

2

Humean Disillusion

Jenann Ismael

C2S1

2.1 Introduction: the Case for Humeanism

C2P1

In philosophy, there tend to be different stages in a debate. One starts with a question or a problem (e.g., what are the laws of nature? Are there objective ethical truths? What is the most just form of society?). In initial stages as positions are being developed, arguments are properly addressed from one side of the debate to the other as objections are used to help refine positions. Once the debate has matured and we have a number of well-developed positions (i.e., where the choice is no longer so much one of consistency or coherence, where people agree that the positions are coherent and it is really a matter of drawing out implications), then I think arguments should be addressed not so much to those on the other side of the debate (the standard should not be ‘will you convince your opponents?’) but to those that have not made up their minds. This is the stage where we have well-defined and well-understood positions, and things have hardened into different viewpoints that typically organize a whole cluster of commitments. That is the phase that the debate over Humeanism has reached.

C2P2

In its early years, the initially attractive aspects of Humeanism were in the foreground. It assumed the existence only of local matters of particular fact. And because it had an epistemology that depended only on the knowledge of such facts, it was tailor-made for the empiricist instincts of science-based metaphysics. The methodology, moreover, was supposed to recapitulate the methods that scientists invoked in choosing between theories. The idea was that science gathered a large and wide-ranging body of information about local matters of particular fact and systematized that body of fact using the methods that scientists actually use. The laws and chances were statements that appeared in a certain role in the systematization. The account provided a reductive, non-metaphysical account of laws and chances that captured the main insights of the simpler regularity and frequency-based accounts of yore, but because it united laws and chances into a single package and allowed systematization to operate on the package, it avoided simple counterexamples to those accounts (Ismael 2015). It was an absolute breath of fresh air for those who wanted a science-based metaphysics, i.e., who wanted to believe in the modal commitments of science without scholastic metaphysics. There were no relations among universals, no irreducible modal forces or

anything added to the Humean mosaic to enforce laws. It was all about systematizing bodies of local matters of particular fact in a way that was itself modelled on science.

C2P3 The epistemology was particularly important for David Lewis, who introduced the Humean view into the literature, and for other empiricists. The non-reductive accounts of what laws and probability are that were on offer in the philosophical literature separated them from the local matters of particular fact that provided evidence for them, rendering them unknowable. Lewis's account had its own objectionable aspects, but those were gradually shed. Later Humeans discarded the 'natural properties' that Lewis appealed to (Loewer 1996), and others developed the account to apply to special science laws (Cohen and Callender 2009). Both of these reforms brought it into closer alignment with science. Some aspects of the account remained schematic; criteria of simplicity, strength, and fit proved difficult to characterize explicitly, but that was because they were meant to designate scientific criteria for choosing among systems that were themselves difficult to characterize. Since its beginnings, moreover, Humeanism has had the advantage of exceptionally eloquent and charismatic defenders: Lewis himself, of course, and since then Ned Hall, Barry Loewer, and David Albert. For the scientifically minded metaphysician or philosopher of science, for a while it seemed there was simply no other game in town.¹

C2S2

2.2 Disillusion

C2P4 What has happened in the years since, however, is that Humeanism's faults have begun to surface. For at least some Humeans, the shoe has begun to pinch. It's a familiar story: you leave an old job or an old lover for a new one with none of its faults, and then over time its own faults begin to surface with increasing clarity.

C2P5 I'll speak of Humeanism in the tradition stemming from Lewis, through Albert and Loewer, and argue that the Humean doesn't have an epistemology that makes sense for embedded agents. The indefinite extendibility of the Humean mosaic, the fact that chances are determined by the pattern over the *whole*, and the explicit commitment to a combinatorial principle for determining what patterns are possible mean that there is no way for a Humean agent to use information about local matters of particular fact to arrive at beliefs about the chances, short of assuming a restriction on priors that is patently at odds with their own metaphysics.

C2P6 Here's how the discussion will go: first, I'll run through the problem that Lewis himself articulated in the article that launched the philosophical literature on Humean chances: the Undermining Problem. The Undermining Problem was

¹ Maudlin has been a vocal and persistent critic of Humeanism and you will see some affinities with his position here. See Maudlin (2007).

that chances contain information from the future and so they are incompatible with certain ways that the future could go. That meant that to assert a theory of chance is to rule out certain ways the future could go, but if we use the chances themselves to guide credence about those futures, they will generally assign a non-zero probability, so we have a contradiction. Then I'll introduce the fix. Undermining teaches us that we need to temper knowledge of the chances to correct for ignorance of the future. If we do that, I will show that the contradiction disappears in a natural way, but now a new problem emerges: no matter how much information you have about local matters of particular past fact, an indefinite portion of the chance-making pattern lies in the future. That means that the idea that observation gives us information about the chances at all was mistaken. Without a boundary condition in the size of the mosaic, you never get any closer to knowing what the chances are.

C2P7 The argument is simple: a combinatorial principle for which Humean mosaics are possible, the claim that laws and chances are determined by a global criterion applied to the mosaic, and the recognition that the mosaic is indefinitely extendible together mean that conditionalizing on local matters of particular fact brings us no closer to knowing what the chances are. If you are inclined to think that it is not in general a problem if your metaphysics of X makes X unknowable to embedded agents, Lewis himself recognized that it is devastating to an account of chance, because the role chances play is to guide belief in the face of ignorance about the future. I'll discuss a response suggested (in conversation) by Albert and Loewer that I argue doesn't work.

C2S3

2.3 The Undermining Problem

C2P8 The problem of undermining and the immense amount of energy and work that went into sorting it out brought a great deal of analytical clarity into discussion of the relationship between a base ontology of local facts and a vocabulary that might contain disguised information about distributed patterns in it.

C2P9 Lewis started out by asking what chances could be, and he introduced the Principal Principle (PP) as an implicit definition of chance that identified chances by the role they play guiding belief. What the Principle said in its original formulation was that one should adjust one's credence to the chances no matter what other information one has, except in the presence of inadmissible information:

C2P10 PP: $cr(A/\langle ch_t(A) = x \rangle E) = x$, provided that E is admissible with respect to $\langle ch_t(A) = x \rangle$

C2P11 Where $cr(A)$ is one's credence in A at some time t and $ch_t(A)$ is the chance of A at t. The restriction to admissible information was needed to discount cases where

PP clearly becomes inapplicable; e.g., when one possesses information from the future of the sort one might get from a crystal ball or a privileged communication from God.²

C2P12 The problem that Lewis noticed that this poses for a Humean account of chance has to do with the possibility of what he called ‘undermining futures’.

C2P13 Undermining futures are futures that are incompatible with the chances being what they are. We know that such futures have to exist on a Humean account of chance to the extent that the correct theory of chance depends on how the future goes. So long as there is some dependence of what the chances are on how the future goes, no matter how small, and so long as a theory of chance assigns some non-zero probability to any future that is nomologically compatible with the past, that was enough to show that there would be some futures assigned a non-zero probability by the chances and yet that were metaphysically incompatible with the chances being what they were.

C2P14 The problem can be put in a nutshell. It is that, on the one hand, chances are those things that play the role of chance in PP, so whatever we assign as reference, it had better be able to play that role. But, on the other hand, no Humean truthmaker *could* play that role because any Humean truthmaker introduces inadmissible information and undermines the applicability of PP.

C2P15 That is the central difficulty of a Humean account. If laws and chances are to be identified with (some function of) the whole pattern of facts, then beliefs about the laws and chances have to be tempered by ignorance of the facts. If presumed knowledge of the chances outruns knowledge of the facts, then we aren’t going to have advance knowledge of them in a way that allows us to use them to guide credence about the future. The reason that seems like disaster is that chances are *there* to serve the epistemic purpose of guiding belief in the face of ignorance of the future.

C2P16 This is how Lewis put the problem, looking back in “Humean Supervenience Debugged”:

C2P17 If I’d seen more clearly, I could have put the core of my reduction like this. According to the best-system analysis, information about present chances is inadmissible, because it reveals future history. But this information is not inadmissible, as witness the way it figures in everyday reasoning about chance and credence.³

² The formula that Lewis wrote down was more complicated than what he said in prose. I’m sticking with the prose formulation, which comes closest to capturing the pre-theoretic idea of how information about chance guides belief. If you prefer the more complex principle, substitute it here.

³ Lewis (1994, p. 486). Ned Hall and Mike Thau were instrumental in getting him to see it this way.

- C2P18 The insight that Lewis took from Hall and Thau is that (i) it is true that if what the chances are depended too sensitively on some particular event *e*, then we can't use beliefs about the chance of *e* to guide expectation about *e* (so, for example, if which theory of chance were true depended sensitively on the outcome of the particular toss of a given coin, then we couldn't use the chances to guide belief about that toss), but (ii) the correct theory of chances don't depend sensitively on the particular events that we typically use them to guide expectation about.
- C2P19 Consider a world consisting of nothing else but a sequence of tosses of a single coin. The chances supervene on *total* histories: the bearing of a single toss on the chances ought to be proportional to the length of the sequence. In a world the length of our world, a single toss isn't going to shift probabilities for which theory of chance is true in any appreciable way, so we can use the chances to guide belief about the outcomes of particular tosses.
- C2P20 And that was enough to reconcile the everyday role of chance with the Humean commitment to the existence of undermining futures. The quantitative disparity between the information contained in a theory of chance about some particular event and the information that that event carries about which theory of chance is correct means that we can ignore the latter. (An analogy: although in principle every object exerts gravitational attraction on every other, the gravitational influence the earth exerts on a piece of dust so outweighs the effect the dust exerts on the earth that we can ignore the latter in calculating the former.) We can get a quantitative measure of that degree of dependence that we might think of as a measure of the degree of inadmissibility of information about chance and we reformulate the PP to reflect that it is applicable only when the degree of inadmissible information is low. It is low in the case of coin tosses and everyday events, making chance information admissible, and high in the case of undermining futures, making it inadmissible. Contradiction dispelled, and consistency restored.
- C2P21 And what we learn from all of this is that the reason that it pays for creatures like us (i.e., creatures that have information about the past, but whose information about the future is always derivative of what they know about the past) to think about laws and chances has to do with the 'balance of information'. The *balance* of information is such that if we have a large enough body of information about the past, we can use that to stabilize beliefs about which theory of chance is correct, and then use our theory of chance to guide credence about everyday events. So even though in *principle* which theory of chance is correct depends on *every* event, *in fact* our theories are largely indifferent to (cannot be undermined by) beliefs about the particular events we want to use them to guide belief about. If the world was too simple, or our actions and the events we are interested in predicting are not highly localized relative to the chance and law-making patterns in the Humean mosaic, it wouldn't work.
- C2P22 All of that seemed to make good sense and recover confidence in the Humean account of what chances are. A very big part of what recommends that account is

that it gives us this very natural epistemology. The account of what the chances are connects them both to what we count as evidence for them, and also to what they guide belief about.

C2S4

2.4 Ignorance about Chances

C2P23 Undermining taught us that beliefs about what the correct theory of chance is, and hence about *what the current chances are*, on a Humean account, are hostage to the outcome of future observations so that any application of PP has to be tempered to reflect our ignorance about what the chances are. That means that the simple PP, which tells us what to do if we know what the chances are, needs to be supplemented with a principle that says how to let beliefs about chance guide credence where one *doesn't* know what the chances are. On this way of understanding it, undermining simply makes explicit the need for some principle about how to form credences in the absence of certainty about the chances.

C2P24 In principle, you could write down any number of these. One could, for example, adopt as credences the chances assigned by the theory of chance assigned highest credence, or divide credence evenly among the top four, or . . . you can think of any number of them.⁴ The simplest suggestion is that one should create a weighted mixture of all theories of chance metaphysically compatible with history so far.

C2P25 GPP: $cr(A) := \sum cr(C_{chi})ch_i(A)$, where ch_i is the chance assigned to A by epistemically possible theory of chance ch_i .⁵

C2P26 Where should the weights come from? They should reflect current credences in the theories of chance in question and ultimately have to be rooted in priors. One can interpret the priors in a Bayesian way or impose additional constraints. But we want to impose the requirement that you can't assign a zero probability to any of the theories of chance that is metaphysically compatible with history so far. If you respect that requirement, you are not going to be in the problematic position—i.e., of plunking for a theory of chance that is metaphysically incompatible with the future going a certain way, and accepting the theory's recommendation to assigning it a non-zero credence—because you are never

⁴ Not all of these are going to behave well diachronically under all conditions; see Pettigrew (2012). But think about what these kinds of arguments show: that there are some ways that the Humean mosaic could turn out to be where such a principle would lead to reasoning that would lead to non-Bayesian conclusions. But the Humean can respond here as below by saying that the success of our methods rests on the mosaic not being 'pathological' in certain ways. If that response is legitimate below in conjunction with induction, it is legitimate here.

⁵ This is the one I proposed (Ismael 2008). I've suppressed the temporal parameters to avoid making the expression unpleasantly wieldy. I now prefer Pettigrew's Aggregate Principle: see Pettigrew (2016).

going to be plunking for a single theory of chance at all. At any given moment in any given history, there is going to be lots of different ways the future could go, and you are going to be mixing chances drawn from theories of chances corresponding to different histories. (That's true even if the world is in fact finite. If the world ends tomorrow, the credences I form on the eve of destruction are going to incorporate chances from lots of theories corresponding to histories in which it continues.) So we are golden.

C2P27 This was a way of recognizing that if theories of chance supervene on *total* histories, we should be as ignorant of which theory of chance is correct as we are of how the future will go. And it seemed to me that it was a point in favor of Humeanism that as soon as you recognize this very natural idea, undermining problems go away. All of this is just getting increasingly precise and explicit about exactly what using information about local matters of particular fact to update beliefs about the laws looks like in a Humean world, assuming no non-Humean necessary connections. Since the laws supervene on the whole mosaic and specifically encode information about global features, the epistemology turns into a matter of using information about local matters of particular fact to form beliefs about global properties of a four-dimensional manifold of such fact.⁶

C2P28 One of the primary arguments *for* Humeanism was that it gave us the connection between chances, on the one hand, and the local matters of particular fact that (i) provide evidence for them and (ii) they guide belief about, on the other. Lewis was quite explicit that whatever analysis one gave of what chances are, it ought to make sense that they should guide credence in accord with PP. So, for example, if the chance of *e* was the degree to which Angela Merkel preferred *e* to occur, it wouldn't make sense for the chance of *e* to guide your credence in it. The complaint he made about non-Humean theories was that they fail this test.

C2P29 Be my guest—posit all the primitive unHumean whatnots you like. [...] But play fair in naming your whatnots. Don't call any alleged feature of reality "chance" unless you've already shown that you have something, knowledge of which could constrain rational credence. [...] I don't begin to see, for instance, how knowledge that two universals stand in a certain special relation N^*N^* could constrain rational credence about the future coinstantiation of those universals.
(Lewis 1994, p. 484)

C2P30 For Lewis and for many of us who liked Humeanism, it was just as much because the epistemology was so clean and straightforward. Patterns in the Humean

⁶ This generalizes to the deterministic case where chances become 1 or 0. In that case, we are using information about local matters of particular fact to guess the laws and then using the laws (in conjunction with the past) to derive predictions for the future.

mosaic were exactly the kinds of thing local matters of particular fact could give one information about. It made perfect sense not only that chances should be what guide credence in the face of uncertainty about local matters of particular fact but that *evidence for what the chances were* would come from information about such things.

C2S5

2.5 The New Problem

C2P31 But now a new problem comes into focus. Undermining showed us that we have to temper our knowledge of the chances to reflect our ignorance of the future. When we correct for our ignorance of the future by tempering our beliefs in which theory of chance is correct, the chances themselves are unavailable for the purposes for which they were designed. We cripple Humean chances for the purposes of guiding belief in the face of ignorance.

C2P32 We can see this by thinking about how the Humean epistemology would work in a bounded universe of definite size. We are going through the world, picking up information about local matters of particular fact. Each observation brings us incrementally closer to knowledge of the whole mosaic. Since the chances are determined by a compressibility criterion applied to the mosaic as a whole, we also get incrementally closer to knowledge of the chances. Of course it's true that we won't have certainty about the chances until the end of time. And indeed, it was crucial to the dissolution of the undermining problem that we never attain certainty. But every time we learn a local matter of particular fact, we fill in one of the tiles in the mosaic. That means: cross off all of the potential ways the world could be incompatible with what we have learned and redistribute our credences over the rest. Over time, since there is a smaller and smaller number of ways the mosaic could be, we get closer and closer to knowledge of the correct theory of chance.

C2P33 As soon as we remove the boundaries that define the mosaic, nothing we observe brings us closer to knowing what the total pattern is and nothing brings us any closer to knowing the chances. Literally, *nothing* that we observe tells us *anything* about the chances.

C2P34 Here's the argument. We start with three premises:

C2P35 (i) The set of possible mosaics is obtained by a combinatorial principle; any assignment physical quantities to spacetime points represents a possible mosaic;

C2P36 (ii) The laws and chances are determined by a global criterion applied to the mosaic; and

C2P37 (iii) The mosaic is indefinitely extendible.

C2P38 Indefinite extendibility means just what it sounds like. It means that the Humean mosaic is open-ended; it stretches indefinitely into the future. Note that it doesn't entail that the Humean mosaic is infinite. It just means that there is no particular finite size that it is constrained to be. Any finite history can be extended indefinitely into the future.

C2P39 Why think the Humean mosaic is indefinitely extendible? There are two reasons. From a Humean perspective, to deny indefinite extendibility would be to hold that the existence of any collection of events was incompatible with the existence of some other. And that would be to deny Humeanism, because Humeanism was precisely the denial that there was any necessary connection between distinct existences. Lewis used to introduce what it meant to be distinct existences by saying the existence of one placed no restriction on the existence of the other.⁷ The reason that indefinite extendibility matters is that the laws and chances supervene on the Humean mosaic as a whole, not on any part of it. And any submanifold, no matter what the spread of events over that submanifold, can be embedded in indefinitely many mosaics whose total spread of events—judged by the criteria of overall fit—supports any chosen theory of chance that you like. This should be intuitively obvious. Fill in any patch P of observations up until a time t , and choose any theory of chance T_w , and it is easy to find a mosaic that would embed P and whose overall pattern would support T_w . And no observation or set of observations could make one theory of chance more probable than another. A nice fit with the pattern in one patch of spacetime can be undone by another and overridden by any larger patch of the same size. Best fit with the pattern over some finite patch has no bearing on best fit over the manifold as a whole, in a manifold of indefinite size. In an indefinitely extendible universe, nothing that you learn from observation will get you any closer to knowing what the best systematization of the whole mosaic will ultimately be, and that means that there is no way of learning from experience what the current chances are.

C2P40 It is important to understand how misleading it is to think of 'the mosaic' as though it is a definite four-dimensional structure of known size. In a finite universe of known size, every tile in the mosaic gives us *some* information about the global pattern. We start out with a finite set of ways the universe could be; for the Humean these are obtained by recombination on assignments of values of physical quantities to spacetime points. Every observation allows us to rule out some of these, and we redistribute credence over the rest. Every observation brings

⁷ Since indefinite extendibility is a modal claim, it comes in different modal strengths: the 'can' of metaphysical possibility and the 'can' of physical possibility. Only the weaker is needed for the argument here, but according to our best current physics, the stronger also holds. In Newtonian mechanics and SR spacetime is infinite in every direction. In GR, the global structure of spacetime depends on the matter distribution a theorem of John Manchak's (2009) shows that any model of General Relativity (finite or infinite) can be extended indefinitely by interspersing volumes of spacetime in a way that preserves the truth of the field equations.

us closer to knowledge of the pattern of events over the mosaic as a whole, and hence closer to knowledge of the chances. Once indefinite extendibility is taken into account, in neither of the senses above do we get closer to knowledge of the chances. Conditionalizing on individual local matters of particular fact doesn't get us any closer to making the *global* assessments of fit that determine the chances. So if the Humean mosaic is indefinitely extendible, i.e., if it spreads along every dimension without limit and is not metaphysically constrained to be some particular finite length, and chances are determined by overall fit with the whole spread of events, we have no intelligible story from the inside of how to go from a body of observed fact to hypotheses about what the chances are.

C2P41 Consider a numerical analogy. You are given a small patch of a two-dimensional array of integers (positive and negative whole numbers) and told that they form part of an array of indefinite extent. (See Figure 2.1.)

C2P42 Now suppose that I ask about global properties of the array as a whole: What do *all* of the numbers sum to? What is the probability that the sum of the top row is greater than 7893? What is the probability that overall the diagonal contains more 5's than 4's? What is the probability that the array as a whole sums to 667?⁸

C2P43 If there are answers to these questions, how does what you've been given here (this section of the mosaic) count as *evidence* for them? And what do you *learn* when another number is revealed? What does an additional bit of the array tell you about the totals in question? How are you supposed to update your probabilities?

C2P44 If the mosaic were bounded and of known size, it doesn't matter what your priors are, you learn from what you see. Conditionalizing on new observations takes you closer to knowing the pattern over the total mosaic, and hence closer to

3	-6	-16	-4	16	6	1
5	9	-14	-10	14	7	-11
15	13	19	-20	17	-23	-21
-2	-8	-18	0	18	8	2
-15	-13	-17	20	-19	23	21
-5	-7	24	10	-24	-9	11
-1	-12	22	4	-22	-12	-3

Figure 2.1

⁸ It would be fair to ask, what is meant by 'the array' here? Just being given an initial segment like this is not enough to specify an array. If reference to 'the array' is to be well defined, some principle must be provided for specifying the array (a recursive function, for example), and under some conditions that principle can be used to fix answers about global properties like sums or lower or higher bounds. In the case of the Humean mosaic, there is no recursive principle or other means of demarcation. We are told only that it includes all events—past, present, and future.

knowing the correct theory of chance. The indefinite extendibility of the mosaic, however, means that nothing that you see takes you any closer to knowing what the full mosaic looks like; nothing that you see makes one theory of chance more likely than another.⁹ Here's a quick hermeneutic argument: Humean credence for what the chances are has to be distributed across all epistemically possible mosaics. Divide the set of possible mosaics into equivalence classes according to which theory of chance they support. Every time we conditionalize on an observation, we cross off mosaics in each of these equivalence classes, but we never reduce the number of classes nor (on any natural way of counting) diminish their relative sizes.

C2P45 If part of what made Humeanism attractive was that it gave us a connection between chances, the local matters of particular fact that we treat as evidence for them, and the future events that they guide credence about, indefinite extendibility severs that connection. By linking the Best System specifically to the *totality* of facts, Humeanism cuts the probabilistic/evidential connection between local matters of particular fact and the correct theory of chance. In a finite manifold of known size, it didn't matter what your priors were, conditionalization on local matters of particular fact would modify your beliefs about chance in prescribed ways. The combination of indefinite extendibility, a criterion for determining laws and chances that applies at the global level, and the absence of constraints on the relationship between one event, or one submanifold, and another means that conditionalization on local matters of particular fact does not generally have any impact on beliefs about the chances. No amount of looking at the world will bring you any closer to knowing what the chances are. There is simply not enough structure on the probability space to tell you how to learn from experience.

C2P46 At least two of these things seem non-negotiable. The defining metaphysical dogma of Humeanism—that there are no necessary connections between local matters of particular fact—entails both the absence of constraints on the relationship between contents of one submanifold and the next, and indefinite extendibility. The remaining one—the global criterion for determining what the chances are—is just the BSA. So it is difficult to know how to resolve the issue. For the Humean, theorizing about chances is theorizing about the global properties of an indefinitely extendible manifold in which every local matter of particular fact is metaphysically independent of every other, and there is no connection between the pattern over some initial segment and the pattern over the whole. Any leaning towards this or that total pattern in a submosaic could be followed by another in which all trends were reversed or overwhelmed.¹⁰

⁹ The likelihood of a hypothesis is the probability of the evidence conditional on the hypothesis.

¹⁰ The vocabulary of submosaics is more accurate relativistic language. Indefinite extendibility applies to the spatial dimensions as well as the temporal, and any submosaic could be embedded in a larger in which all observed trends were reversed or overwhelmed.

- C2P47 Identifying laws and chances with distributed patterns in the mosaic of fact works really well when dealing with a finite mosaic of known size. It gets the connection with past facts that made it possible to use observed facts as information about the laws and chances, and the connection with future facts that made them good guides for belief in the face of ignorance. Every local matter of particular fact takes you incrementally closer to knowing the pattern over the total mosaic and so it is inherently information about the chances. Indefinite extendibility severs the connection between local matters of particular fact and the chances.
- C2P48 Global supervenience works well for many things. It is plausible that facts about beauty or value supervene globally on the physical facts. In those cases, reductions are to quell any worry that the facts in question introduce something metaphysically strange or something that can't be fit into a naturalistic view of the world. But it does *not* work well for chance since (as Lewis painfully pointed out in connection with undermining) the epistemic role of chance is to *guide belief in the face of uncertainty of the future*. Whatever chances are, they have to be the kinds of things that could play the epistemic role of chances, which means that they have to be the kinds of thing that situated agents could know about and use to guide expectation.¹¹ But once indefinite extendibility is taken into account, the whole neat epistemology that I described in the first half of the chapter, in which every local matter of particular fact takes you quantifiably closer to knowing what the chances are, while strategically avoiding undermining by not committing itself to the non-existence of undermining futures, falls apart. The problem is a tension between the Humean account of what the chances *are* (distributed patterns in the manifold) and what they *do* (guide belief about the future in the face of uncertainty).¹² Agents need to be able to learn about the chances, and to form reliable beliefs about them in the situation in which they need them to guide belief about the future. Unless there is a sensible account of how agents can have reliable beliefs about Humean chances in advance, Humean chances can't be the sorts of things we use to guide belief. It's really the same problem that Lewis saw behind the worry about undermining.
- C2P49 Everything I've said up until now should be uncontentious. It is the epistemological upshot of yoking the Humean account to the pattern over the whole Humean mosaic where the mosaic is indefinitely extendible and intrinsically

¹¹ We can be unfussy and non-committal about what 'know' amounts to here, beyond saying that it falls short of certainty and involves the ability to learn from experience. Any way of firming up the notion that captures its central philosophical uses and falls within this range will do.

¹² Did Lewis consider indefinite extendibility? I don't think so. Lewis did not note that the step from 'all that there is is one damn thing and then another' to 'there is a totality of local matters of particular fact' is a substantive one. As far as I know, the only place that Lewis discussed it is in "Postscript to 'Things Qua Truth-makers': Negative Existentials," written with Gideon Rosen (Lewis and Rosen 2003). They argue in that piece that the world should be treated as a concrete particular and provide the truthmaker for negative existentials. This addresses truthmaker problems associated with indefinite extendibility but doesn't help with the epistemic considerations here.

unconnected. The situation for the Humean agent trying to form beliefs about chances is the precise analogue of someone given what they are told is the initial segment of an indefinitely extendible string of integers in which each is stipulated to be independent of the others, and they are forced to make judgments about the sum. In that setting, what they have been given tells them nothing about the sum, and the revelation of additional numbers tells them nothing more: it brings one no closer, rules nothing out, narrows nothing down.¹³

C2P50 And it is worth pointing out that indefinite extendibility not only problematizes the epistemology of chances but also problematizes the role chances play guiding belief, and for something like the same reason. Chances defined by global fit with an indefinitely extendible manifold have no definite, quantifiable bearing on what happens next in the here and now.

C2S6

2.6 A Humean Response

C2P51 Here is a response suggested by Loewer and Albert.¹⁴ They say that the problem I've pointed out is just the problem of induction. The solution is to put constraints on priors that favor worlds which are inductively hospitable. What that has to mean in this setting is that the laws and chances that would be derived from systematization from a large enough initial submosaic are reflective of those in the whole. They go on to say that everybody has the problem of induction, so this isn't specific to Humeanism. And they add that the problem is double for non-Humeans. Humeans need induction to generalize from knowledge of part of the manifold to the whole, but then they are done. Non-Humeans, by contrast, need induction to go from part of the manifold to the whole, and then on top of that they have an unbridgeable gap between the pattern of events in the manifold and the laws. Since non-Humeanism is a non-reductive view, even knowing the whole Humean mosaic won't fix the chances.

C2P52 Let's look a little more carefully at what this response amounts to. In probabilistic terms, it amounts to recognizing that there is no link internal to the probability space between beliefs about local matters of particular fact and beliefs about chances. The logical structure of the space leaves beliefs about chances unaffected by conditionalization on local matters of particular fact. The proposal here is to impose a link *externally* by constraining priors so that they correlate the pattern of events over the submosaic from which observation is drawn to the pattern over the whole. It is non-trivial to make this precise in a way that would

¹³ Integers, of course, are positive and negative whole numbers, so one can take away with the next number what one added with the one before, or take away with the next string what one added with the string preceding. Indefinite extendibility is compatible with the universe being *actually* finite, but credence for what the chances are has to be distributed across all epistemically possible mosaics.

¹⁴ In conversation.

achieve the desired effect, but let's suppose that it could be done. Let's review the situation: When you have a finite mosaic of known size there is a link between local matters of particular fact and facts about the total mosaic that is built into the probability space so that, when you conditionalize on observation, you automatically update your credences for laws and chances. In that setting, it doesn't matter what priors you start with, conditionalizing on local matters of particular fact will take you incrementally towards knowledge of the chances because it will take you incrementally closer to knowing what the whole mosaic is like. That link is severed in an indefinitely extendible setting. Conditionalizing on particular observations has no intrinsic effect on the global pattern that determines the chances. The Humean suggests that we supplement Humeanism with constraints on priors that tie the systematization of the past to the systematization of the whole. The availability of such priors tells us nothing. One can *always* choose priors that will set up a link between facts that have no intrinsic link. I could choose priors that link the value of Genentech stock to the price of tea in China. I could choose priors that link the color of my socks on any given day to the color of the Queen's underwear. That doesn't mean that I can learn about the Queen's underwear by looking at my socks.

C2P53 In physical terms, constraints on priors embody assumptions that one makes about the way the universe is *before* you have any evidence and they guide how you update on what you see. The constraints that the Humean would need to impose in order to make scientific practice as it stands a way of learning about chances would be the assumption that the laws and chances of the best systematization of the patch of the universe that falls within our own past reflects those laws and chances that one would get from systematizing the whole. The problem with this for the Humean is that putting constraints on priors that heavily discount mosaics where systematizations of the past do not reflect systematizations of the whole is discounting the bulk by far of what they regard as possible mosaics. Consider, for example, the full set of random strings of integers of some particular length—say 20,000 integers long—obtained by a combinatorial principle. The overwhelming majority of those strings will not be ones that satisfy this condition. So the Humean is committed to pairing a combinatorial principle for obtaining future possibilities with an epistemic principle that says: even though there are as many ways the future could go as there are ways of assigning physical quantities to spacetime points, in forming beliefs about the future heavily discount all but a very small sliver of those.

C2P54 If we follow this line, for the Humean induction is a heavy bet that is *pre-empirical, essential* to being able to draw any conclusions from experience about the chances, and at odds with a metaphysics that is explicitly committed to recombination at the level of local matters of particular fact. It is also arbitrary, since we could just as well, and with no less justification, adopt priors that assume that the global pattern will form an American flag or that it will surround our

little patch of spacetime in a sea of bland uniformity. The intuitive naturalness of the assumptions embodied in these constraints should not mask how unnatural they are from a Humean point of view. Think about the connection between Humean chances and the kind of regularity that is being assumed to support induction. The existence of this kind of regularity makes summarizing *possible*, but few summarizable mosaics would be supported by the Humean epistemology with the inductive premise. There are innumerable ways that the Humean mosaic could be that would make it easy to summarize but not hospitable to induction. If the mosaic formed a giant pattern of the American flag, for example, or if after a brief period of apparent complexity everything turned into a simple uniform gas, or if it followed a plan laid out in two pages of the Book of Mormon. It makes sense if you live in a world with the sort of structure that supports inductive practices that you would exploit it in summarizing, but it makes no sense to *assume* it in a Humean world looking forward. Indeed, it is at odds with the defining metaphysical dogma of Humeanism. Humeanism is defined by the denial that there are necessary connections between local matters of particular fact. It is defined by the belief that no matter what has happened up until this moment, moving forward there are as many ways the future could go as there are ways of assigning events to spacetime that lie in the future. Betting heavily on priors that favor induction for a Humean is like me saying: “look, I know the color of the Queen’s underwear is metaphysically independent of the color of my socks. I know, that is to say, that any combination of colors for my socks and her underwear is possible. But I can see my socks, so I’m going to assume priors that heavily discount all combinations except those in which they are the same and use observations of my sock to update my credences for the color of her underwear.” If you are a Humean, any regularities that emerge over the course of history can be exploited after the fact to give a compact description, but there is nothing that rationalizes assuming this sort of regularity looking forward.

C2P55 Dustin Lazarovici (2020) has recently argued that typical Humean manifolds won’t permit systematization. Exactly the same argument will show that *even if we restrict attention to those manifolds that permit systematization*, the assumption that Loewer and Albert are recommending we build into the pre-empirical constraints will typically fail and fail quite badly. Following the Humean epistemology in a Humean universe in which there *are* Humean chances and laws will typically lead us *away* from (not towards) them.

C2P56 This proposed supplement or revision to Humeanism comes at considerable cost in the elegance of the view. Canonical Humeanism says: looking forward, there are in fact as many ways the world could be as you would get by extending history without constraint. You wait until the mosaic is complete and then you systematize the full body of fact. Systematization is just informational compression. It exploits whatever regularities the manifold possesses and combines it with whatever assumptions of fact pack enough predictive punch to warrant their

inclusion. The laws and chances are just axioms in this systematization. There is no more metaphysics than that. The epistemology is very simple: God solves for the laws and chances and delivers them to agents to use as guides to belief.¹⁵

C2P57 Here's the new view. It's the same metaphysics as above: looking forward from any point in history, there are as many ways the world could be as we would get by assigning values of physical quantities to spacetime points in the future. The divine epistemology and metaphorical sit-down with God are discarded. It is now agents who have to discover the chances. Even though the metaphysics says that looking forward from any point in history, there are as many ways the world could be as we would get by assigning values of physical quantities to spacetime points in the future, the epistemology says that you must take as a pre-empirical assumption that the laws and chances derived from any large enough submanifold would reflect the laws and chances derived from a global systematization. This amounts heavily weighting your priors to ignore all but a small sliver of Humeanly possible completions of the mosaic. Since the metaphysics is explicitly committed to combinatorial possibilities for the future, the only thing that keeps this from being flat-out inconsistent is that one reserves *nominal possibility* that the future might be among the vast majority of worlds whose overall systematization is different from that of the initial segment.

C2S7

2.7 What's the Alternative?

C2P58 At this stage, the Humean will ask: what's the alternative? He will say that one has to assume induction works for scientific practice to make sense. Nobody has a solution to the problem of induction. There is no metaphysical guarantee that the future will be like the past. And if one assumes that induction works, the Humean account is still the best game in town.

C2P59 There are two things to say about this. The first is that if it's a question of choosing priors that make sense of scientific practice, why not start with priors that expect no correlations unless there are connections? If those are the priors that you start with, everything in your experience will be telling you that the world is *not* an intrinsically unconnected pattern of fact. You will find a wide body of correlations—some local and temporary, some deep and pervasive—and

¹⁵ Readers of this literature will be familiar with the primordial myth motivating the Humean account in which David Albert relates a sit-down with God where you ask God to tell you about the world and He begins to list every event one by one. Albert says,

"[Y]ou explain to God that you're actually a bit pressed for time, that this is not all you have to do today, that you are not going to be in a position to hear out the whole story. And you ask if maybe there's something meaty and pithy and helpful and informative and short that He might be able to tell you about the world which (you understand) would not amount to everything, or nearly everything, but would nonetheless still somehow amount to a lot. Something that will serve you well, or reasonably well, or as well as possible, in making your way about in the world" (Albert 2015, p. 23).

scientific theorizing will take the form of systematizing those connections. Once those connections are systematized, they provide the basis for prediction, intervention, design.¹⁶ In the old days the assumption was that the ‘connections’ would ultimately take the form of deterministic laws. Although most theorizing is probabilistic (we have large bodies of data, look for correlations among the values of measured quantities, and assess likelihoods of hypotheses about the structure of the source), the presumption was that stochastic regularities would be underwritten by deterministic laws and eliminated in the fundamental theory. Since quantum mechanics we have at least become comfortable with fundamentally stochastic laws, and statistical mechanics has made the role of probability even in deterministic settings clear. In either case the information that provides the basis for theorizing is probabilistic, and theories that are the product of systematization will involve probabilities in various guises.

C2P60

What is there to recommend these priors over the Humean priors? I’m Bayesian enough to think that priors are priors. There are no rational constraints on priors and no *a priori* argument for one set of priors over another. Descriptively, this is close to the set of priors that we employ in everyday reasoning. It is a generalization of what the child does when approached with a new toy or an opaque box with handles and buttons (Gopnik et al. 2004), and I suspect that it is close to the unreflective, default view of most scientists.¹⁷ It does a better job hermeneutically of making sense of scientific practice. Observed correlations suggest connections.

¹⁶ Saying that they provide the basis for intervention and design is a way of saying that we don’t just rely on them to make predictions about the future; we rely on them in the kind of hypothetical reasoning and assessment of possibilities that is characteristic of those contexts. That is all that their modal force comes to. The remarks above were focused on chances since there is a well-developed discussion that puts the epistemology of chance front and center, but I switch now to talking about laws and chances together. These are the joint product of global systematization; together they provide the fundamental modal outputs of theorizing.

¹⁷ Consider an analogy. Unlike people, countries can last indefinitely long into the future. Credit ratings are assigned to countries, as they are to people, to assess creditworthiness and assist lenders in deciding whether to issue loans. Suppose I asked you: what is the credit rating of a country? Is it a summary of the borrowing behavior over the course of its existence? If it was, you’d never be in a position to so much as offer a guess at the credit rating of a country. It would always depend on how things go, and since countries can continue indefinitely into the future, you wouldn’t even be able to say whether what you knew from past behavior had any probative value and what that value was. Credit ratings weren’t meant to summarize behavior over complete history (past, present, and future) but to extract from past behavior information relevant to future behavior. They are needed to tell banks whether to lend money in the here and now. Credit ratings take into account payment history, amount owed, length of history, types of credit used, things like that. And the reason that credit ratings can play the role they play, i.e., the reason that they are a good predictor of future behavior, is that there are stable features of the system to which they are assigned that is manifested in their past history and is a good indicator of future behavior. The whole practice of assigning credit ratings is predicated on the idea that there is some kind of regularity that guides what people do that we can separate for all of the contingencies of their situation and that will provide some guidance about how they’ll behave in the future. In the case of credit ratings, this is all quite informal and seat of the pants, but it is not different in kind from what we do in scientific settings. The whole business of looking for laws and chances is predicated on the idea that there is some kind of regularity that guides the behavior of physical systems which we can separate from all of the contingencies of their situation and that will provide some guidance about future behavior and behavior across a range of hypothetical situations.

Connections are analyzed and tested. Once established, the connections are relied on in prediction, design, and intervention.¹⁸

C2P61 What are these ‘connections’? Formally they are the symbolic expressions that can restrict the joint values of parameters, or the probabilities of values (e.g., $f = ma$, Maxwell’s Equations, Born’s Rule). They often take the form of laws of temporal evolution (Schrödinger’s equation, Hamilton’s equations). Physics offers no ‘analysis’ of what these connections are. Philosophers became interested because certain ways of thinking about laws made them seem problematic. The way we present physical theories makes it sound like we start with a space of possibilities and pare them down by adding laws. This gives us a sense that laws act as restraints that *keep things* from happening. And the notion of law, of course, is borrowed from the human domain and suggests restrictions on freedom. Hume problematized the idea of physical necessity with his critique of causation and forced a reckoning with the idea that there is anything in the world over and above events and their regularities. It became common among empiricists to hold that laws were regularities and to deny that there is anything in the world that enforces them.

C2P62 I want to suggest that this was the wrong way to think from the beginning. Notions of physical necessity are dual to notions of physical possibility and we should focus instead on the notion of possibility. Physical theories give us a notion of physical possibility that comes from a kind of analysis and synthesis that is familiar since Newton. We are accustomed to thinking in terms of analysis of complex systems into spatial components, but the real basis of systematization is analysis of complex *motions* or behaviors into simple ones and then recombination of those simple motions to derive the possibilities at the higher levels.¹⁹ The result of this process is that there’s a radical reduction of degrees of freedom looking forward from any point in spacetime relative to what we would get if we followed the Humean prescription of treating every local matter of particular fact as an independent degree of freedom. If we view possibility as the more basic notion and treat it in this constructive way, the idea of laws as restraints goes away. Theories are an attempt to discern latent possibilities in the world. It is about

¹⁸ The fact that they are relied on in design and intervention is what makes them modal. In design and intervention we are assessing not only what will happen but what would happen under a range of merely possible conditions. Laws are meant as much to capture the full modal latitude inherent in the world as to generate predictions. See Hicks (2018) for a nice discussion of the difficulty that traditional Humeanism has accounting for this aspect of practice. Hicks’s article shares much of my own sense that understanding Humeanism means abandoning the atemporal God’s-eye view epistemology and giving a situated account that gets the role of laws and chances from an embedded perspective right. See also Jaag and Loew (2020).

¹⁹ ‘Motions’, here, doesn’t mean simple change of position, but change of state. Although Newton was primarily interested in change of position, nowadays we recognize internal properties alongside position; the state of a system is represented by a point in a phase space, and ‘motion’ means change of state represented by movement through phase space. In the philosophical literature what I’ve been calling ‘ranges of motion’ have sometimes been called powers. See Demarest (2017).

looking at the phenomena with an eye to separating the dimensions along which the phenomena can vary. Laws are expressions of restrictions not in the sense of ‘restraints’ but in the sense of ‘boundaries’.

C2P63 The problem with Humeanism is that it starts with a metaphysics that is explicitly committed to recombination at the level of local matters of particular fact. If we start from a more neutral standpoint with no pre-empirical commitment to how many ways the world could be and think of science as in the business of discovering the immanent possibilities in the world, Humeanism will look like it is overreaching. The more conservative view will turn out to be the one that recognizes only such immanent possibilities as are implicit in the phenomena. Through a combination of theory, observation, and experiment designed to push nature to exhibit her full capacity for independent variation, the scientist is trying to establish the limits of what is possible. That still leaves the notion of possibility, but it is a notion that has every claim to be conceptually rock-bottom independent of any connection to science. A view like this involves a reworking of the way that we think about physical modality. It has the virtue of involving none of the extraneous metaphysical machinery invoked by traditional anti-Humean accounts,²⁰ and I think it is closest to the one that most working scientists adopt. If I had to give it a name, I’d call it the ‘natural nomological attitude’ (echoing Arthur Fine’s ‘natural ontological attitude’ perhaps only in name).

C2P64 The second thing to say is that I agree that there’s no solution to the problem of induction in the form of a rational or metaphysical guarantee that induction works. But precisely because we can easily imagine worlds in which those practices fail, the success of the inductive practices embodied in science is a datum or a clue that we can use in understanding how our world is structured. Instead of grafting an inductive assumption onto a metaphysics for science that is at odds with it, we should lead by asking *why* science works. And we should be looking for an understanding of how the world is structured that makes the inductive methods embodied in our scientific practices sensible and non-arbitrary.

²⁰ Anti-Humean views are typically modalist, in a sense characterized nicely by Michael Hicks: “The regularity theory holds that laws of nature are merely generalizations.... Modalist views are less metaphysically perspicuous than the regularity theory because they claim that to be laws, a generalization must be backed, made true, or associated with a relation between properties (Dretske [1977]; Tooley [1977]; Armstrong [1983]), the essences of properties Shoemaker [1980]; Ellis [2001]; Bird [2007]), *sui generis* facts about production (Maudlin [2007]), or irreducible counterfactuals (Lange [2009]).... Although each of these metaphysical machines is distinct, they are united in holding that facts about nomic necessity, or facts with modal implications of some sort, are fundamental. In contrast to these views, the regularity theorist holds that the laws are not backed or made true by anything beyond their instances and they are made laws by nothing more than the sum of non-nomic facts at a world” (Hicks 2018, p. 984). These kinds of metaphysical machines, as Hicks calls them, are foreign to science, and Humeans rightly oppose them. Laws need no metaphysical backing. The modal force of laws just means that we rely on them not simply in predicting what *will* happen but in the kind of hypothetical reasoning we rely on in decision, design, intervention, and control. See my Ismael (2017), and John Norton (2021).

- C2P65 We don't have to look very far to find such an understanding. Here are some unguarded descriptions from physicists about how theorizing works in the kind of fundamental theory for which the BSA was originally designed. The first is from James Hartle:
- C2P66 Identifying and explaining the regularities of nature is the goal of science. Physics, like other sciences, is concerned with the regularities exhibited by particular systems. Stars, atoms, fluid flows, high temperature superconductors, black holes, and the elementary particles are just some of the many examples. Studies of these specific systems define the various subfields of physics — astrophysics, atomic physics, fluid mechanics, and so forth. But beyond the regularities exhibited by specific systems, physics has a special charge. This is to find the laws that govern the regularities that are exhibited by all physical systems — without exception, without qualification, and without approximation. The equality of gravitational accelerations of different things is an example. These are usually called the fundamental laws of physics. Taken together they are called informally a “theory of everything”. (Hartle 2002, p. 4)
- C2P67 The second is from Gerard T’Hooft, a very different type of physicist:
- C2P68 What is a “Theory of Everything”? When physicists use this term, ... we have a deductive chain of exposition in mind, implying that there are ‘fundamental’ laws describing space, time, matter, forces and dynamics at the tiniest conceivable distance scale. Using advanced mathematics, these laws prescribe how elementary particles behave, how they exchange energy, momentum and charges, and how they bind together to form larger structures, such as atoms, molecules, solids, liquids and gases. The laws have the potential to explain the basic features of nuclear physics, of astrophysics, cosmology and material sciences. With statistical methods they explain the basis of thermodynamics and more. Further logical chains of reasoning connect this knowledge to chemistry, the life sciences and so on. (T’Hooft 2017, p. 2)
- C2P69 These are relatively generic. One can find similar remarks in Feynman, Wigner, Hawking. The sort of picture is guiding the imagination when they are looking for theories is very different from the Humean one. The scientist doesn't think of herself as in the business of summarizing a global mosaic of intrinsically unconnected fact. She sees herself as in the business of trying to uncover an immanent substructure in the composition of material systems. The idea that is guiding the practice is that if you look closely enough at the fine-scale structure of the matter, you will find that systems everywhere are made of the same components with a limited range of motion. Observed regularities are used as guides to this substructure. The laws describing components are identified and then combined

with facts to yield predictions for systems built from the same components. Scientific theories involve analysis is to isolate the fundamental components of matter and the laws that describe their behavior, and then synthesis to proscribe the bounds of what is possible. Those boundaries are then viewed as genuine constraints on what can happen, not just retrospective summaries of what did.

C2P70 That is important because it addresses another source of discomfort about Humeanism. On the view here, the components of a complex system have modal profiles that it is the task of theory to discern and experiment to test. They can be arranged into arbitrary configurations and we can ask about such configurations how they would behave under conditions that may or may not be actual. In so doing, we are drawing conclusions that are rooted in the intrinsic structure of the system, and that we rely on not simply to *predict* but to *engineer* the future: to design and steer and avoid and forestall. One can see how the pieces of this view fit together to make sense of that practice. Not everybody shares this discomfort, but I have never been able to arrive at a way of thinking about modality on a Humean view that seems halfway adequate to its practical role. For the Humean, systematization is about informational compression. The things that end up as laws, and that govern our beliefs about what can and can't happen, are just the axioms in the system that achieves the best global fit. Ideally, the laws of the Best System would be so strong that they predict everything that happens. As it is, some things get squeezed out in the process of compression and are undetermined by the laws, leaving us with a range of possibilities. These appear as simply the regrettable overflow of systematizing a complex domain. But when an engineer is designing an airplane or we are exploring new ways of producing nuclear energy, one wants to think of the laws as genuine constraints on what *can* happen looking forward.

C2P71 None of this is offered as a philosophical *analysis* or *metaphysics* of laws. It really just amounts to taking science at face value and declining to see physical modality as problematic. Humeanism gets a lot of traction by opposing certain philosophical views of what the laws are, but Humeanism doesn't have enough structure on the domain to make it possible to learn from data without wheeling in constraints on priors that are at odds with its own metaphysics. The scientist relies on the idea that there are laws that can be discovered by investigating the material substructure of the part of the universe they have access to and projected into other parts of the universe, on the defeasible assumption that they are configurations of the same components. Where the Humean sees induction as separate from and prior to the discovery of laws (and chances),²¹ she sees it as the whole point of theorizing. If one wants a modern empiricist position tailored to science, why not simply accept the world of science at face value? It is not a world

²¹ I'll omit the '(and chances)' below for ease of expression.

of intrinsically unconnected matters of particular fact arranged in a mosaic that—if nature is kind—can be retrospectively summarized. It is a world of components with limited ranges of motion that combine into larger configurations to form the objects we see around us and whose behavior is derivative of that of their parts.

C2P72 Is this an anti-Humean position?²² The laws, so understood, don't govern or guide events; they are not transcendent relations between universals; they don't *produce* their instances or enforce behavior; and they aren't 'backed' by special modal truthmakers, so it is nothing like the familiar anti-Humean positions. On the other hand, it does recognize immanent constraints in the manifold of events. On this view, laws aren't retrospective summaries of what *did* happen; they are genuine constraints on what *can* happen looking forward. I'm torn between calling it neo-anti-Humean and anti-neo-Humean. The latter seems more appropriate because it is not a descendant of anti-Humeans in the philosophical tradition that Humeans oppose.²³ It is defined, rather, in opposition to Humeanism in the style that it takes in the hands of its most influential proponents.

C2S8

2.8 Conclusion

C2P73 Humeanism is an attractive view if one looks from a God's-eye view and one asks, 'what kinds of thing' would it make sense for God to recommend as guides for the beliefs of creatures like us (limited creatures with no specific information about the future)? But when one demands an epistemology that embedded agents can employ, these three things together leave us without a sensible story about how to learn about what the chances are from what you observe:

C2P74 (i) Chances are determined by a global criterion applied to the mosaic as a whole;

C2P75 (ii) There are as many ways that the mosaic could be as there are assignments of local quantities to spacetime points; and

C2P76 (iii) The Humean mosaic is indefinitely extendible.

C2P77 A fix for the problem suggested by Albert and Loewer was to presuppose that induction works. I argued that this was an unstable position for the Humean because it says effectively: even though the defining doctrine of Humeanism is that there are as many ways the future could go as there are assignments of

²² There are different ways to elaborate the metaphysics of a view like this. We could take laws or powers as primitive, for example. I don't myself have a preference.

²³ See Carroll (2016), also Armstrong (1983) and Schaffer (2016).

physical quantities to spacetime points, one should nevertheless assign nominal probability to all but a tiny sliver of them.

C2P78 The Humean can, and I suspect will, stick to his guns. But I think that when the Humean is forced to be explicit about how embedded agents are supposed to learn about the chances, some of the sheen comes off Humeanism, and empiricists might be persuaded to look elsewhere.²⁴

C2S9

References

- Albert, D. (2015). *After Physics*. Cambridge, MA: Harvard University Press.
- Armstrong, D. M. (1983). *What Is a Law of Nature?* Cambridge: Cambridge University Press.
- Bird, A. (2007). *Nature's Metaphysics*. Oxford: Oxford University Press.
- Carroll, J. W. (2016). 'Laws of Nature', *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition), ed. Edward N. Zalta, <https://plato.stanford.edu/archives/fall2016/entries/laws-of-nature/>.
- Cohen, J., and Callender, C. (2009). 'A Better Best System Account of Lawhood', *Philosophical Studies*, 145(1), pp. 1–34.
- Demarest, H. (2017). 'Powerful Properties, Powerless Laws', in Jacobs, J. (ed.), *Causal Powers*. Oxford: Oxford University Press, pp. 38–53.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., and Danks, D. (2004). 'A Theory of Causal Learning in Children: Causal Maps and Bayes Nets', *Psychological Review*, 111(1), pp. 3–32.
- Hartle, J. (2002). 'Theories of Everything and Hawking's Wave Function of the Universe', <https://arxiv.org/abs/gr-qc/0209047> (accessed 2 Mar. 2021).
- Hicks, M. (2018). 'Dynamic Humeanism', *British Journal for the Philosophy of Science* 69(4), pp. 983–1007.
- Ismael, J. (2008). 'Raid! Dissolving the Big, Bad Bug', *Noûs*, 42(2), pp. 292–307.
- Ismael, J. (2015). 'How to be Humean', in Loewer, B., and Schaffer, J. (eds.) *A Companion to David Lewis*. Oxford: John Wiley & Sons, pp. 188–205.
- Ismael, J. (2017). 'An Empiricist Guide to Objective Modality', in Slater, M., and Yudell, Z. (eds.), *Metaphysics and the Philosophy of Science: New Essays*. New York: Oxford University Press, pp. 109–25.
- Jaag, S., and Loew, C. (2020). 'Making Best Systems Best for Us', *Synthese*, 197, pp. 2525–50.

²⁴ This paper was originally presented at a conference organized by Barry Loewer. I want to thank him for the invitation and thank him and David Albert for really helpful discussion at the conference. I am also indebted to Christian Loew for the invitation to include my paper here and to him and Siegfried Jaag for very helpful comments.

- Lazarovici, D. (2020). 'Typical Humean Worlds Have No Laws', <http://philsci-archive.pitt.edu/17469/> (accessed 2 Mar. 2021).
- Lewis, D. (1994). 'Humean Supervenience Debugged', *Mind*, 103(412), pp. 473–90.
- Lewis, D., and Rosen, G. (2003). 'Postscript to 'Things Qua Truth-makers': Negative Existentials', in Lillehammer, H., and Rodriguez-Pereyra, G. (eds.), *Real Metaphysics: Essays in Honour of D. H. Mellor*. London: Routledge, pp. 39–41.
- Loewer, B. (1996). 'Humean Supervenience', *Philosophical Topics*, 24(1), pp. 101–27.
- Manchak, J. B. (2009). 'Can We Know the Global Structure of Spacetime?', *Studies in History and Philosophy of Modern Physics*, 40(1), pp. 53–56.
- Maudlin, T. (2007). *The Metaphysics Within Physics*. Oxford: Oxford University Press.
- Norton, J. (2021). 'How to Make Possibility Safe for Empiricists', https://www.pitt.edu/~jdnorton/papers/Empiricist_possibility.pdf (accessed 2 Mar. 2021).
- Pettigrew, R. (2012). 'Accuracy, Chance, and the Principal Principle', *Philosophical Review*, 121(2), pp. 241–75.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Schaffer, J. (2016). 'It is the Business of Laws to Govern', *Dialectica*, 70(4), pp. 577–88.
- T'Hooft, G. (2017). 'Free Will in the Theory of Everything', <https://arxiv.org/abs/1709.02874> (accessed 2 Mar. 2021).

Author Query

AQ1: Please write a caption for Fig. 2.1.