, even if they say so truly, doesn't make its future occurrence likely. Some lottery winners say they're going to win before buying a ticket. Their confident, accurate declarations don't make their subsequent victories likely. Granted, the oracle doesn't just accurately say you will order coffee. She *knows* you will. But it's not obvious why this would make it likely in the macro-sense that you will order coffee when you try not to.

Along these lines, imagine an oracle predicts that you will flip a fair coin on heads ten times in a row. Her prediction might make you confident the coin will land on heads ten time in a row. Indeed, it *should* make you confident. But it doesn't make the coin landing on heads ten times in a row likely in the macro-sense. She has predicted something unlikely will happen. Giving you confidence that the coin will land on heads ten times doesn't make you better at flipping coins so they land on heads. Granted, getting a pep talk from an oracle might make some tasks easier, such as cartwheels and shooting free throws. But confidence is useless for coin-flipping. The coffee case is similar. By predicting you will order coffee the oracle in effect reveals that when you try not to order coffee you will fail. She thereby reveals something unlikely will happen. The prediction doesn't make it any harder for you to refrain from ordering coffee. It's still easy. Just as an oracle's prediction of an unlikely sequence of coin tosses doesn't make that sequence any more likely, the oracle's prediction that you will order coffee doesn't make it any more likely you will order coffee when you try not to. You will order the coffee when you try not to, but it's

unlikely. For these reasons, we stand by our claim that if there were an oracle who shared many predictions with people about their behavior there would be many unlikely failures to do easy tasks. Our argument stays intact.

5. Easy Arguments

It's interesting, albeit disappointing, that there likely won't be any time machines or infallible predictors. Just as interesting is how we've reached these conclusions. We did so, in a word, easily. You might have thought that in order to know whether there will be a time machine you should consult an engineer or a physicist. You might have thought that in order to know whether there will be an infallible predictor, you should consult a theologian (in case the predictor is a supernatural oracle or deity) or an engineer (in case the predictor is a device or algorithm built by Amazon or some other technology company). But our reasoning hasn't relied on engineering, physics, or theology. We know little about those subjects. We're just humble country philosophers.[6]

We've relied on obvious psychological claims: that people would try to change the past or defy an oracle. And we've relied on trivial empirical claims about how easy it is to do certain tasks, such as wish someone Happy Birthday or not order a coffee. Nothing hard. You might suspect that there must be something wrong with our reasoning. It might seem too easy.
But it would be hasty to reject our reasoning on these grounds alone. After all, there are other arguments in philosophy that argue we can reach conclusions from trivial empirical observations in surprisingly easy ways. We'll now consider some of these easy arguments.

[6] We borrow this bit from Mark Balaguer.

First, as Elliott (2018) notes, there is Hempel's (1945) Raven Paradox. The statement "all ravens are black" is logically equivalent to "all non-black things are non-ravens." This suggests that seeing a white sofa is at least some evidence for all ravens being black, given that it is a non-black non-raven. Doing ornithology indoors might also feel too easy.

Another example is Bostrom's (2003) simulation argument. Bostrom argues that one of the following three claims is very likely to be true. (1) all human-like civilizations go extinct before being able to create simulations that contain conscious beings; (2) no human-like civilization capable of creating such a simulation creates any; (3) we are currently in a simulation. A fourth possibility—that we're in the base level of reality and will create simulations that might in turn contain simulations—seems extremely unlikely. It seems extremely unlikely to be in base reality if the vast majority of conscious beings are simulated.

You might have thought that in order to discern whether you're in a simulation you should consult a physicist or engineer who, for example, could tell you whether there are any "glitchy" occurrences that are evidence of this world being artificial. But, if Bostrom is correct, you can get evidence about whether you're in a simulation just from keeping up with current events. If a war involving governments with nuclear weapons calms down that suggests that it's more likely that humanity will survive long enough to build a simulation, which could increase your confidence in the proposition that you are currently in a simulation. Just as the fact that it is easy to wish someone Happy Birthday might feel like it shouldn't be evidence against time machines, it might feel like how a war is going should be irrelevant to the question of whether you are currently in a simulation. But that's where Elliott's and Bostrom's arguments lead us.

A related easy argument is Leslie's (1990) Doomsday argument. It goes roughly like this. We should assume, unless given evidence to the contrary, that we're likely typical: somewhere in

the middle of humans who will ever exist, not particularly close to the beginning or end. Supposing there have been 100 billion humans before us (which is in line with some estimates), we should therefore expect that there will be about 100 billion more. Give current birthrates, we may conclude that we have evidence humanity will last only hundreds or thousands of years more—not millions or billions. You might have thought if you wanted to know when humanity will go extinct you should consult political scientists about the chances of nuclear war, or epidemiologists about the chances of pandemics, or geologists about the expected effects of climate change. But this argument suggests that one's birth rank is evidence for how long humanity will last in a surprising way.

The last easy argument we'll mention is David Builes and Caspar Hare's (2023) Whale Argument against mereological universalism: the view that any collection of material objects composes another object. Suppose this view is true. Then, not only do a whale's atoms compose an object (namely, the whale) but every smaller collection of those atoms composes an object that is part of the whale. Plausibly, many of those composite objects are conscious, just like the whale is. For example, an object composed of all of the particles in the whale minus a single atom in its eye is conscious. It turns out, given how big whales are, there are way more whale-like conscious things than human-like conscious things.[7] But, now, behold the fact that you are conscious but not a part of a whale! What are the odds of that being true if mereological universalism is in fact true? Much lower than if mereological universalism were false. We thus have an easy argument against mereological universalism that doesn't require pouring over books of a priori metaphysical investigation. All you need to do is think about the empirical fact that you are not a part of a whale.

[7] This is an understatement. See Builes and Hare (2023) for discussion of the math.

We won't defend here responses to Hempel's Paradox that insist that a white sofa is evidence (albeit tiny) for all ravens being black. Nor will we defend the Simulation Argument, the Doomsday Argument, or the Whale Argument. All of this is controversial. Nor will we argue that these easy arguments stand or fall together. Some of them might be better than others. We raise these examples merely to suggest that if you're comfortable with easy arguments in any of these contexts, you should not be discouraged by the easiness of easy arguments about time machines and oracles.

In all of these cases, there's an argument that some piece of relatively trivial empirical information is evidence for a conclusion in a way that might seem too easy. For some of these arguments the relevant empirical observations are so trivial that they almost feel like a priori arguments for contingent truths. Our argument about oracles argues that the fact that people would try to disprove an oracle by trying to do easy things is evidence that it's unlikely there will be oracles. It's reasonable at first glance to be suspicious of any or all of these arguments, given how easy they are. But we hope both that our defense of the oracle-argument and the comparison to other easy arguments will inspire open-mindedness.[8]

References

Bernstein, Sara. (2017) "Time Travel and the Movable Present" in John Christopher Adorno (ed.), Being, Freedom, and Method: Themes from the Philosophy of Peter van Inwagen: 80-94.

Builes, David and Casper Hare. (2023) "Why Aren't I Part of a Whale?" Anaylsis 83, 2: 227-234.

[8] Thanks to Meghan Cenci, Katrina Elliott, Don Fallis, Peter Lewis, Ned Markosian, Amy Seymour, the audience at LanCog 2024, and an anonymous referee for helpful discussion.

Cuypers, Stefaan E., and Stefan Rummens. (2009). “Determinism and the Paradox of Predictability.” Erkenntnis 72, 2: 233-249.

Elliott, Katrina. (2018) “How to Know that Time Travel is Unlikely Without Knowing Why.” Pacific Philosophical Quarterly.

Hempel, Carl. (1945), “Studies in the Logic of Confirmation”, Mind, 54: 213.

Horwich, Paul. (1987). Asymmetries in Time: Problems in the Philosophy of Science. Cambridge: MIT Press.

Ismael, J. T. (2016). How Physics Makes Us Free. Oxford.

Lewis, David. (1976.) “The Paradoxes of Time Travel.” American Philosophical Quarterly 13: 2. 145-152

Leslie, John. (1990) “Is the End of the World Nigh?” Philosophical Quarterly 40: 158, 65-72.

Scriven, Michael. (1965). “An essential unpredictability in human behaviour.” In B. B. Wolman and E. Nagel (Eds.), Scientific psychology: Principles and approaches (411-425). New York: Basic Books.