

Laws, Melodies, and the Paradox of Predictability*

Chris Dorst

Abstract

If the laws of nature are deterministic, then it seems possible that a Laplacean intelligence that knows the initial conditions and the laws would be able to accurately predict everything that will ever happen. However, it would be easy to construct a *counterpredictive device* that falsifies any revealed prediction about its future behavior. What would then occur if a Laplacean intelligence encountered a counterpredictive device? This is the *paradox of predictability*. A number of philosophers have proposed solutions to it, though part of my aim here is to argue that the paradox is more pernicious than has thus far been appreciated, and therefore that extant solutions are inadequate. My broader aim is to argue that the paradox motivates Humeanism about laws of nature.

1 Introduction

Suppose the laws of nature are deterministic. Imagine a Laplacean intelligence that knows all the laws of nature, and also knows the initial conditions of the universe. Also imagine that this Laplacean intelligence has unlimited computational capabilities. Using its knowledge of the deterministic laws and the initial conditions, this intelligence should then be able to predict, without error, everything that will ever happen in the entire history of the universe.

Now suppose we construct a *counterpredictive device*, which works as follows. The device has a red and a green light on it. It takes as input a prediction about which color light it will turn on, and then turns on the color opposite of the prediction. If it is fed the prediction “Red,” it turns on the green light; if it is fed the prediction “Green,” it turns on the red light. Clearly, it would be easy to construct such a device in a universe like ours.

What will happen if the Laplacean intelligence makes a prediction about the output of the counterpredictive device, and then we feed this prediction into the device? Of course, we have seen that the intelligence ought to be a perfect predictor. Short of computational error, it is hard to see how it could fail. But the device is designed to frustrate any prediction that it receives about its action. Short of malfunction, it is hard to see how *it* could fail.

*Thanks to John Biro, Rodrigo Borges, Taylor Cyr, Finnur Dellsén, Kevin Dorst, Cameron Gibbs, James Gillespie, Jenann Ismael, Nic Koziol, Greg Ray, and two anonymous referees for helpful advice, comments, and suggestions. Thanks also to an audience at the 2021 Central APA for very valuable feedback.

This is the *paradox of predictability*. It received a good deal of philosophical attention in the 1960s and 1970s, before largely falling out of vogue. Lately it has seen a moderate resurgence in interest, first in Rummens and Cuypers (2010) and more recently in two insightful pieces by Ismael (2016, 2019).¹ A variety of different solutions have been proposed, but here I will primarily focus on the two that Ismael has suggested. In *How Physics Makes Us Free* (2016), Ismael uses the paradox to motivate Humeanism about laws of nature—the view that laws of nature are patterns in the totality of the particular matters of fact. By contrast, Ismael (2019) proposes an entirely different solution based on the unknowability of certain facts about initial conditions, which I will call the “Nomic Preclusion” solution. My aim here is to defend Ismael’s earlier solution. I think that the paradox can be turned into a compelling argument for Humeanism.

This paper proceeds as follows. In §2 I clarify a few points about the paradox and then classify the general kinds of solutions that have been proposed. In §3, I explicate the Humean solution to the paradox, and in §4, I explain the Nomic Preclusion solution. In §5, I suggest a problem for the Nomic Preclusion solution, and in §6 I argue that this problem also affects any solution other than adopting Humeanism about laws. In §7 I explore a potential objection on behalf of the anti-Humean, and I argue that it is unsuccessful. In §8 I consider just how thoroughly Humean the paradox requires us to be: as it turns out, the paradox leaves room for *some* anti-Humean necessary connections in nature. I conclude in §9.

2 The Paradox and the Types of Solutions

Three clarifications are in order about the paradox. First, it is not really a paradox. Or rather, it is only a paradox *if* Laplacean intelligences and counterpredictive devices are compossible.² Really, then, it is a demand for explanation: an explanation of the fact that Laplacean intelligences and counterpredictive devices cannot coexist.

But you might think that this demand is misplaced. We already know that there is no possible world where a genuine Laplacean intelligence meets a non-malfunctioning counterpredictive device in the prescribed way. Such a situation is logically impossible. Why should we demand an explanation of a logical impossibility?³

¹My presentation of the paradox here follows Ismael (2019). Representative papers from the earlier discussions are Mackay (1960), Scriven (1965), Boyd (1972), and Landsberg and Evans (1970).

²When I talk of Laplacean intelligences and counterpredictive devices not being compossible, I mean to imply that (i) the intelligence makes a prediction about the output of the device, (ii) the device is provided that prediction as input, and (iii) the device does not malfunction. It is the specific situation where all of (i)-(iii) obtain that is impossible. If any of these conditions does not obtain, then Laplacean intelligences and counterpredictive devices may perfectly well coexist.

³Garrett and Joaquin (2021) present a compelling case for the idea that we need no explanation of this logical impossibility. To see the point more clearly, we might compare our

I think it is right that we shouldn't *demand* an explanation. It may be that there is nothing more to say here beyond noting that the imagined situation is logically impossible. Nevertheless, we should still prefer to have an explanation if there is one to be found. In a structurally parallel situation, mathematicians often distinguish between explanatory and non-explanatory proofs. One may have a proof of theorem T that demonstrates T 's truth without *explaining why* T is true. In such cases, it is common practice to continue to search for an explanatory proof that elucidates why T is true. Sometimes no such explanatory proof exists, in which case we may have a "mathematical coincidence." But in other cases, we may find a proof that explains why T holds; if so, we prefer it on these grounds to the non-explanatory proof.⁴

Of course, if a mathematical theorem T is true, then it is mathematically necessary. What we have here, then, is a methodological principle that tells us to prefer explanations of at least some sorts of mathematical necessities if they are available. I am recommending that we adopt essentially the same methodological principle in our present case. We know *that* a certain situation is logically impossible: a Laplacean intelligence will never accurately predict the behavior of a non-malfunctioning counterpredictive device, so long as its prediction is fed into the device. Still, we should prefer to have an explanation of this impossibility if one is available. For example, perhaps the true impossibility somehow lies in the idea of a Laplacean intelligence itself. If such beings are logically impossible, then of course it is logically impossible that such a being would encounter a counterpredictive device in the prescribed manner. But we would miss out on this insight if we stopped our analysis at the logical impossibility of the encounter itself. So asking for an explanation of why the encounter is impossible seems eminently reasonable, even if there is no guarantee of an illuminating answer.

The request for explanation may feel more compelling if one adopts a first-personal perspective. Imagine yourself in the position of the intelligence. If there is a counterpredictive device in your future, then you cannot *now* know the initial conditions and laws of nature and correctly calculate their consequences. But why not? What could be preventing you from doing so? Logical impossibility, no doubt. But this logical impossibility is based on something that happens in the *future*: some yet-to-occur event is making it logically impossible that you put yourself in a particular epistemic situation in the present.

situation to Turing's proof that no Turing machine constitutes a general solution to the Halting Problem. Turing showed that if there were such a "general halting machine," then we could easily create another Turing machine that was bound to frustrate it. Given a description of this "counterpredictive machine" with an arbitrary input, we ask the supposed general halting machine if the counterpredictive machine will halt with that input. If it says yes, the counterpredictive machine goes into an infinite loop and never halts. If it says no, the counterpredictive machine halts. So a general halting machine is logically impossible.

That situation is structurally analogous to our present one. But Turing and other computability theorists didn't think we needed an *explanation* of that impossibility. So why should we think we need one here? (Gijssbers (2021) argues that the structural analogy between Turing's proof and the paradox of predictability is in fact relevant to the paradox's solution.)

⁴See Lange (2016, especially chapters 7-9) for some illustrative examples.

That, I think, is quite mysterious indeed. This perspective thus motivates the idea that simply observing the logical impossibility is not enough: we should prefer to have an explanation of the impossibility—something that reveals why an embedded agent should not expect to be able to know all the laws and initial conditions. When I talk of finding a resolution to the paradox, this kind of explanation is what I have in mind.⁵ (I’ll motivate the search for explanation further in §7.)

The second clarification is that the *embeddedness* of the Laplacean intelligence is important. A Laplacean intelligence that was external to the universe and was somehow “watching it play out” could, on the basis of its knowledge of the laws of nature and the initial conditions, accurately predict the entire future, even one that contains counterpredictive devices. That is because this external intelligence’s prediction about the output of the counterpredictive device is not a physical event that occurs in the history of the universe, and therefore it is not one that a counterpredictive device existing in that universe could operate on and be guaranteed to falsify. The fact that the prediction’s being a physical event is crucial to the paradox is emphasized in a very helpful paper by Rummens and Cuypers (2010). Note, however, that this point does not resolve the paradox. What we want is an explanation of why an embedded Laplacean intelligence cannot coexist with a counterpredictive device. Observing that the prediction has to be a physical event in order to generate the paradox does not explain why those two entities cannot coexist.⁶

The third clarification is that it is not essential that the Laplacean intelligence knows the initial conditions of the entire universe. The intelligence only needs to know the antecedent conditions at some moment in the past to be able to predict all future states.⁷

How, then, should we attempt to explain this impossibility? Following the lead of others who have written on the paradox, we may note that counterpredictive devices are *obviously* possible, and take this to indicate that the real impossibility lies in the idea of an embedded Laplacean intelligence. To recap, a Laplacean intelligence is a being that knows the total state of the universe at some past instant, as well as the laws of nature, and has unlimited computational power. Such a being should be able to plug the total state into the laws of nature and compute everything that will ever happen. Strategies for explaining the impossibility of embedded Laplacean intelligences have therefore fallen into

⁵I take it that this is also the kind of explanation that has been sought by other authors writing on the paradox, e.g. those cited in footnote 1. None of them thought that the paradox could be solved simply by noting the relevant logical impossibility.

⁶C.f. Gijssbers (2021), p. 6: “But of course we need to say more in order to dispel the paradox. We need to explain *why* or *in what sense* determinism fails to deliver predictability.”

⁷I am assuming here a particular form of determinism, namely forward determinism, where the initial state plus the laws determines all subsequent states. Other kinds of determinism would make the intelligence’s task even easier. If in addition the laws were reverse deterministic, the intelligence would arguably have an easier time discerning the total state at some particular time. Stronger forms of determinism might require only that the intelligence knows the boundary conditions inside some highly circumscribed region of spacetime (even just a single event) in order to use the laws to infer the content of the entire spacetime manifold.

three categories:

1. Deny that an embedded agent can know the total state required for its predictions.
2. Deny that an embedded agent can compute the consequences of the laws.
3. Deny that an embedded agent can know the laws.

A few words are in order about each of these strategies.

Consider strategy 1. Why should we think that an embedded agent cannot know the total state of the universe at any past instant? There are at least two ways that one might motivate this idea. First, it might be suggested that an embedded agent cannot measure the values of physical variables to infinite precision. If so, then there is no way to eliminate a small degree of uncertainty about the exact value of any physical variable, and even minor imprecision in the initial values of some physical variables can amplify into extreme imprecision down the road. The rate at which this error amplification occurs will vary depending on both the nature of the laws and the nature of the physical system in question. In chaotic systems, for example, any measurement errors will amplify at an exponential rate. For other systems, the error amplification rate will be lower, but even a linear growth rate will eventually produce predictions that are wildly divergent from actuality.⁸

A second way to motivate this strategy is to suggest that even if one could somehow measure the values of physical variables to infinite precision, it would not be possible to *represent* them for calculational purposes. An embedded agent cannot perform calculations on numbers that are only representable by infinite strings of digits. Thus for the purposes of calculation (and hence prediction), the numerical representations of variable values will have to be truncated somewhere. This is again going to introduce a degree of error into any predicted values, and this error will amplify over time, as described above.

One reason to be skeptical of these ways of motivating this strategy is that they rely on error amplification *over time*. The general approach here is to start with an embedded intelligence that is subject to the given epistemic restrictions (it cannot measure or represent values to infinite precision) and then to find some future event that is beyond the predictive scope of that intelligence. But what happens if the event of the counterpredictive device's operation is *within* the predictive scope of the intelligence? Nothing has been said to preclude such a possibility. So while we may have been given an explanation of why a perfect, embedded Laplacean intelligence cannot exist, we have not been given an explanation of why an embedded intelligence, operating within the relevant epistemic constraints, cannot accurately predict the outcome of a counterpredictive device that is within its predictive horizon. Thus we have not really gotten to the heart of the matter.⁹

⁸See Bishop (2003) and Stone (1989) for more in-depth discussion of these points.

⁹When I talk of the intelligence's "predictive scope" or "horizon" I mean to denote everything that is entailed by its knowledge of the physical variables in question. You might

Consider now the second strategy: denying that an embedded agent can compute the consequences of the laws. If determinism is true, the laws of nature define a function f_L mapping from the state of the universe at a given time t_1 to the state of the universe at any given future time t_2 , i.e. $S_2 = f_L(S_1)$. Call f_L a *nomic function*. The idea here, suggested by Boyd (1972), is that it may turn out that f_L is not an effectively computable function, and therefore that an embedded Laplacean intelligence would have no method available by which it could compute the future states of the universe.¹⁰

This strategy is only worth brief consideration here because Boyd’s task was not to show why embedded Laplacean intelligences are impossible, but rather to show that determinism does not entail predictability in principle. Thus, while this strategy may show why embedded Laplacean intelligences are impossible in universes where f_L is not effectively computable, surely there are plenty of universes where f_L is effectively computable. Indeed, we should hope that ours is such a universe if we want to be able to use the laws to make predictions. This strategy thus has nothing to say about the impossibility of embedded Laplacean intelligences in universes where f_L is effectively computable.¹¹

Lastly, one might pursue the third strategy and argue that an embedded intelligence could not possibly know the laws of nature. Admittedly, this might seem to be a rather strange approach, for it is not clear why the laws should be unknowable in principle. Moreover, physicists have long sought to discover the laws of nature, so a blanket denial of the knowability of the laws would fly in the face of scientific practice.

But it turns out that we do not need a *blanket* denial of the laws’ knowability to achieve our goals here. I originally framed our task as follows: explain why embedded Laplacean intelligences and counterpredictive devices cannot coexist. Given that counterpredictive devices manifestly *can* exist, we took this to imply that our task was to explain the impossibility of an embedded Laplacean intelligence. This was a natural reaction, but it may have been too coarse-grained. Perhaps what we need to explain is not the impossibility of embedded Laplacean intelligences *simpliciter*, but rather the impossibility of embedded Laplacean intelligences whenever there is a (non-malfunctioning) counterpredictive device in

then think that the intelligence’s predictive scope is never going to include the output of a counterpredictive device—no matter how nearby in the future that device operates—because any uncertainty at all in the initial variables will result in *some* degree of uncertainty about the output of the counterpredictive device. But it would not be difficult to design a counterpredictive device whose operations were “resilient” to this kind of uncertainty. For example, we might imagine that so long as the position of a lever is within a certain allowed range, the device will definitely turn on its green light. In that case, as long as the initial uncertainty does not ramify into uncertainty greater than the allowed range on the position of the lever, the output of the counterpredictive device falls within the intelligence’s predictive scope.

¹⁰Roberts (2008, p. 345) also suggests that the mathematical problems the intelligence would have to solve may be uncomputable.

¹¹Moreover, even if f_L is not effectively computable, it is likely that there are approximations of f_L that are effectively computable, and are accurate within limited timescales. Such approximations would then have a built-in predictive horizon beyond which they are insufficiently accurate, but again, we would only need to posit that the counterpredictive device operates within that horizon to generate the paradox.

the intelligence’s future.

As it relates to this third strategy, the idea is that an embedded intelligence cannot know the laws of nature *if* there is a counterpredictive device in its future. Ismael pursued a strategy along these lines in her (2016), before abandoning it in more recent work. My aim here is to defend it. The key to this strategy is to adopt Humeanism about laws of nature.

3 The Humean Solution

Humeanism about laws of nature is the view that laws are certain sorts of patterns in the totality of the particular matters of fact. Nowadays the canonical Humean view is David Lewis’s Best System Account (“BSA”), according to which laws are the regularities of the simplest and strongest systematization of all the particular facts that obtain in the total history of the universe.¹² There is disagreement among Humeans about how to pick out exactly which patterns count as laws, but for our purposes we can set this issue aside. What matters for us is that according to Humeanism, the laws are essentially a *summary* of the total history of the universe.

In contrast to Humeanism, anti-Humean views take the laws to be elements of the universe that *govern* or *produce* the patterns in the particular matters of fact. Characteristic anti-Humean theories are those of Armstrong (1983) and Maudlin (2007). Armstrong regards laws as relations of nomic necessitation between universals, whereas Maudlin treats laws as ontologically fundamental aspects of the universe that produce subsequent states from its current state. Another anti-Humean view is that of Bird (2007), who argues that laws of nature are consequences of the metaphysically necessary connections between properties.¹³

According to an anti-Humean like Maudlin, Armstrong, or Bird, the laws are currently-existing elements of the universe. If one were to enumerate the “current ontology,” to make a list of everything that exists *right now*, the laws of nature would be on that list, either as independently-existing entities, or in the case of Bird’s view, they would at least be *implied* by things on that list, namely by the properties instantiated by objects in the universe. According to such views, the laws either *produce* future states from the present state, or they *constrain* the way future states can be in light of present states. In this vein, Ismael (2016, p. 111) analogizes anti-Humean laws to “iron rails” leading the current state of the world into the future.

¹²See Lewis (1973, 1986, 1994) for canonical statements of the view. Various developments of the BSA have been proposed over the years. See, for example, Loewer (1996, 2007), Cohen and Callender (2009), and Demarest (2017) for modifications of the account of the particular matters of fact that get systematized. And see Callender (2017), Hicks (2018), Dorst (2019), Jaag and Loew (2020), and Loewer (2020) for modifications of the systematizing standards that get applied to those facts.

¹³Carroll (1994) develops an anti-Humean view similar to Maudlin’s; Tooley (1977) and Dretske (1977) advocate views much like Armstrong’s; Shoemaker (1980), Ellis (2001), and Mumford (2004) advocate dispositional essentialist views similar to Bird’s.

Not so for the Humean. According to her, laws are a function of all the particular matters of fact, past, present, and future. They are not, then, determined by the totality of facts obtaining at any particular instant in that universe's history. Ismael (2016, pp. 176-177) likens them to the property of "being a tragedy." *Hamlet's* being a tragedy is not something that drives its story forward. Rather, it is a property of the overall story arc, and *Hamlet* does not count as a tragedy until the whole story has played out, as it were. Alternatively, as one of my students recently suggested, Humean laws are like a melody. A melody is determined by all the notes that compose it, but the notes themselves are not "governed" or "produced" by the melody. The melody is not "there already" as each of the notes is being played, just like Humean laws are not "there already" as the history of the universe is playing out.¹⁴

Of course, this does not imply that an agent that finds themselves *in* that history, sometime before the end, could not know the laws. One may have good inductive evidence supporting some theory of what the laws are, and that theory may in fact end up being the true theory of laws. There is nothing mysterious about this (problems of induction aside). It is just like recognizing that a given melody is "My Heart Will Go On" after only the first four notes: we can have good inductive evidence that this is true, even though part of what makes it true are facts about the future that have not yet obtained.

For either the Humean or the anti-Humean, then, an embedded intelligence could know the initial conditions and the laws of nature. That much is straightforward. But things are *not* so straightforward when there are counterpredictive devices in the future of that intelligence. If a would-be Laplacean intelligence makes a prediction about the output of a counterpredictive device, and we feed that prediction into the counterpredictive device, it will falsify the prediction. Assuming that the intelligence did in fact know the initial conditions and correctly calculated the implications of those conditions from what it took to be the laws of nature, it follows that those cannot have been the actual laws.

All of what I just said in the last paragraph was true on both Humeanism and anti-Humeanism. Where the views come apart is when we ask *why* the intelligence must have been wrong about the laws.

The Humean has an elegant answer to this question: the intelligence had to be wrong because of what laws are. They are summaries of what happens rather than rules that determine what happens, so they *depend* on what happens. Of course, part of what happens in this universe is that the intelligence predicts the outcome of the counterpredictive device, and the device falsifies the intelligence's prediction. And so, still assuming that the intelligence knew the initial conditions and calculated correctly, there is no way that a true deterministic summary of that universe would allow that the intelligence used the laws (read: used *that very summary*) to generate a false prediction. For the Humean, then, a counterpredictive device is essentially a mechanism for *rigging the summary* against the predictor: no embedded intelligence whose prediction

¹⁴See Beebe (2000, pp. 578-579) for a similar way of contrasting Humeanism and anti-Humeanism.

is fed into such a device could have known the correct summary and used it to generate its prediction. Since that summary constitutes the laws, we thereby have an illuminating explanation for why the intelligence must be wrong about the laws.

No analogous answer is available for the anti-Humean. For someone like Maudlin, Armstrong, or Bird, the laws are not determined, in part, by the prediction made by the intelligence, nor by the behavior of the counterpredictive device. They are already *there to be known* by the intelligence; as we said, they are part of the universe's "current ontology." In short, the laws are currently-existing elements of the universe at the time the intelligence goes to make its prediction. So unlike on the Humean view, on the anti-Humean view there is no explanation of why the intelligence must be wrong about the laws if there is a counterpredictive device in its future. In this situation, it is just a *brute impossibility* that the intelligence knows the laws.

Here is another way to make the point. Imagine that at the beginning of the universe, God writes down, in English, an exhaustive enumeration of everything that will ever happen. She puts this on a stone tablet, which you happen to find in your backyard. Moreover, suppose that there is a counterpredictive device in your future, and that by reference to the stone tablet, you predict which color light that device will turn on. (It says this, of course, in plain English on the tablet.) And suppose I subsequently feed your prediction into the device, so it lights up the color opposite of your prediction. What has gone wrong?

Well, there is no possible world where (a) you read the tablet correctly, (b) your (unaltered) prediction is input into the device, and (c) the device does not malfunction. Given (b) and (c), then, you must have misread the tablet. Granted, it is hard to see how you *could* have misread "red" for "green," but there really is no other possibility here. No matter how carefully you read and now matter how many times you double-check, you still misread it. Stranger still, if I tell you beforehand that I'm going to feed your prediction into a reliable counterpredictive device, you should *expect* to misread it, no matter how careful you are.

In this analogy the tablet plays the role of the anti-Humean laws. It is a currently-existing element of the universe that fixes the future outcome of the counterpredictive device. If the laws are conceived that way, our resources for explaining the impossibility of a correct prediction of the output of a counterpredictive device are extremely limited. If the intelligence knows the initial conditions and calculates correctly, it must—*somehow*—have been wrong about the laws, no matter how good its evidence, and no matter how thorough its investigations. But in contrast to the Humean explanation, here we have no account of the source of the error. All we can say is that the intelligence must be wrong, end of story. The anti-Humean therefore has no principled explanation of the impossibility.

Insofar as we find Ismael's Humean solution compelling, it looks like we will have to eschew eternalism and adopt some form of presentist or growing block metaphysics of time. For an essential ingredient to the Humean solution is the denial that, at the time of the intelligence's prediction, there are facts determin-

ing the output of the counterpredictive device. But according to eternalism (or even some sort of moving spotlight theory), all such events already exist, and thus there already are facts, both about what the counterpredictive device does in the future, and more generally about what the best systematization of the mosaic is. Facts about the Humean laws will then already be decided when the intelligence goes to make its prediction, and we cannot appeal to their status as summaries of what happens to explain why the intelligence cannot know them. Instead we are left with a brute impossibility, just as on anti-Humeanism. Only a presentist or growing block view of time will suffice for the Humean solution to work, because only on such a view are there genuinely no facts about what the laws are when the intelligence goes to make its prediction.

It is an interesting question whether Humeanism about laws can be plausibly conjoined with a presentist or growing block metaphysics of time. Certainly the combination would be nonstandard, as the Humean mosaic is most easily interpreted in eternalist terms. But I do not see any decisive impossibility here. Why not allow that the mosaic grows over time, for example? Indeed, the Humean should probably hope that their view is compatible with presentism, if only because an incompatibility here would somewhat decrease the probability that their view is correct. To be sure, a presentist Humeanism engenders plenty of questions that would need attention (e.g. is it consistent with relativity?), but it seems fair to proceed on the assumption that the view can be made coherent and plausible enough. At the very least, our results below should motivate further exploration of this question.

It is worth noting, in passing, that Beebe and Mele (2002) have given a similar argument aiming to show that adopting Humeanism about laws allows us to escape van Inwagen's (1983) Consequence Argument against compatibilism. Briefly put, the Consequence Argument is as follows:

If determinism is true, then our acts are the consequence of laws of nature and events in the remote past. But it's not up to us what went on before we were born, and neither is it up to us what the laws of nature are. Therefore, the consequences of these things (including our present acts) are not up to us (van Inwagen, 1983, p. 56).

Beebe and Mele's basic thought is that the Humean has reason to question the premise that the laws of nature are not up to us. If we adopt the Humean view that the laws are determined by the totality of particular facts, and we note that our actions are part of that totality, then it is no longer obvious that we have no influence over the laws. It is interesting that roughly the same idea is operative in both arguments: namely, for the Humean, the laws are not currently-existing elements of the universe driving it toward the future. This is what allows the Humean to escape the paradox of predictability, and also what apparently allows the Humean to escape the Consequence Argument.

4 The Nomic Preclusion Solution

Ismael has since changed her diagnosis of the solution to the paradox, and she no longer advocates Humeanism about laws as the way out. Rather, in her more recent work (2019) she advocates a version of strategy 1: deny that an embedded agent can know the total state required for its prediction. The idea is roughly that the laws of nature and accompanying spacetime structure may systematically preclude an embedded agent from knowing everything that would be required to generate a reliable prediction of the output of a counterpredictive device. Ismael makes this argument for both Newtonian and Special Relativistic settings, and we will consider both in turn.

In a Newtonian regime, the laws are deterministic¹⁵, but knowing the total state of the world and the laws is actually not sufficient for knowing all future states: the history of the world up to a certain time t , plus the Newtonian laws, does not entail the state at $t+1$. This is because it is not merely the collection of facts in the history that, coupled with the laws, is required for the subsequent state. Rather, it is the collection of those facts plus the totality fact that *there is nothing else*. In other words, the fact that there are no additional forces operative at t is relevant to what happens at $t+1$. So an embedded agent would have to know not only all of the forces that are operative, but also the fact that those are all the forces. Without this additional knowledge, there is no entailment to future states of the universe. And the key is that, as Ismael puts it, the fact that *those are all the operative forces* “is not information that is available to any embedded subsystem of the world...” (*ibid.*, p. 483). The idea here is that the information that is available to an embedded agent about the total state is comprised of particular matters of fact: that there is a yellow pencil on the desk in my office, that there are five chairs in the hallway, etc. But no amount of knowledge of particular matters of fact is going to entail anything about a totality fact. An embedded agent, then, will never be in position to know that its knowledge of the universal state is complete, and therefore will never have information that entails the behavior of a counterpredictive device in its future.

The story is slightly more complicated in a Special Relativistic spacetime. In such a setting, the factors that can influence a particular event e are limited to the past lightcone of e .¹⁶ So if one wants knowledge that entails what occurs at e , one has to know a cross-section of e 's past lightcone. Now suppose that e is the output of the counterpredictive device, and suppose that a would-be Laplacean intelligence is trying to predict that outcome in advance, and that the prediction will be fed into the device. If all that is true, then the intelligence's prediction must be in the past lightcone of the output of the counterpredictive device, else it could not be fed into it. The prediction itself will be an event—call it p —that occurs in the timelike past of e , i.e. $t_p < t_e$. Now, p itself has

¹⁵Setting aside unusual situations like Norton's dome; see Norton (2008).

¹⁶This is the standard view of relativity (see e.g. Malament (1996)), though some theorists argue that the speed of light should be treated as simply an invariant quantity rather than as a speed limit for causal influences (see e.g. Maudlin (2011)).

a past lightcone, and again the only factors that can influence it are found inside of it. And the key is that the past lightcone of e is at all times “wider” than the past lightcone of p : the set of factors that can influence p is a proper subset of the factors that can influence e . That is to say, the factors that can influence the intelligence’s prediction at t_p do not contain all of the factors that can influence the action of the counterpredictive device at t_e . This means that the information available to the intelligence is insufficient to determine the outcome of a counterpredictive device in its future. The “missing” information is distributed in the area between p ’s past lightcone and e ’s past lightcone. Or equivalently, that information is distributed across any timelike curve connecting p and e . As long as p occurs before e —that is, as long as p is a genuine prediction—there will be timelike worldlines between them that contain the information necessary to fully determine e that was not available at p .¹⁷

In both Newtonian and Special Relativistic settings, then, the laws of nature preclude an embedded agent from knowing enough about the antecedent conditions to be in a position to perfectly predict the output of a counterpredictive device in its future. This new solution to the paradox is, of course, friendlier to anti-Humean theories like Maudlin’s, Armstrong’s, and Bird’s. It gives them a way to explain the impossibility of a Laplacean intelligence even while maintaining that the laws are currently-existing elements of the universe when the intelligence goes to make its prediction. The intelligence may very well know the laws, but it cannot know all of the initial conditions necessary to entail what the counterpredictive device will do.¹⁸

5 Knowing and Guessing

There is, however, a problem with this “Nomic Preclusion” solution, which we will build up to here.¹⁹ To begin with, note that this solution appears potentially sensitive to the nature of the relevant laws and spacetime structure. What we have seen is that given a Newtonian universe, there is no way for an embedded agent to come to know the requisite totality fact. And likewise, given a Special Relativistic universe, there is no way for an embedded agent to know a complete cross section of the past lightcone of the future action of a counterpredictive device.

But are *all* conceivable sets of laws like that? Do they all preclude an

¹⁷See Geroch (1977) and Manchak (2008) for further discussion of the difficulties of prediction for agents embedded in relativistic spacetimes.

¹⁸It is worth noting that nothing about this solution is incompatible with Humeanism about laws. The important point is just that it is compatible with anti-Humeanism.

¹⁹One problem that I will only mention, but not pursue: the Nomic Preclusion solution seems to imply much too strong a result about the limits of our knowledge, namely that any contingent universally-quantified claim of unrestricted scope (e.g. “There are no gold cubes with a mass greater than 100,000,000 kilograms”) is unknowable to an embedded agent, since it relies on a totality fact. This conclusion seems untenable. An embedded agent may have good enough evidence to count as “knowing” such a claim, even if its evidence does not *entail* that claim. In the ensuing discussion, however, I will grant that an embedded agent could not know such things, because I want to focus on a more revealing problem with the proposal.

embedded agent from coming to know everything required to predict the counterpredictive device's action? This is not so obvious. Imagine, for example, a universe that operates according to laws very much like those of John Conway's game of Life. The space of such a universe is comprised of a grid, where each cell is in one of two states: "occupied" or "unoccupied." The laws for the evolution of the grid over time are quite simple. Whether a given cell will be occupied at t_{n+1} depends solely on the states of its eight adjacent cells at t_n . Moreover, suppose that in this Life-like universe there are very clear "edges" to the grid. (Perhaps they are even labeled: "This is the edge of the universe.") Now, what reason is there to think that an agent embedded within that universe could not navigate around the entire space, gradually come to know the states of each individual cell at some time or other, and use that information together with the laws to piece together the total state of that universe at some past instant, and moreover to know that this was the total state?²⁰

I can see no principled reason to believe that this is impossible. One might suggest that in a universe with such simple laws, there could be no embedded intelligence capable of storing all of the requisite information about the states of various cells at various times. But I have not even specified exactly what the laws of this universe are, beyond saying that they are Life-like, so it is hard to see what motivation there could be for thinking that a sufficiently complex embedded agent could not exist in such a universe. Indeed, Life universes can support quite complex entities—such as universal Turing machines—even though their dynamical laws are quite simple.

Alternatively, one might suggest that, while such an agent is indeed possible in a Life-like universe, they would not be in a position to truly *know* that their information about the state is complete. Rather, at best they would have some evidence that they know the total state, but no definitive proof that that was indeed the total state—maybe because, as Ismael argues, they cannot come to know the requisite totality fact solely through knowledge of particular matters of fact. Perhaps, for example, they would not be able to know that what look like edges (and are labeled as edges) really are genuine edges, and therefore they would not be able to know that there are no cells outside of them.

Here, however, things start to seem rather fishy. What we are now suggesting is that while the embedded agent might indeed be *right* about what the total state was at some past instant, it cannot *know* that was the total state. In other words, its epistemic position with respect to the total state falls short of knowledge; it is a mere true belief, or something analogous thereto. The issue—and this is really the crucial point—is that it is difficult to see why the *strength* of the intelligence's epistemic position with respect to the total state should have any bearing on the paradox. As long as the intelligence is correct in its belief about the total state—no matter how it came to have that belief, and no matter the strength of its justification for that belief—the paradox can

²⁰If we consider deterministic theories (like those mentioned in fn. 7) in which the intelligence only needs to know a very limited region of spacetime, such as a single event, to use the laws to infer the contents of the rest of the spacetime manifold, then the predictive task becomes even easier.

get a foothold.

To make this clearer, return to our Newtonian universe. In that universe, the total state at all past times has already obtained. Now suppose that an agent navigates around the universe and gathers a large amount of information about particular matters of fact at some particular time t . It does this by sometimes observing a system before t and predicting, via its knowledge of the deterministic Newtonian laws, what that system's state will be at t , sometimes observing a system after t and retrodicting what its state was at t (again via its knowledge of the reverse-deterministic laws), and sometimes by observing systems nearby the target system and using the laws to figure out what the target system is like at t , given the states of its nearby systems. Eventually, at some time after t , it comes to know all of the particular matters of fact about the total state at t . Still, as per Ismael's argument, it does not know the totality fact, namely that its knowledge of the state at t is complete. In other words, it knows the total state, but it does not know that this is the total state.

Suppose, however, that it just guesses that indeed its knowledge of the state at t is complete. This may be an educated guess: perhaps it sees that if any other forces were present, they would be influencing the behaviors of the systems it already observed in ways that it probably would have noticed. But it does not have to be an educated guess; perhaps the intelligence finally gets tired of gathering further information about particular matters of fact and thinks to itself, "I'm just going to hope that's all of them."

For all that's been said so far, this guess could be correct. If t is in the past when the intelligence makes its guess, all of the facts about the state at t have already obtained, so they are there to be guessed correctly or incorrectly. Suppose that the intelligence's guess, namely that "I know all of the information about the total state at t ," is correct. And suppose that the intelligence plugs *what just so happens to be the total state* into the deterministic Newtonian laws, and evolves that state forward to figure out the output of a counterpredictive device in its near future.²¹

The prediction then ought to be correct. But unless the counterpredictive device malfunctions, it will falsify the prediction. So something, again, has gone wrong. The Nomic Preclusion argument therefore does not rid us of the paradox. The problem is that the Nomic Preclusion argument explains the impossibility of too narrow a range of circumstances. If it works, it shows why an embedded agent cannot *know* the total state information required to predict the output of the counterpredictive device. But it does not explain why an embedded agent cannot *guess* that information, or believe it with reasonable-but-not-definitive justification. And these alternative epistemic states would work just as well for generating the paradox; what matters is not that the intelligence knows the requisite information, but merely that it somehow, by hook or by crook, acquires it.²²

²¹"Near" future, to guarantee that the counterpredictive device is within the predictive horizon of the intelligence, taking into account measurement errors and calculational shortcuts that it must take as an embedded agent.

²²To be clear, it is perfectly obvious why the intelligence cannot correctly guess *the output*

The paradox is thus more pernicious than has so far been appreciated. The intelligence does not have to be a full-fledged Laplacean intelligence in order for the paradox to arise. Indeed it does not have to know the laws, nor the total state, nor does it have to be a perfect calculator. All that matters for the paradox is that the intelligence *as a matter of fact* just so happens to possess correct information about the laws and the total state, and that as a matter of fact it calculates correctly this time. Perhaps it comes across its correct information about the laws and the total state by guessing. If the laws and the total state are already-existing elements of the universe when it goes to make its guesses, then it may indeed guess correctly. It does not have to know that its guess is correct; what matters is only that its guess *is*, in fact, correct.

It may be tempting to retreat to brute logical impossibility here. If the intelligence correctly guesses the total state and computes the output of the counterpredictive device, then the counterpredictive device must malfunction. If the device does not malfunction, the intelligence cannot correctly guess the total state. There is no possible world in which the intelligence guesses correctly and yet the counterpredictive device works correctly, for then we would have a logical contradiction. So there is really no mystery here.

But if this is an adequate solution to the paradox, we might as well have stopped much earlier. Remember, we knew from the get-go that a Laplacean intelligence and a counterpredictive device are not compossible: there are no possible worlds where the one encounters the other in the prescribed manner. What we wanted, what we've been looking for, is an *explanation* of this impossibility.

Ismael's Nomic Preclusion argument was supposed to provide such an explanation: an embedded intelligence is nomically precluded from having knowledge of the information required for its prediction. In other words, the Nomic Preclusion argument attempts to explain the impossibility of the coexistence of a Laplacean intelligence and a counterpredictive device by showing that (embedded) Laplacean intelligences themselves are impossible. Its goal is to explain an impossibility of coexistence that is otherwise mysterious.

But as I have just argued, there is another impossibility of coexistence lurking here that is just as mysterious, namely the impossibility of (i) an intelligence that correctly guesses the total state and the laws and (ii) a counterpredictive device that does not malfunction. This impossibility is not explained by the Nomic Preclusion argument, but if we wanted an explanation in the first case, we should also want one here. Yes, there is no possible world where the one encounters the other. But each of the ingredients *by itself* is perfectly possible. In the absence of the counterpredictive device, the intelligence might well guess correctly; and in the absence of the intelligence, the counterpredictive device might be perfectly reliable. What then prevents the intelligence from guessing

of the device (assuming that the device is given the prediction, does not malfunction, etc.). That is not what's happening here. Rather, the intelligence is guessing the total state at t , and then *calculating* the device's output based on the laws and its guess about the total state. What we want an explanation of is why, in that situation, the intelligence would be guaranteed to guess wrong.

correctly when we suddenly introduce a reliable counterpredictive device into its future? Again it may be helpful to adopt a first-person perspective here. Imagine yourself in the position of the intelligence: what is stopping you from guessing correctly? Some event in the future that hasn't happened yet? One almost gets the sense that some sort of backward causation is at work here. It is right to find this mysterious. It would be good to have an explanation of your guaranteed failure. And the Nomic Preclusion argument does not provide one.

6 Guessing Laws and Melodies

It is worth noting that this “guessing” gambit also works against many of the solutions discussed in §2. For example, even though an embedded intelligence might not be able to perfectly measure the values of all of the physical parameters in the total state, it can still guess them, and it may happen to guess correctly. Similarly, even if the intelligence could not compute the nomic function, it could still guess the output of that function, and happen to get it right. And even if it could not know the laws, it could still guess them and turn out to get lucky. As long as there are matters of fact about the laws and antecedent conditions that entail the future output of the counterpredictive device, and as long as those facts are available at the time that the intelligence goes to make its guess, it seems that the intelligence may correctly guess those facts and thereby be in a position to perfectly predict the outcome of the device.

This way of putting it, however, shows why Humeanism about laws still gives us a way out. According to Humeanism—coupled with some variety of presentism—the laws are not set in stone at the time that the intelligence goes to make its guesses. It can guess the laws and the total state, and it can calculate their consequences, but if there is a counterpredictive device in its future it must be wrong about one of these. The total state has already occurred and is there to be guessed, so suppose that the intelligence has guessed the total state correctly, and suppose it has calculated correctly. Then it must be wrong in its guess about the laws.

Why? The story here is essentially the same as we saw in §3; it works just as well when we replace *knowing* with *guessing*. The intelligence must guess wrong because of what laws are, namely summaries of what happens. Given that metaphysics of laws, a counterpredictive device is a mechanism for rigging the summary—i.e. the laws—to ensure that any guess about the laws that was used to generate a prediction fed into the device must have been wrong.

Humeanism thus shows us why the intelligence cannot even guess the laws correctly if there is a counterpredictive device in its future. In such situations, the laws are unguessable, because the guess is part of what determines the laws. It would be like asking someone to guess what melody is being played on a piano by playing a note, on the piano, to indicate their guess. If the correspondence between notes and guesses is rigged, the note itself will change the melody being played, and the guess is guaranteed to be wrong. Even a highly experienced musician who would otherwise be a flawless melody guesser couldn't get it right

here.

7 Is an Explanation Necessary?

Humeanism about laws is sometimes regarded as a radical view. And if the successful explanation also requires presentism, then we will have to confront the fact that presentism itself looks to be incompatible with at least some of our best physical theories (Putnam, 1967). Moreover, the prospects of combining these two views are at best underexplored. So finding ourselves pushed toward a presentist Humeanism in response to the paradox of predictability might reasonably prompt a re-evaluation of how we got to this point. On reflection, there is one place in particular that we might scrutinize more carefully, and that is our insistence that we really ought to have an explanation of the impossibility here. Perhaps it actually wouldn't be so bad to just rest content knowing that there is no possible world where a Laplacean intelligence encounters a non-malfunctioning counterpredictive device. It was in seeking an explanation of this fact that we were led to Humeanism. But do we really need an explanation? Is it not enough to point out the impossibility and leave it at that?

Again, I do not think we *need* an explanation, but that we should prefer to have one if it is available—just as mathematicians prefer to have proofs that are explanatory as opposed to proofs that merely demonstrate the truth of the target proposition. In the case at hand, Humeanism gives us an explanation, whereas anti-Humeanism does not. This is a *prima facie* reason to prefer Humeanism.

The aim of this section is to further motivate the request for an explanation. To that end, first consider an analogy. What happens when an unstoppable force meets an immovable object? Well, nothing happens. There is no possible world where those two things meet, and we know this without delving any deeper into the relevant concepts. But what is important in this case is not so much the fact of the impossibility, rather it's the explanation of the impossibility. Our pretheoretic ways of thinking about forces suggest that an unstoppable force should be possible, and also that an immovable object should be possible. Why then can those things not coexist? The insight comes when we learn that our pretheoretic way of thinking about forces is inadequate. Forces are vector quantities with attendant numerical magnitudes. There is no such thing as an “unstoppable force” because there is no greatest numerical magnitude that a force can have. And there is no such thing as an immovable object, because an object's movability is just a function of the forces exerted upon it.

The discovery that we have to modify our naïve conception of forces is an important insight, but we could not reach it just by resting content with the impossibility of the coexistence of an unstoppable force and an immovable object. It is only in seeking an explanation of that impossibility that we make this discovery. Analogously, merely resting content with the impossibility of a Laplacean intelligence coexisting with a counterpredictive device may not be enough. It is only in seeking to explain the impossibility that we achieve any real insight.

Another way to highlight the benefits of an explanation is with a thought experiment. Imagine that Avi is a would-be Laplacean intelligence. She has studied a great deal about the universe and has come up with an educated guess about the total state at some prior time t as well as the deterministic laws of nature. She has some confidence that this guess is right, but she is not positive. Moreover, she has no intentions of calculating all of the future consequences of that state using the laws.

Avi's colleague Trey works just across the hall. Trey has long been envious of Avi's success: she got tenure before him, has a better publishing record, and consistently gets better teaching evaluations. When Trey now learns of Avi's aspirations, he realizes that he could frustrate them by doing the following: (1) constructing a counterpredictive device, (2) calculating the future consequences of her guess, and (3) feeding the prediction generated thereby, about the action of the counterpredictive device, into the device itself. Suppose he plans to do this, and reveals these plans to Avi. What should happen to Avi's confidence in her guess about the laws and total state?

It seems that it ought to plummet. For if Trey succeeds in constructing a counterpredictive device, and if he calculates all the future consequences of her guess, he will obtain a prediction about the output of the device. If he then feeds that prediction into the device, the prediction is guaranteed to be wrong. And that implies, assuming that Trey calculated correctly, that Avi's initial guess had to be wrong. Moreover, her guess is guaranteed to be wrong *regardless of whether Avi modifies it in response to learning about Trey's intentions*, and also regardless if *how* she modifies it.

Suppose Trey does construct the device, and informs Avi of this. Avi should now be even less confident in her guess about the laws and total state. Of course, there is still a chance that Trey will not calculate a prediction on the basis of Avi's guess and feed that prediction into his device, so it is not yet guaranteed that her guess is wrong. Indeed, we might imagine that Trey's envy is now fading (he recently had a run of luck at the journals, and also just received tenure), so he is vacillating about whether he will follow through with his plan. If Avi wants her guess to be right, she then has reason to try to talk Trey out of it. In fact she would have reason to be as nice as possible to him: flatter him, bake him cookies, and so on.

We could extend this thought experiment further, perhaps imagining Avi publishing her guess in a journal, which then receives a multitude of "falsification" reports from people looking to get a quick publication.²³ But the point should already be clear. From the perspective of an anti-Humean, this situation is tremendously bizarre. From such a perspective, Avi's guess concerns matters of fact that have already obtained. It should then strike us as puzzling *in the extreme* that her guess cannot be correct, and that this is so regardless of how she modifies it. It should also be puzzling how Trey's intentions with the counterpredictive device have any evidential relation to her guess about the total

²³Given the amount of calculation involved, it wouldn't actually be that quick. Rather, it would be the most laborious publication, in terms of total work required, in the history of humankind.

state and the laws, for Trey’s intentions are not in any obvious way *about* the total state or the laws. The only connection is that if he succeeds, Avi’s guess is logically guaranteed to be wrong.

This puzzlingness motivates the search for an explanation. It hardly seems sufficient here to rest content with knowledge of the logical impossibility that Avi guesses correctly and Trey feeds a prediction extracted from that guess into a functioning counterpredictive device. If you do not agree, try putting yourself into the position of Avi, and ask yourself: “Why can’t I guess right?” It’s as if your guess is somehow affecting the things that you are guessing about, but from the anti-Humean perspective there is no way your guess *could* affect such things.

By contrast, there is nothing mysterious here from the point of view of Humeanism. Avi cannot guess right about the laws and total state because her guess helps to determine part of what is being guessed about, namely the laws of nature.

8 How Humean Must We Be?

If the Humean resolution of the paradox is correct, there is still a question of exactly how Humean we must be about the laws. There are two ways we might be less-than-fully Humean that are worth considering here: first, we might try being Humean about the global laws while being anti-Humean about the local ones, and second, we might try being Humean about some global laws while being anti-Humean about others.

The distinction between global laws and local laws is due to Ismael (2016, pp. 178-179). The global laws of temporal evolution, or *GLOEs*, link the present state of the entire universe with its future states. The local laws, on the other hand, govern how the matter in a localized region of spacetime interacts with its environment. I say “govern” here because Ismael’s suggestion is that while we might be Humean about the GLOEs, we could still maintain some form of anti-Humeanism about the local laws. The idea is that the on-the-ground facts at a particular time would include the anti-Humean local laws but not the Humean GLOEs, which are a function of the total history of the universe. The local laws and initial conditions, Ismael observes, do not themselves entail the GLOEs without the addition of a totality fact. So even if we imagine an embedded intelligence that knows the total state and the local laws at some instant, it would not know the totality fact required to entail the future behavior of a counterpredictive device. Of course, if it knew the GLOEs, then it would know enough to entail the behavior of the device, but since the GLOEs are Humean and therefore depend on a totality fact that the intelligence cannot know, it cannot know the GLOEs either.

The problem here is that, as before, even if the intelligence cannot know the requisite totality fact, it may still guess it correctly. And if it guesses correctly about the total state, and if it also knows all the local laws governing how matter in localized regions of spacetime evolves and interacts with matter in nearby

regions, then it possesses enough information—whether it knows it or not—to entail the behavior of a counterpredictive device in its future. Put differently, if we allow that the intelligence can know the local laws, then it might guess the total state correctly, in which case it effectively has all of the information contained in the GLOEs, even though it does not know that it does. With that information in hand, it should be able to accurately predict the behavior of a future counterpredictive device, and the paradox re-arises.

So it does not seem to me that we can be Humean about the global laws while being anti-Humean about the local laws. Such a position allows too many facts to be set in stone at the time that the intelligence goes to make its prediction to prevent the paradox from arising.

Suppose instead that we maintain Humeanism about some of the global laws, and anti-Humeanism about others. More specifically, imagine a view like Maudlin's, where some of the laws of temporal evolution are fundamental elements of the universe driving it toward the future. However, not all of the laws are like that. Rather, some of the laws of temporal evolution are Humean: they are determined by how the entire history of the universe plays out. The anti-Humean laws, here, constrain that history, but they do not fully determine it by themselves.

This “mixed” view could avoid the paradox of predictability. As long as the anti-Humean laws do not fully determine the output of the counterpredictive device, an embedded intelligence in the past of that device would be guaranteed to be wrong about the extra *Humean* laws that are needed to entail the output of the device, for the same reasons as we saw above.

This view is somewhat reminiscent of Cartwright's (1980) position, where she suggests that there might not be “superlaws” in God's book of laws that cover every situation that ever arises. If we think of these superlaws as Maudlin-style laws of temporal evolution, and we stipulate that the remaining situations not covered in God's book are instead covered by Humean laws, then we have essentially the view described here.

Admittedly, this is a rather odd and disjointed view. Aside from the fact that it avoids the paradox of predictability and preserves some anti-Humean necessary connections, it is hard to see what motivation there could be for conjoining anti-Humeanism about some laws with Humeanism about others. Perhaps, however, there is some motivation for this position that I am not seeing. At any rate, the upshot is that avoiding the paradox requires only a partial commitment to Humeanism. We then have to rely on considerations of parsimony and theoretical elegance to take us the rest of the way: once you are committed to the existence of some Humean laws, it is far more elegant to allow that all laws are Humean than to steadfastly cling to some anti-Humean laws as well.

9 Conclusion

Philosophers who have worked on the paradox of predictability have typically assumed that what is needed is an explanation for why an embedded agent lacks a reliable method for figuring out the total state and laws of nature. This is a reasonable assumption, but I have argued that it is false: the paradox persists even if we allow that the intelligence merely guesses the requisite information. In hindsight, this should perhaps not be too surprising, for if the paradox hinged on the *knowledge* possessed by the intelligence, then the appropriate analysis of “*S* knows that *P*” would become relevant, as indeed would arguments for external world skepticism and the like. But it seems clear that such arguments are not relevant. The paradox arises for the Cartesian skeptic just as well as it arises for common-sense Moorean. Arguments about the nature of knowledge are red herrings here.

Once we recognize that the paradox does not hinge on the strength of the intelligence’s justification for their beliefs, it becomes clear that if an explanation is to be had here, it must be more flexible, able to explain even the impossibility of *guessing* correctly. Humeanism about laws of nature is the natural remedy. If we should prefer to have an explanation here, then we should prefer Humeanism.

References

- [1] Armstrong, David. (1983) *What is a Law of Nature?* Cambridge: Cambridge University Press.
- [2] Beebe, Helen. (2000) "The Non-Governing Conception of Laws of Nature." *Philosophy and Phenomenological Research*. 61: 571-594.
- [3] Beebe, Helen and Alfred Mele. (2002) "Humean Compatibilism." *Mind*. 111: 201-223.
- [4] Bird, Alexander. (2007) *Nature's Metaphysics*. New York: Oxford University Press.
- [5] Bishop, Robert C. (2003) "On Separating Predictability and Determinism." *Erkenntnis*. 58: 169-188.
- [6] Boyd, Richard. (1972) "Determinism, Laws, and Predictability in Principle." *Philosophy of Science*. 39: 431-450.
- [7] Callender, Craig. (2017) *What Makes Time Special?* New York: Oxford University Press.
- [8] Carroll, John. (1994) *Laws of Nature*. New York: Cambridge University Press.
- [9] Cartwright, Nancy. (1980) "Do the Laws of Physics State the Facts?" *Pacific Philosophical Quarterly*. 61: 75-84.
- [10] Cohen, Jonathan and Craig Callender. (2009) "A Better Best System Account of Lawhood." *Philosophical Studies*. 145: 1-34.
- [11] Demarest, Heather. (2017) "Powerful Properties, Powerless Laws." In Jacobs, Jonathan (Ed.), *Putting Powers to Work: Causal Powers in Contemporary Metaphysics*. Oxford: Oxford University Press.
- [12] Dorst, Chris. (2018) "Towards a Best Predictive System Account of Laws of Nature." *British Journal for the Philosophy of Science*. 70: 877-900.
- [13] Dretske, Fred. (1977) "Laws of Nature." *Philosophy of Science*. 44: 248-268.
- [14] Ellis, Brian. (2001) *Scientific Essentialism*. Cambridge, UK: Cambridge University Press.
- [15] Garrett, Brian and Jeremiah Joven Joaquin. (2021) "Ismael on the Paradox of Predictability." *Philosophia*. <https://doi.org/10.1007/s11406-021-00381-z>.
- [16] Geroch, Robert. (1977) "Prediction in General Relativity." In Earman, J., Glymour, C., and Stachel, J. (eds.) *Foundations of Space-Time Theories*, Vol. 8, pp. 81-93. Minneapolis: University of Minnesota Press.

- [17] Gijssbers, Victor. (2021) “The Paradox of Predictability.” *Erkenntnis*. <https://doi.org/10.1007/s10670-020-00369-3>.
- [18] Hall, Ned. (ms) “Humean Reductionism about Laws of Nature.” <https://philpapers.org/archive/HALHRA.pdf>
- [19] Hicks, Michael. (2018) “Dynamic Humeanism.” *British Journal for the Philosophy of Science*. 69: 987-1007.
- [20] Ismael, Jenann. (2016) *How Physics Makes Us Free*. New York: Oxford University Press.
- [21] Ismael, Jenann. (2019) “Determinism, Counterpredictive Devices, and the Impossibility of Laplacean Intelligences.” *The Monist*. 102: 478-498.
- [22] Jaag, Siegfried and Christian Loew. (forthcoming) “Making Best Systems Best for Us.” *Synthese*.
- [23] Landsberg, P. T. and D. A. Evans. (1970) “Free Will in a Mechanistic Universe?” *The British Journal for the Philosophy of Science*. 21: 343-358.
- [24] Lange, Marc. (2016) *Because without Cause: Non-causal Explanations in Science and Mathematics*. New York: Oxford University Press.
- [25] Lewis, David. (1973) *Counterfactuals*. Malden: Blackwell.
- [26] Lewis, David. (1986) *Philosophical Papers, Volume II*. New York: Oxford University Press.
- [27] Lewis, David. (1994) “Humean Supervenience Debugged.” *Mind*. 103: 474-90.
- [28] Loewer, Barry. (1996) “Humean Supervenience.” *Philosophical Topics*. 24: 101-127.
- [29] Loewer, Barry. (2007) “Laws and Natural Properties.” *Philosophical Topics*. 35: 313-328.
- [30] Loewer, Barry. (2020) “The Package Deal Account of Laws and Properties.” *Synthese*. <https://doi.org/10.1007/s11229-020-02765-2>
- [31] Mackay, D. M. (1960) “On the Logical Indeterminacy of a Free Choice.” *Mind*. 69: 31-40.
- [32] Malament, David. (1996) “In Defense of Dogma: Why There Cannot be a Relativistic Quantum Mechanics of (Localizable) Particles.” In Rob Clifton (ed.), *Perspectives on Quantum Reality*. Dordrecht: Springer.
- [33] Manchak, John. (2008) “Is Prediction Possible in General Relativity?” *Foundations of Physics*. 38: 317-21.

- [34] Maudlin, Tim. (2007) *The Metaphysics within Physics*. Oxford: Oxford University Press.
- [35] Maudlin, Tim. (2011) *Quantum Non-Locality and Relativity, Third Edition*. Malden: Wiley-Blackwell.
- [36] Mumford, Stephen. (2004) *Laws in Nature*. London: Routledge.
- [37] Norton, John. (2008) "The Dome: An Unexpectedly Simple Failure of Determinism." *Philosophy of Science*. 75: 786-798.
- [38] Putnam, Hilary. (1967) "Time and Physical Geometry." *The Journal of Philosophy*. 64: 240-247.
- [39] Roberts, John. (2008) *The Law-Governed Universe*. New York: Oxford University Press.
- [40] Rummens, Stefan and Stefaan E. Cuypers. (2010) "Determinism and the Paradox of Predictability." *Erkenntnis*. 72: 233-249.
- [41] Scriven, Michael. (1965) "An Essential Unpredictability in Human Behaviour." In B. B. Wolman and E. Nagel (Eds.) *Scientific Psychology: Principles and Approaches* (pp. 411-425). New York: Basic Books.
- [42] Shoemaker, Sydney. (1980) "Causality and Properties." In Peter van Inwagen (ed.), *Time and Cause*, Dordrecht: Reidel, pp. 109-35.
- [43] Stone, M. A. (1989) "Chaos, Prediction, and Laplacean Determinism." *American Philosophical Quarterly*. 26: 123-131.
- [44] Tooley, Michael. (1977) "The Nature of Laws." *Canadian Journal of Philosophy*. 7: 667-698.
- [45] van Inwagen, Peter. (1983) *An Essay on Free Will*. Oxford: Clarendon Press.