

A Treatise of Humean Nature: Being an Attempt to Introduce the Experimental Method of Reasoning Into The Metaphysics of Laws; And, Dialogues Concerning Natural Philosophy

Neil Dewar

*Faculty of Philosophy
University of Cambridge*

James Owen Weatherall

*Department of Logic and Philosophy of Science
University of California, Irvine*

Abstract

This paper contends that Humeanism suffers from a major flaw: it is unable to explain why the practice of experimental science—i.e., the practice of creating novel circumstances and seeing what happens—should be expected to lead to scientific knowledge. In particular, it is unclear how the Humean can make sense of local, concrete inquiry yielding knowledge which is global in scope, lawlike in character, and counterfactually robust. As a result, Humeanism is left without a plausible account of scientific epistemology.

If we see that knowing is not the act of an outside spectator but of a participator inside the natural and social scene, then the true object of knowledge resides in the consequences of directed action.

Dewey (1929)

1. Introduction

Humeanism is a prominent—we dare say leading—family of views concerning laws of nature, according to which those laws are (nothing but) summaries or systematizations of

Email addresses: nad42@cam.ac.uk (Neil Dewar), weatherj@uci.edu (James Owen Weatherall)

Draft of October 4, 2024

PLEASE CITE THE PUBLISHED VERSION IF AVAILABLE!

the regularities encountered in nature. As we discuss below, there are many varieties of Humeanism, with significant differences in motivation and substance. But certain tenets tend to be shared among its adherents. One is the claim that two possible worlds that are identical in all of their “occurrent” matters of local fact *ipso facto* share their laws. This is because the best system concerns only those facts that occur in the world, and not additional modal facts about what could have occurred under different circumstances. (This view is often called “Humean supervenience” and the collection of occurrent facts is called the “Humean mosaic”.) Another is the claim that laws of nature do not *govern*, in the sense that the laws do not bring about, or otherwise causally constrain or necessitate, the local matters of fact in the actual world or in (metaphysically) possible worlds. Both of these views are closely related to another doctrine often endorsed by Humeans, which is that the laws are the axioms, and for some also the theorems, of the best systematization of the local matters of fact in the world. The best systems view, coupled with Humean supervenience, can be seen as a strategy for reducing laws to non-modal facts.

Humeanism has many defenders and elaborators—and also many critics. We are among the critics. We sympathise with many others’ worries.¹ But as we argue here, much of the extant literature criticizing Humeanism misses or underplays the most important problem for the view. We claim that any adequate account of the laws of nature must be able to make sense of the epistemic practices by which we come to discover those laws. That is: an adequate account of laws must be compatible with scientific practice, in the sense that if the account is true, then scientists’ methods for learning about laws should be epistemically well-motivated, plausibly reliable, and compatible with their goals. More importantly still, insofar as one thinks scientists have been successful in discovering some laws of nature, or at least plausible candidates for such laws, we must accept that laws

¹Maudlin (2007) is particularly clear and persuasive on these issues; see also Hall (2015), Fernandes (2023) and Guo (2022) for recent views that resonate with ours. Note, too, that many criticisms of Humeanism are advanced by advocates of the view.

are the sorts of things that scientists could have discovered using the methods that they in fact use. Humeanism fails all of these criteria. Worse, the failures seem irremediable, as they implicate the central tenets of the view, shared by virtually all of its adherents.

So what is the problem for Humeanism? As we discuss in more detail below, the issue concerns the basic logic of experiment, which is a bedrock part of successful science. Experiments consist of active interventions in the world designed to test whether certain regularities obtain. Very often, the goal of experimental science is to create arrangements that would not otherwise have existed, and may not exist elsewhere in the world, to determine whether those arrangements evolve in ways predicted by or compatible with the laws. More, the particular arrangements that are brought about are chosen precisely because they are challenge cases for the laws as hypothesized at a time. That is, experimental science crucially involves intervention for the purposes of probing the limits of the laws. So, as a first pass, the problem with Humeanism is that insofar as experiments are successfully used to discover and test laws of nature, the laws cannot be mere summaries or systematizations of local matters of fact, because any such summary will not generally be robust under the full range of experimental interventions scientists might concoct.

The Humean will have a ready response to this, along the following lines: experimenters, their setups, and their results are all part of the mosaic of local matters of fact. The laws systematize those occurrent facts, too. This is correct—but not to the point. The issue is that experimenters design the experiments they do because they expect that if the laws are true or approximately true, then the outcomes of the experiment will be in accord with the laws—more, that insofar as the laws are correct, they will “govern” the outcomes—and they wish to see if those expected results obtain. If the laws are nothing but a systematization of what happens in the world, including the experimental results, then the practice of articulating causal conditionals about what would happen if an experiment were set up in such and such way, and then actually bringing about that setup to see

if it does happen as expected, would have no justification. Humean laws are not the kind of thing one can probe by intervention, for two reasons. First, they are not projectable: there is no reason to expect that the axioms of a best system as discovered in some region of space and time—say, near Earth over the lifetime of the human species—will have any bearing at all on how a system that has never been encountered in that region of space and time will behave. And second, they are not robust. Were the mosaic in fact different, or made to be different by an experimental intervention, the laws would be different as well.²

Humeans will have responses to this as well, but the outline of our arguments should now be clear, and so we will postpone further engagement for the sequel. The remainder of the paper will be organized as follows. First, we will introduce Humeanism with more care, presenting several different versions of Humeanism and identifying key differences in their motivation and execution that will matter in what follows. From that we will extract the core tenets of Humeanism that we seek to refute. Next, we will say more about the logic of experimental science. Our goal is not to give a systematic account of experimentation; that would clearly be impossible. But we will highlight some key features, drawing on examples from physics, that will illustrate our point in what follows. Section 4 will present the main argument and respond to likely replies by the Humean. Then, in section 5, we will step back and make some more general comments about the methodology of Humean metaphysics in light of successful science. We will also sketch some minimal features we think an account of laws would need to have in order to adequately meet the criteria we claim Humeanism fails. Section 6 concludes.

²“We have a semantics for counterfactuals!” cries the Humean at this juncture. True, but it is no help. We return to this point in section 4.

2. The Humean Species

There is some sense in which Humeanism can be traced back to Hume, though the connection is tenuous and not especially informative regarding contemporary Humeanism, and so we will not pursue that.³ Instead, we wish to highlight three distinct strands of thought that often go by the name Humeanism, with different motivations and different details; and then we will try to identify some common commitments among them.⁴ The arguments we wish to make target core commitments shared across virtually all of these strands.

The first strand of thought originates with Mill (1843), who introduced the idea that the laws of nature are best seen as the axioms of a logical system describing the observed regularities of nature.⁵ Mill was not obviously committed to a reductionist metaphysics about laws, and while he certainly claimed that there are systematic regularities in the world, he does not take a position about nomic necessity. He seems to be agnostic about why there are regularities. For Mill, the main question is which, of the many true generalizations we might discover in nature, should get the special designation “law”, as opposed to some other title. His answer is that when we try to systematize the many generalizations we discover, we will find some subset of those from which the others may be derived. Those generalizations will have a special status, which Mill suggests is sufficient to count them as laws. Note that for Mill, it is the *axioms* of the best system, not the axioms and theorems, that count as laws.

Mill is largely silent on many of the metaphysical claims that animate the contemporary

³Earman (1984) begins with a brief summary of different claims in Hume’s opus that can be seen to motivate Humean views of laws. More frequently, authors who wish to invoke Hume directly argue that Hume claimed there are no necessary connections between distinct existences, and this rules out more robust views of laws. But this strikes us as an oddly metaphysical reading of Hume, who we read as denying only that we can *know* that there are necessary connections between distinct existences.

⁴We do not necessarily claim that these strands exhaust the varieties of Humeanism, but they do reflect the vast majority of the literature.

⁵See especially (Mill, 1843, Vol 1, Book 3, Ch. IV, §1). Ram (1928/1978) also advocated for a view like this one.

literature, and it is not clear that we have a quarrel with him. One can read him as a deflationist who would maintain that there are apparently some regularities that can be observed in nature, and which can be described, summarized, and systematized, but which we have no grounds for expecting to continue to obtain upon future observations. But on the other hand, given the richness of his writings on scientific methodology, and his view that one can, via his own methods, discover causal structure in the world, there is also a plausible reading on which his version of Humeanism is compatible with the views we defend below. However, insofar as this is the case, Mill's best systems account is only tangentially related to contemporary Humeanism.

The second strand of Humeanism is most clearly articulated by Earman (1984). This version of Humeanism, like Mill's, is not committed to a reductionist account—but it also is not necessarily committed to a best system analysis. Rather, its core is a rejection of modal facts distinct from “occurrent” facts. This position is best captured by Earman's “empiricist loyalty test”, which demands commitment to the claim that two worlds that agree on all of their occurrent facts also agree on their laws. The view is less about saying what laws are than saying what they are not: laws are not the sorts of things that can support modal facts extending beyond what actually occurs in the world, and they do not govern occurrent events.⁶ Earman and Roberts (2005b,a) develop a version of the Humean mosaic that can support this view, which involves a notion of “reliably measurable” quantity. They endorse a form of Humean supervenience, where laws supervene on the mosaic, but deny a reductionist view on which all nomic claims reduce to non-nomic (or non-modal) ones, since they acknowledge that “measureability” has a modal character.

The final strand is the most famous. This is the view most strongly associated with

⁶Beebe (2000) offers an extended and influential argument that “governing” conceptions of laws of nature are in tension with the Humean view; see also Emery (2023) for a recent response.

Lewis (1973). Lewis can be seen as combining aspects of the Mill-Ramsey best systems account and the empiricist-inspired views from Earman. But he also adds further positive claims to both, and he packages his view on laws with compatible views of chance and counterfactuals. Focusing on the account of laws: Lewis's view is that what exists is a mosaic of "local matters of fact" consisting of distinguished properties distributed throughout space and time. This is his version of the mosaic. The laws are the best systematization of this mosaic, where "best" is defined as simplest and strongest. This account is fully reductionist, in the sense that the mosaic consists entirely of non-nomic, non-modal structure. Laws do not constrain the mosaic in any way; in fact, Lewis endorses a principle of "free recombination" on which any distribution of properties at all is (metaphysically) possible.

There are many other defenses, versions, refinements, and attempted improvements of these views—particularly Lewis's—in the literature. Some focus on improving the characterization of "best" system, to allow for virtues more closely connected to the virtues employed in scientific practice (Hicks, 2018; Dorst, 2019; Jaag and Loew, 2020). Others liberalize Lewis's view of the mosaic to allow laws that supervene on different properties, or properties at different levels, to make better sense of laws in the special sciences (Cohen and Callender, 2009). Yet others broaden the structure that gets fixed by systematization, to include not only laws but properties, spatiotemporal structure, or even the quantum wavefunction (Hall, 2015; Loewer, 2021; Huggett, 2006; Miller, 2014; Callender, 2015; Esfeld et al., 2014; Bhogal and Perry, 2017). And so on.

Since our target is the whole family of views, these details will not matter. For our purposes, it suffices to draw out certain commonalities between these three strands and their successors. First is the idea that whatever the laws are, they supervene on the actual regularities in nature, in the sense that the laws could not be different unless the occurrent events were different. Second, laws do not contain modal information that

could determine whether one kind of modification of the actual world is possible as opposed to another. Third, the laws do not govern in the sense of enforcing how the world must be given any partial specification of how the world is—and so the laws do not determine the future given the past. More generally, the laws do not constrain the mosaic.⁷ The mosaic occurs or exists and the laws summarize it. Finally, the Humean is committed to the claim that the laws are not projectable, in the sense that the best system of some local region need not be the best system in other regions or globally.

These claims are clearly related. Even so, we grant that perhaps not every Humean accepts all of them. There may even be some philosophers in the Humean tradition, such as Mill on a plausible reading, who would deny all of these. But we think the vast majority of Humeans are committed to most, if not all, of these claims. We will presently give multiple arguments intended to draw out the tension between experimental science and these theses. The goal will be to show a broad tension between Humeanism and scientific practice, not to refute any specific variant of the view. As Hume would say: commit them all to the flames.

3. The Logic of Experimental Science

We take the essential idea of experimental science to be the following, We sometimes find ourselves in a situation where the truth of some hypothesis H about the behaviour of a certain kind of system under certain circumstances is in question. We realise that it is within our power to put a system of that kind into such circumstances. We do so, and we observe whether or not H obtains. In either case, we learn something about the system in question that will guide future inquiry. Very often, we are interested in H because

⁷This points toward two ways that non-Humeans might develop their views: laws as constraints, or laws as mechanisms for producing the future from the past. See Adlam (2022) and Chen and Goldstein (2022) for defences of the former picture, and Maudlin (2007) and Sebens (2024) for defences of the latter. We do not take a stand on this debate here.

it follows from some postulated laws of nature, or more generally from some scientific theory. In those cases, we investigate H because if the law or theory were true, then H would obtain were the necessary circumstances brought about. Actually bringing those circumstances about allows us to test the laws or theory. If H obtains, we take that to provide evidence for those laws or that theory; if does not, then that is evidence against the laws or theory.

Of course, this is hardly a novel or sophisticated account of experimental practice (c.f. Bacon, 1267; Mill, 1843; Hacking, 1983). But despite its familiarity, we contend that it highlights crucial features with which Humeanism must struggle: that *circumstances brought about to test H are very often entirely novel*, i.e., have not been brought about before; and that the goal of bringing those circumstances about is to determine whether some generalization extends to those novel circumstances. In other words, scientists identify subjunctive counterfactual conditionals of the form “if this law were true, then under these (novel) circumstances, such and such would occur”, and then bring about those novel circumstances to seek evidence that the law is true.

Take Hans Christian Ørsted’s pioneering experiments in electromagnetism as an example.⁸ In April 1820, Ørsted found that if a current is run through a wire, a compass needle placed parallel to the wire would turn to orient itself perpendicularly to the wire. The insight derived from this observation was that there is a relationship between electrical phenomena (the current in the wire) and magnetic phenomena (the behavior of a compass needle). This discovery was an important step towards the development of a unified theory of electromagnetism later in the 19th century. On Ørsted’s account of the event, others had tried to produce an effect by placing a needle perpendicular to a wire and trying to detect a deflection as a current passed through the wire, but no deflection

⁸What follows is primarily based upon the account in Altmann (1992).

had been observed.⁹ Ørsted's contribution was to hypothesize that in fact a needle that was not already perpendicular would become perpendicular in the presence of the current, whereas one that was already perpendicular would not change its state. And to check that hypothesis, Ørsted placed the compass needle differently from how others had done before and observed the result.

Despite his own description of the experiment, it is a matter of some historical dispute whether Ørsted was initially engaged in deliberate experimentation, or observed the effect through mere happenstance.¹⁰ Either way, two things are agreed by all parties: first, prior to 1820, no experimenter had thought to place an electrical current parallel to a magnetic needle—at most, such a current had been placed perpendicular to the needle. So the crucial step was the bringing-about of new, hitherto unforeseen, experimental circumstances. And second, the publication of Ørsted's discovery led to a series of further, clearly deliberate, experiments by others, most notably by the French physicist and mathematician André-Marie Ampère in Paris, who showed, for instance, that when current ran through parallel wires, they would attract or repel, depending on the relative orientation of the currents.¹¹

In the Ørsted and Ampère examples, new experimental circumstances were brought about in the absence of a well-developed theory that predicted what would happen. In other cases, the theorists have already been hard at work, and there is a specific prediction

⁹Ørsted's 1830 description of his experiment notes that attempting "to produce a magnetic effect in the direction of the current [i.e. such that a magnetic needle placed perpendicular to the current would be deflected] . . . had been so often tried in vain" (Ørsted, 1830, p. 575), but does not specify who had been making such attempts.

¹⁰In publishing the first German translation of Ørsted's work in the *Annalen der Physik*, of which he was the editor, Ludwig Wilhelm Gilbert asserted that "that which all researches and efforts had failed to yield, was given by chance to Professor *Oersted* of Copenhagen" ["Was alles Forschen und Bemühen nicht hatte geben wollen, das brachte ein Zufall Hr'n Professor *O e r s t e d* in Kopenhagen" (Gilbert, 1820, p. 292).] This was strenuously denied by Ørsted. More recent discussions vary: Altmann (1992) takes the view that Ørsted's discovery likely was indeed an accident, whereas Martins (2003) suggests an interpretation more consistent with Ørsted's own account of the matter.

¹¹Hacking (c.f. 1983, pp. 161-2)

that has been made as to what the result of the experiment will be. The successful detection of the Higgs boson, under the circumstances that the Standard Model predicted ought to make the Higgs detectable, is perhaps the highest-profile example in physics of a recent experiment that was theory-led in this way. Of course, the dual to this case—where a specific prediction is made and then refuted—is equally informative, if not more so. In yet other cases, the story will be somewhere in between: perhaps the theorists are confident that *something* will happen if we do such-and-such, and even have some good guesses as to what form the something might take, but do not know its specifics. The discovery of the Hall effect is a good example of this kind: some offhand (and rather cryptic) remarks of Maxwell’s suggested that it would be worth seeing what happens if a current-carrying conductor is placed in a magnetic field. Having done so, Hall found that an electric potential is produced—a result that Maxwell did not predict, at least not to any great degree of specificity.¹² This third kind of case falls between the extremes of testing exact predictions and unguided exploration. Indeed, one might doubt whether those extremes are ever instantiated: even in the case of the Higgs boson, the mass of the boson had not been predicted in advance, and was filled in as a result of the experiment.

It should be emphasised that active intervention to bring about novel experimental circumstances is, at least for physics, the source of vastly more scientific information than passive observation of naturally occurring events. The vast bulk of scientific effort in historical and contemporary physics, from the creation of the air-pump to the construction of the LHC, has been devoted to making new things happen. This point is well-made in Hacking’s illuminating discussion of the “creation of phenomena”, where a phenomenon is some noteworthy and discernible event or process that, at least in most cases, “occurs regularly under definite circumstances”.¹³ As Hacking observes, in many ways the natural

¹²This is based on the description of Hall’s work in Hacking (1983), which in turn makes use of Buchwald (1979).

¹³(Hacking, 1983, p. 221)

world is surprisingly empty of phenomena in this sense, at least in anything approaching a “pure” form: most of what we see in the world around us (with the notable exception of celestial motions) is the result of very many laws of nature, cumulatively acting on top of and interfering with one another.

As a result, most phenomena are the result of human ingenuity:

Every time I say that there are only so many phenomena out there in nature to be observed—60, say—someone wisely reminds me that there are more. But even those who construct the longest lists will agree that most of the phenomena of modern physics are manufactured. The phenomena about the species—say the one that a pride of lions hunts by having the male roar and sit at home base while the females chase after and kill scared gazelle—are anecdotes. But the phenomena of physics—the Faraday effect, the Hall effect, the Josephson effect—are the keys that unlock the universe. People made the keys—and perhaps the locks in which they turn.¹⁴

As such, an epistemology of science that neglects the significance of experimental intervention, and takes instead the basic epistemic unit of science to be the (passive) observation, will be drastically limited in its capacity to make sense of scientific practice. But—so we will claim—Humeanism is committed to just such a “spectator theory” of scientific knowledge.

Even in “observational” experiments, the creation of novel circumstances plays a crucial role. Consider the LIGO experiment. Certainly, in some sense, what was being observed (the gravitational radiation produced by a black hole merger) was something that existed naturally, independently of any human action. However, making such an observation required bringing about essentially novel experimental circumstances. That is, in order to “observe” gravitational waves, the LIGO team devised a highly sophisticated device, consisting of lasers reflected between parallel mirrors, such that small variations in the distance between the mirrors could be detected in the interference pattern of the

¹⁴(Hacking, 1983, p. 228)

lasers. That this device works at all depends on certain regularities of both gravitation and optics. Thus, though scientists do not intervene on the distant astrophysical systems that generate gravitational waves, they do intervene to create novel circumstances in our local environment that make those gravitational waves detectable.

4. What Humeans Can Say, But Cannot Do

As we just argued, experimentation involves identifying subjunctive counterfactual conditionals implied by hypothesized laws, and then bringing about situations in the world that realize the antecedents of those conditionals. The goal in doing so is to check whether their consequents obtain. Hypothesized laws imply that if a flame were introduced to a vacuum chamber, it would be exhausted; and so we build a vacuum chamber, light a candle, and see what happens. We expect that if the laws are correct, then the new system we have concocted—never before in nature has there been a candle in a vacuum!—will behave as the laws imply, whereas if it does not do so, we infer that the laws are incorrect or we have misunderstood some aspect of the system. In other cases, like that of Ørsted, the experimenter marches ahead of the theorist, and sets up circumstances for which no hypothesis has yet been formulated (or at least, no hypothesis with any substantial theoretical scaffolding). In both cases, the experimenter is not merely observing regularities in the world. Experimentation is a process of active intervention to probe the limits of those regularities. And it is a practice that has been enormously successful, reliably supporting a wide range of scientific and engineering endeavors.

As we indicated in the introduction, we now wish to argue that the Humean cannot make sense of this practice. The reason is that experimentation presupposes that there are laws, holding independently of whether the experiment is actually conducted, that will constrain the experiment's outcome. In the case where experiments are intended to probe laws hypothesized on the basis of past observation and experiment, the epistemic moti-

vation for conducting those experiments is the expectation that the laws are projectable in the strong sense that if a law is true, and the law implies that under as-yet-unrealized circumstances some event will occur, then bringing about those circumstances will produce the event.¹⁵ But the Humean should deny that laws can play this role. The core reason is that the Humean denies that laws “govern”. We will now present three arguments for this claim, each starting from different Humean commitments. In the next section, we will reflect on their significance.

4.1. Local-to-global inference

The first argument concerns what we will call “local-to-global” inference. This argument is primarily targeted at best-systematizers. This species of Humean, recall, takes the laws to be the axioms of the best systematization of the regularities encountered in the world. As elaborated by Lewis, the best system is specifically that system that best balances simplicity and strength, in the sense of allowing one to derive a maximal number of true general (and specific) statements about the world from a minimal number of simple premises. It is meant to be an objective fact what this best system is, not some reflection of human preferences or cognitive capacities or perspective in the universe.

Most importantly, the best system is supposed to be the best *global* system, i.e., the best system for the totality of the regularities in the world over all of space and time. Presumably this global system will also support successful inferences about local matters of fact.¹⁶ But the Humean has no account of how to reason in the opposite direction: the

¹⁵Humeans should reject other commitments of experimental science as well, such as that experiments should be reliably repeatable.

¹⁶Although even this is not entirely trivial: for example, if some existentially quantified claim were to feature in the best system, then its truth at the global level would not support its truth at the local level. More generally, the desire to be able to make local inferences is bound up with the fact that laws are typically conceived of as universally quantified: indeed, consider the Łos-Tarski theorem, which states that a first-order formula is preserved under taking arbitrary substructures if and only if it is \forall_1 (Hodges, 1993, p. 295). See Hicks (2018) and Dorst (2019) for an argument that the standard best-system criteria (simplicity

best system concerning those regularities constrained to some local region of space and time—say, a single person’s life and environs, or the Earth over the history of modern science—need not be, or have anything to do with, the best system overall. This, in turn, means that the Humean cannot infer, from such and such being an axiom of the best system for the regularities observed to date, that the world will exhibit similar patterns in the future. But this is precisely the sort of projectability that experimental science apparently requires. If there is no reason to expect future events to reflect the regularities observed in the past, and no reason to think that the results of future experiments will bear on yet-later experiments, there is no epistemic motivation for doing experiments at all. At best one will enrich the mosaic by bringing about regularities that would not otherwise have been manifest. But doing so is a way of potentially *changing* the laws, not learning what they are.

One might immediately object that the Humean cannot possibly be committed to this sort of global dependence. Surely the Humean must believe there is *some* relationship between local regularities and global ones, or else they face a far greater epistemological challenge than the one we present. But no: Lewis, at least, directly acknowledges this global character of the view. Recall Lewis’s Big Bad Bug, which he identifies as a conflict between a Humean account of chances and the principal principle. Part of the problem there is that the true chance of an event for a Lewisian Humean depends not only on the frequency with which events of that sort have occurred in the past, but also on future events.¹⁷ And those future events could involve chancemaking patterns that result in chances that differ markedly from the frequencies in any particular region of space and time. The laws have the same character—as, indeed, they must, since for Lewis some of

and strength) do not sufficiently prioritise the need to be able to make local inferences from global laws, and that the Humean should therefore adjust their criteria.

¹⁷There is another problem that is also important in the Bug argument, which is the undermining problem: “there is some present chance that events would go in such a way as to complete a chancemaking pattern that would make the present chances different from what they actually are” (Lewis, 1994, p. 482).

the laws involve specifying chances of various events, such as radioactive decay.

Another possible response is to grant that the Humean cannot reliably infer the future from the past, but reply that this is just the problem of induction. Everyone has that problem. But as Bhogal (2021) has argued—contra Beebee (2011)—the problem of induction is worse (far worse!) for the Humean. One way to put the point is that the problem of induction we all face is an epistemic one: one cannot infer, with certainty, what future events will be like based on patterns in past events. Regularities we have observed before may have depended on unknown contingencies, or inferred laws may have been merely apparent, etc. But for the Humean, the problem is a metaphysical one. They positively deny that the world has the sort of structure—the “glue”, as Bhogal puts it—needed to ensure that future events *are* like past ones. The problem is not merely that we could be *wrong* about the laws as inferred from past events, as on the traditional problem, but that in fact the laws did not figure in bringing about the arrangement of those events—they did not govern them—and likewise they will not figure in the arrangement of future events. From a God’s eye view, there will be a best system that strongly and simply describes these events, but the Humean ideology blocks the use of those laws in inferences.

At this point, we encounter a common Humean refrain, which will come up again in the sections that follow. Like any self-respecting empiricist, Humeans can speak with the masses: there are claims that Humeans *can* make that sound very close to the sorts of claims others make, including experimental scientists. But on reflection, the Humean versions of these claims do not mean what they seem to, in a way that undermines the epistemic motivation for experimental science. In the present case, the Humean will say that they certainly *can* make predictions for future experiments, in the sense that if some law L is true (in the Humean sense), and L implies that some event will occur given some experimental set up, then that event will occur. Laws can play the required logical role in claims about future experimental outcomes, contingent on the truth of the laws.

This is true. But it is important that the implication involved here is *material* implication. Recall that what it means for the law to be true for the Humean is that the implications of the law are, in fact, reflected in the pattern of regularities in the world. So the implied outcomes of the experiment are also reflected in the actual regularities. This reasoning reduces to the following: if the experiment in fact has such and such outcome, then the experiment has such and such outcome. But of course, this does nothing to assuage the worry. On the one hand, the experimenter has no reason to think laws based on their local best system are the right antecedents in these inferences, because the inference only works for the global laws. And on the other hand, the sort of implication needed to make sense of experimental practice is *causal* implication, not material implication. The experimenter needs to reasonably believe that by bringing about such and such circumstances, ones which otherwise would not have obtained, certain events will occur. Otherwise, experiment would not play the right epistemic role, as a way of probing the laws. If Humeanism is correct, one can *talk* about the sorts of counterfactual, nomic dependencies apparently involved in experiment, but there is no point in ever *doing* experiments.

Before proceeding, we observe that the failure of local-to-global inference about laws is a problem for many aspects of scientific methodology, not just experiment. Scientists very often suppose that locally discovered laws apply elsewhere in the world, and draw inferences—often very successful ones—on the basis of that assumption. To give just two examples, it is standard in cosmology (and astrophysics) to assume that the same laws that we have identified from observing local phenomena in the Milky Way also apply elsewhere in the universe (Smeenk and Weatherall, Forthcoming): for instance, that distant galaxies are subject to the same gravitational influences as our own, that hydrogen burns in distant stars just as in our sun, and that in the early universe, the same sorts of particle interactions occurred as the ones we find in our particle accelerators

on earth. None of these inferences should be reliable or even reasonable if Humeanism is true. Similarly, paleontologists routinely assume that ancient animals were subject to the same kinds of constraints, including lawlike ones concerning, for instance, how efficiently blood can pump and how fast nerve signals can travel, as modern organisms, and thereby draw inferences from skeletal remains about how various systems in living creatures must have been (Currie, 2018). Again, such inferences would not be reasonable if Humeanism is true.

We can, of course, guess, based on our current best system, what patterns will be exhibited going forward. We might even have credences that reflect our expectations about the global best system. A Humean could even be a kind of Bayesian, updating their credences that such and such is the best system in light of new evidence.¹⁸ But what are these about? For a Humean, at least on a best systems view, they are credences that such and such system really is the strongest and simplest systematization of the occurrent events over all space and time. They are *not* credences about what the best systematization would be if some local changes to the mosaic were brought about, especially if they were changes intended to create novel situations. In other words, the best system concerns the actual mosaic, not the mosaic one might seek to bring about in an experiment. The Humean's running credences about the best system have nothing to say about that—for all they know, such interventions could change the best system, even from a global perspective.

¹⁸But even here there are complications. As we noted above, the conditionals involved in inferring the probability of some evidence from a hypothesis are material conditionals. Moreover, to calculate the likelihood that you would see some evidence given that such and such system depends not only on whether the system predicts that evidence, but also how strong and simple you expect the best system to be, whether it will be exceptionless, etc. These are not the sorts of considerations that usually go into Bayesian analyses in science, and so again, we find a mismatch between scientific practice and Humeanism. But we will not develop this point further.

4.2. Occurrent-to-nomic inference

For our second argument, we move to the question of how the Humean is to infer the laws from the occurrent facts. In doing so, we leave behind the best system account and focus on a different aspect of Humeanism, which is the supervenience thesis: two worlds that agree in all of their occurrent facts must also agree in their laws. Recall that this is what Earman (1984) calls the Empiricist Loyalty Test. We claim that the empiricist who passes this test does not get to say that experimental knowledge has a justified role to play in science.

A famous example due to Tooley (1977) can be adapted to illustrate the point. Suppose a world that is very much like our own, except that it ends suddenly and inexplicably at some future time. (For all we know, our own world is precisely like this!) Suppose further that, in this world, there are three species of particles, called rock, paper, and scissor. Rock and paper particles have interacted many times during the history of this world, as have paper and scissor particles, and those interactions exhibit certain regularities. Every time a rock particle meets a paper particle, the rock particle is annihilated and the paper particle gains its mass and energy. Scientists in that world come to say “paper beats rock” to describe this well-documented phenomenon. Meanwhile, every time a paper particle meets a scissor particle, the paper particle is annihilated and the scissor particle gains its mass and energy: scissor beats paper. But by whatever twist of fate, rock and scissor particles never meet before the world ends.

What should the Humean say about the laws of this world? *Prima facie*, they should say there is no law at all governing the interactions of rock and scissor. Temporarily adopting the best system approach, this is because including such a law among the axioms of the system would decrease simplicity with no increase in strength, at least as far as deriving the regularities in this world are concerned. But even if one does not accept the best system account, whatever the Humean might say here, it is clear that there

are two distinct possible worlds, both identical with the world described thus far up to the end of that world, but which do not end and have different Humean laws. In one of those worlds, rock and scissor eventually do meet, and rock beats scissor. In the other world, rock and scissor meet and scissor beats rock. These worlds differ in their occurrent facts, and both differ from the first world, so all three can (and do) have different laws according to the Humean. So plausibly, the Humean should maintain that in the world that *does* end, there cannot be any fact about which of the two extensions would have been the “correct” one, from a nomic point of view, because the restriction of these two worlds to their initial segments yields two worlds that are identical in all of their local matters of fact.¹⁹

Now, suppose an enterprising experimenter, Lucy, living in one of these worlds, becomes curious about what happens when rock and scissor particles meet. She secures funding from her world’s NSF to isolate rock and scissor particles and bring them together in her lab. She does the work to prepare the experiment. And then the day comes, and just before she can bring the particles together, the world ends. Her experiment never occurs. Alas, she was in the first of the worlds we described. One immediate reaction might be that this is an *epistemic* tragedy: Lucy never learned what rock and scissor particles do when they meet. But for the Humean who denies the existence of the rock-scissor law, the tragedy is instead metaphysical. In her world, there was no law about rock and scissor particles, no regularity unobserved, because in fact they never met. There was nothing for Lucy to have learned in the first place.

Now observe that it is not necessary for the world to have ended to get the same result.

¹⁹As we will discuss shortly, this reasoning is a little quick: the Humean may have the resources to argue that in the truncated world, there is in fact a law concerning the rock-scissor interaction. However, in order to determine whether it is worthwhile their doing so, we should consider what happens if they take this more straightforward path and deny that there is such a law. So that will be our first concern in this section: to show that this denial has untenable consequences for the Humean, especially with regards to their analysis of scientific practice.

There are any number of reasons why Lucy may not have completed her experiment. She could have been distracted on the way to work that day and forgotten all about the experiment, or she could have decided that morning to quit science and sail around the world. Whatever the case may be, suppose the experiment is never done, and as a result rock and scissor never meet. Once again, we are in a situation where there is no law about what happens when rock and scissor meet. The question that motivated Lucy makes no sense, in retrospect, because there is no law for her experiment to have probed.²⁰

So supposing that there is no rock-scissor law gets the Humean into trouble. But, the Humean might retort, that supposition rests upon an overly simplified picture of science. A more realistic model—one that is closer to the physics we actually have—is one in which the interactions between the different kinds of particles are not fully independent of one another; rather, those interactions are all embedded in some kind of unifying framework. On this proposal, the nature of the rock-scissor interaction does not float free of the other scissor-paper and paper-rock interactions in the manner that the original thought experiment proposes. As a result, the best system of the world up to the point at which Lucy's experiment is about to start will include a law that describes how rock and scissor particles interact—even despite the absence of any interactions that exhibit this regularity.

Let us suppose, then, that in Lucy's world there is such a law; let us say, a law that rock beats scissor. Even so, it does not resolve the problem. For now the question is: why should Lucy's experiment be supported? Of course, the non-Humean has a natural answer: those previous interactions might give us a hypothesis—even a hypothesis of which we are very confident—as to what will happen in the event of a rock-scissor interaction; but a direct test of that hypothesis is still epistemically valuable, as it might

²⁰A natural thought here might be to appeal to counterfactuals: perhaps if she *had* done her experiment, there would have been a law? We will discuss this idea further in the next section.

reveal our hypothesis to be mistaken. Yet this story is unavailable to the Humean. *Ex hypothesi*, the particle interactions that have taken place previously suffice to determine the (or at least, a) rock-scissor interaction law. And so long as nobody adds a rock-scissor interaction to the Humean mosaic, that law is guaranteed to still hold true.²¹

Moreover, as we will discuss in more detail in the following section, the truth of that law determines that the counterfactual “were rock and scissor to interact, then rock would beat scissor” is true. Hence, an NSF reviewer who assessed the proposed experiment (and who has been correctly instructed in the truth of Humeanism) may reason as follows: were this experiment to be performed, the outcome would be that rock would beat scissor. But that would merely be further confirmation of a law that is already known. So this experiment lacks value for money, and should not receive funding.²²

Two responses to this line of reasoning need refuting. The first points out that for all Lucy and the NSF bureaucrat know, there have been or will be rock-scissor interactions elsewhere that suffice to fix the law; so the experiment could be justified as a way of learning what that law is. But we have already shown (in §4.1) that it is deeply unclear why the Humean should expect any systematic relationship between the result of a rock-scissor experiment conducted here and the results of rock-scissor interactions elsewhere in the universe. Moreover, as described in §3, we have grounds for thinking that at least some laboratory experiments realise conditions that *never* occur naturally—at least not in the “pure” form that they arise in the lab. Surely those experiments are just as deserving of funding (if not more so).

The second response argues that Lucy’s funding is only in trouble if the physicists in her world have been astute enough to spot that the existing pattern of paper-rock and

²¹For added poignancy: this is arguably the situation that LHC scientists were in when they sought to find the Higgs boson.

²²(Hall, 2015, p. 275) also raises this problem for the Humean. We see here the basis for a powerful argument for why grant agencies should redirect funds from experimental science to philosophy. Funding administrators are encouraged to contact the authors to discuss.

scissor-paper interactions are best systematized such that rock beats scissor. If that is not the case, then it may be true that rock would beat scissor, without that counterfactual being known by the NSF. But this response mislocates the problem: it turns what is clearly an experimental gap into a theoretical one. As a result, it still cannot explain why the NSF should fund Lucy rather than mathematical physicists; either would be equally good ways of finding out what the rock-scissor law is, and the latter would require much less outlay.²³

Finally, consider the case where Lucy *does* perform the experiment. We contend that even here there is a problem for the Humean! In this case, we may suppose, there is indeed a law that describes rock-scissor interactions, and that Lucy learns what it is. (One might worry that a single experiment is not sufficient to ground a Humean law. To avert this concern, let us suppose that Lucy's successful experiment establishes a new paradigm of rock-scissor testing, and so many other rock-scissor interactions are performed by her fellow-scientists—with results that agree with Lucy's.) The problem is that the kind of knowledge Lucy has gained is of the wrong sort. She has not gained knowledge of a fact that exists independently of her inquiry, but rather of a fact that depends, in part, on the fact that that enquiry was performed. We might say that Lucy has a form of *maker's knowledge*: the knowledge that one possesses of something in virtue of having produced it. Maimonides gives a classic description of the difference between maker's knowledge and empirical knowledge:

There is a great difference between the knowledge which the producer of a thing possesses concerning it and the knowledge which other persons possess concerning the same thing. Suppose a thing is produced in accordance with the knowledge of the producer. The producer was then guided by his knowledge in the act of producing the thing. Other people, however, who examine this

²³After all, as the old joke goes, mathematicians (who only need pencils, paper, and a waste-paper basket) are a lot cheaper than experimental physicists—even if they are dearer than philosophers (who only need pencils and paper).

work and acquire a knowledge of the whole of it, depend for that knowledge on the work itself. For instance an artisan makes a box in which weights move with the running of water and thus indicate how many hours have passed [...] His knowledge is not the result of observing the movements as they are actually going on; but on the contrary, the movements are produced in accordance with his knowledge.²⁴

Yet this seems to radically mischaracterise the nature of scientific inquiry. As Langton puts it in her discussion of maker's knowledge, "The difference between a mirror and a blueprint is a difference in direction of fit: a mirror conforms to fit the world; the world conforms to fit a blueprint."²⁵ And knowledge of the natural world is, at least for non-divine knowers such as ourselves, usually taken as the paradigmatic example of knowledge that is a mirror not a blueprint.

4.3. Actual-to-counterfactual-to-actual inference

Finally, we turn to the third argument, which in some ways brings the previous two together. Consider, again, the reasoning that motivates Lucy. To her knowledge, no rock and scissor particles have ever met in the history of the universe, and she wants to know what would happen if they did. To assess that, she decides to bring these two particles together to observe the outcome. Her purpose in doing so is to learn about this final type of particle interaction. If rock beats scissor in the experiment, she would hypothesize that that was what would happen in future experiments (which she would then run, to check); and *mutatis mutandis*, if scissor beat rock. Moreover, she would conclude that had rock and scissor particles met in the past, rock would have beaten scissor (or vice versa) as well.

We claim that this reasoning, which is ubiquitous in the sciences, involves subjunctive conditionals and inferences about them that the Humean cannot handle adequately. Lucy

²⁴Maimonides (1946), quoted in (Langton, 2009, p. 303)

²⁵(Langton, 2009, p. 299)

wonders about the truth value of a certain counterfactual conditional—say, “were it the case that rock and scissor met, then rock would beat scissor”—and then reasons that there is an experiment she could perform whose outcome would be informative about the laws, and that there is a matter of fact about what that outcome would be (even if she never actually performs the experiment). But (we claim) the Humean cannot make sense of her reasoning about this counterfactual—at least not in a way that coheres with her actions as an experimenter. So just as the Humean struggles with the inferential connections between the local and the global, and between the occurrent and the nomic, the Humean struggles with the inferential connection between the actual and the counterfactual.

This might seem like a surprise. After all, the Humean has a perfectly respectable account of counterfactuals, thanks to Lewis. Indeed, this is one of the great virtues of the view. For our purposes, it will be sufficient to work with a greatly simplified version of Lewis’ semantics, according to which a counterfactual of the form “were it the case that P , then it would be the case that Q ” is true at a world w if and only if at all the worlds “closest” to w in which P is true, Q is also true. Closeness is determined by the “similarity ordering” for w : a centred preorder \lesssim_w on the set of possible worlds, i.e. a reflexive and transitive binary relation such that for all $v \in W$, if $v \lesssim_w w$ then $v = w$.²⁶ This is, intuitively, the relation of *comparative similarity to w* : to say that $v \lesssim_w u$ is to say that v is more like w than u is. Note that we have assumed that there is such a thing as *the* set of worlds closest to w in which P is true; assuming this is true for all w and P means making what is known as the *Limit Assumption*.²⁷

We will also work with Lewis’ own account of what determines similarity between worlds. He says the following:

²⁶Lewis’ treatment uses a partial order rather than a preorder; the advantage of a preorder is that we can distinguish between ties ($v \lesssim_w u$ and $u \lesssim_w v$) and incomparabilities ($v \not\lesssim_w u$ and $u \not\lesssim_w v$).

²⁷For more comprehensive discussion of the semantics of counterfactuals, see Lewis (1973, 1981) or Starr (2022).

[...] a similarity relation that combines with [the semantics given above for counterfactuals] to give the correct truth conditions for counterfactuals [...] must be governed by the following system of weights or priorities.

1. It is of the first importance to avoid big, widespread, diverse violations of law.
2. It is of the second importance to maximize the spatio-temporal region throughout which perfect match of particular fact prevails.
3. It is of the third importance to avoid even small, localized, simple violations of law.
4. It is of little or no importance to secure approximate similarity of particular fact, even in matters that concern us greatly.²⁸

This is how, on Lewis' account, the connection between counterfactuals and laws of nature comes about: counterfactuals are evaluated by considering what happens at the nearest possible worlds, and part of what makes a world nearby is that it should do its best to satisfy the laws of nature.

So what of the counterfactual conditionals used in Lucy's reasoning? In the first instance, suppose that Lucy does not actually perform the experiment, for whatever reason; and suppose that no one else does either. Suppose, moreover, that—as we initially supposed in the previous section—the best system of Lucy's world does *not* contain a law specifying what happens in the case of a rock-scissor interaction. We presume her reasoning is still cogent even in that case, since she may well have done the experiment, and since there are many experiments in the actual world that never get performed, but could have and which one thinks would have had determinate results. This much seems to be borne out by Lewis's account. Consider, for instance, the counterfactual “if Lucy conducted her experiment, rock would beat scissor or scissor would beat rock”. Plausibly, even with the above supposition, it *is* a law that when two particles meet, one beats the other; and so one would expect that in all of the worlds closest to Lucy's, if rock did

²⁸(Lewis, 1979, p. 472)

meet scissor, one would beat the other. Thus, it seems that if Lucy had conducted her experiment, she would have learned an answer to her question, just as she expects.

There are several subtleties to this analysis, however. First, although this counterfactual works out the way one might expect, if one asks the question in a more fine-grained way, the conditional comes out very differently. For instance, consider the counterfactual “if Lucy had conducted her experiment, rock would have beaten scissor”. In the absence of a law specifying what will happen in this interaction, it would seem we have a tie regarding whether rock-beats-scissor or scissor-beats-rock worlds are closer to Lucy’s. It follows, at least on Lewis’ account, that it is false that rock would have beaten scissor, and also false that scissor would have beaten rock. This is odd: it suggests that while it is true that Lucy would get an answer to her question, it is not because there *is* any objective or determinate fact of the matter that she happens not to know. So the fact that the first counterfactual comes out true does not support the reasoning Lucy does based on that counterfactual.

One might worry that we are mislocating the problem here: surely the issue just that Lewis’ account of counterfactuals denies the validity of Conditional Excluded Middle (CEM)?²⁹ No. The basic problem is that from Lucy’s world’s perspective, there is nothing to distinguish the world where rock beats scissor from the world where scissor beats rock. One can of course insist nevertheless that one of the relevant counterfactuals is true and the other is false; but the problem is then how to get such an asymmetry of truth-values out from the Humean’s symmetric metaphysical structure.

To see the second subtlety, consider counterfactuals whose consequents are not categorical statements but rather are themselves nomic or counterfactual claims. For example, Lucy is likely to be interested in counterfactuals such as “had the experiment been per-

²⁹CEM is the principle that for any P and Q , the following is true: “were it the case that P it would be the case that Q , or were it the case that P , it would not be the case that Q ”. See (Lewis, 1973, pp. 79–80) for discussion.

formed, the result would have reflected the laws of nature”, or “had the experiment been performed, then had it been performed again, the result would have been the same both times”. The Humean, we claim, faces a dilemma with respect to such claims. The two horns of the dilemma concern what the term “the laws of nature” should take as its referent in the above counterfactual. Should it refer to the *actual* laws of nature, i.e. those that hold in Lucy’s original world? Or should it refer to the *counterfactual* laws of nature, i.e. those that hold in the world where the counterfactual is being evaluated? In other words, should “the laws of nature here” be taken as a rigid or a non-rigid designator?

If we grasp the first horn, then we obtain an unsatisfactory treatment of the two counterfactuals just introduced. First, it is false that had the experiment been performed, the result would have reflected the laws of nature: there is no actual law of nature that fixes what happens when the rock and scissor particles interact, and so neither result is reflective of such a law. Second, it is at least not clear that had the experiment been performed, then had it been performed again, the second performance would have yielded the same result. From (say) the rock-beats-scissor world, there are two relevant worlds in which the experiment is repeated: one where rock beats scissor again, and one where scissor beats rock. Now, in the absence of an actual law about whether rock beats scissor, one could claim that these two worlds are just as close to the rock-beats-scissor world—and hence the nested counterfactual comes out false. In effect, the nested counterfactual receives the same treatment as the clearly false counterfactual “were I to toss this coin, then were I to toss it again, the result would be the same”.

There is scope for the Humean to push back here, by claiming that it *is* an (actual) law that the result of particle interactions are consistent over time—made true, presumably, by the consistency of the results of scissor-paper and paper-rock interactions. In that case, we run into a problem similar to that discussed above: it would be true that had the experiment been performed twice it would have yielded the same result both times, but

false that had it been performed twice then rock would have beaten scissor twice, and false that had it been performed twice then scissor would have beaten rock twice.³⁰

Let us consider the other horn of the dilemma, then. Perhaps we should evaluate the inner conditional with respect to the laws that hold at the world(s) picked out by the antecedent of the outer conditional. Let us suppose that at the rock-beats-scissor worlds, the best system does include a law of nature to the effect that rock beats scissor, and similarly for the scissor-beats-rock world. This does yield a satisfactory treatment of our two counterfactuals. Had the experiment been performed, then the result would indeed have reflected the laws of nature—whatever that result was. And had it been performed, then had it been performed again, the similarity metric for evaluating the inner conditional would prioritise preserving that law of nature; hence, the world where rock beats scissor twice is closer to the world where rock beats scissor (with respect to the similarity metric for the rock-beats-scissor world), and the same, *mutatis mutandis*, for the case of scissor beats rock. So had the experiment been performed, then had it been performed again, the result would have been the same.

However, choosing this horn does not yield a satisfactory treatment of such counterfactuals in general. This point is made forcefully by Lange (2009). Lange first notes that we expect the laws to be robust under counterfactual supposition. Indeed, he proposes that counterfactual robustness is both necessary and sufficient for lawhood, in a principle he calls NP (for “Nomic Preservation”): “*m* is a law if and only if *m* would still have held under any counterfactual (or subjunctive) supposition *p* that is logically consistent with all of the laws (taken together).”³¹ This principle, of course, is consistent with Humeanism: indeed, it is precisely in order to secure the counterfactual robustness of laws that Lewis’

³⁰Here, “had the experiment been performed twice...” is shorthand for “had the experiment been performed once, then had it been performed again...”.

³¹(Lange, 2009, p. 13). Strictly speaking, Lange does not endorse this principle, but instead endorses a more nuanced variant; we neglect this for ease of exposition.

system of weights places such importance on avoiding violations of law.

However, Lange goes on to argue that not only should the laws' *truth* be preserved under counterfactual supposition, but also their *lawhood*. Otherwise, he thinks, a kind of explanatory gap arises:

Now according to NP, energy would still have been conserved had p been the case, as long as p is logically consistent with the first-order laws. In the "closest p -world," then, why is it the case that energy is conserved? What is the scientific explanation there for the fact that energy is conserved? Presumably, the closest p -world is like the actual world in that energy is conserved there not by accident, but because a law requires it: energy conservation is a law in that world, too. If it is not the case that had someone tried to build a perpetual-motion machine, energy conservation would still have been a law, then why is it the case that had someone tried to build a perpetual-motion machine, energy would still have been conserved? Without the principle of energy conservation remaining a law under this counterfactual supposition, there is no need for the principle to remain true under that supposition.

However, this stronger kind of counterfactual robustness is not satisfied by Humeanism—at least not if we take talk of "the laws of nature" embedded in counterfactuals to refer to the laws of the counterfactual world rather than the actual world.

To apply this to our running example, let us consider the world in which Lucy's experiment *is* performed; with, let us suppose, the result that rock beats scissor. In this world, it is uncontroversially a law that when rock and scissor particles meet, the former beats the latter. Lange's observation is that we at least have the intuition that in that world, the counterfactual "had rock and scissor never met, it would nevertheless have been a law that rock beats scissor" is true. Moreover, that intuition seems crucial to explaining why Lucy might regard herself as fortunate in having received the funding to perform the experiment; for, she might reflect, had rock and scissor not met, then it would still have been a law that rock beats scissor—but a law that she would not have had the good fortune to discover. If the counterfactual is false, then (as discussed in the previous section) we seem to have no sense in which Lucy has met with an epistemic misfortune.

This horn of the dilemma provides no more satisfactory a treatment of nested counterfactuals than the first horn. Again, in reflecting on her successful rock-scissor experiment, Lucy might consider the counterfactual “suppose rock and scissor were to have never interacted; then, it would still have been true that had they had the chance to interact, rock would have beaten scissor”. This nested counterfactual, we claim, should be regarded by Lucy as true. It expresses, after all, exactly what Lucy has learned from her experiment. Yet for the Humean it is false—at least, it is false for the Humean who takes the second horn of the dilemma, and who thinks that there is no rock-scissor law if there are no rock-scissor interactions.³²

What about the case where there is a law concerning the rock-scissor interaction, even in the absence of such an interaction? Of course, such cases cannot be guaranteed to *always* happen on the Humean account, and any time such a law is lacking, the concerns above will apply. But let us consider this kind of case in its own right. That is, let us once again consider a world in which no rock and scissor particles ever interact—but this time, suppose that it is nevertheless a law that rock beats scissor.

It seems to us that Lucy, in reflecting upon her scientific career, might say something like the following: “You know, it’s a real shame that we were never able to perform that rock-scissor experiment. For had we performed that experiment, the results might have been different to what we expected.” We think Lucy’s claim here is intuitively correct. But, it seems, the Humean must deny this. By construction, the counterfactual, “had the experiment been performed, the result would have been what was expected” would be true—because, of course, what the other experimental data leads one to expect to happen in the counterfactual situation *constitutes*, for the Humean, what would happen in that counterfactual. In other words, the Lewisian theory of counterfactuals, when combined with the Humean account of laws, denies that unperformed experiments might have had

³²Nested counterfactuals are discussed further in Lange (2009); Loew and Jaag (2020); Lange (2022).

surprising results. Again, this makes a mystery of the significance of such experiments.

In response, the Humean might note that it is of course *possible* that the experiment was performed and had the result that scissor beats rock. It is not nomically possible, for it goes against the laws of nature; but it is at least a metaphysical possibility. Is that sufficient? We think not, for reasons related to those we discussed at the end of the previous section. The point is that for Lucy, the following is *prima facie* an epistemically possible world:³³ a world where the experiment is not performed, but it is a law of nature that scissor beats rock (and hence the counterfactual “had the experiment been performed, then scissor would have beaten rock” is true). But on the Humean account, that world is a metaphysical impossibility, and so provided that Lucy knows that it is so, it is not an epistemic possibility either—and its absence is not compensated for by the presence of a metaphysically possible world where scissor does in fact beat rock (and hence where it is a law of nature it should do so). That world cannot be the one that accounts for the fact that for all Lucy knows, the experiment might have had different results—for Lucy knows that the experiment was never performed, and so that world is not an epistemic possibility for her!

The analysis above requires that Lucy recognises the metaphysical impossibility in question; if that is not the case, then it could remain an epistemic possibility for her. Perhaps she does not know that the best system of her world is one that entails that rock beats scissor, or perhaps she does not know that the truth of the counterfactual is determined by that best system. But this is not good enough, as it only permits the epistemic possibility by virtue of an epistemic failure on Lucy’s part. The epistemic possibility is closed off if Lucy is well-informed enough about physics to know what the best system is, and well-informed enough about philosophy to know what something’s being the best system entails for laws and counterfactuals (assuming that Humeanism is

³³That is to say, a world that is consistent with what Lucy knows to be the case.

true). By contrast, we think that Lucy could be as good a physicist and philosopher as may be—whilst still being correct in recognising it as a live epistemic possibility that she might have been surprised by an experiment she never performed.

We have now gone through a large number of permutations of views the Humean might adopt. It is worth taking a moment to step back and consider the moral that we take to arise from consideration of these cases. That moral is simple. Experimental reasoning involves inferences of the following form: if one is uncertain about what would happen in the world under various circumstances, one can learn about it by intervening to bring about those scenarios and observing the result. For that reasoning to make sense, it needs to be the case that the meaning of the counterfactuals the experimenter entertains are related to facts about events that would happen were she to perform the experiment and that those events are informative about the laws in her world—and not what would happen in a different world, possibly with different (additional) laws. Finally, whether the counterfactuals an experimenter contemplates end up being true depends (as we have seen) on details about how one elaborates the Lewisian semantics, and *not* on the sorts of modal facts that she needs them to depend on for her reasoning to go through. This is what the Humean cannot deliver.

5. Dialogues Concerning Natural Philosophy

We have argued that Humeanism is incompatible with experimental science, in the sense that anyone truly committed to Humeanism would find little epistemic value in intervening to change the mosaic. This may seem odd. Many scientists and philosophers of science are attracted to some form of empiricism—or at least the denial of rationalistic, a priori knowledge of the natural world. And Humeanism is advertised as the empiricist account of laws, the only account that passes Earman’s empiricist loyalty test. After all, what empiricist would deny that they have common purpose with Hume? Besides, it

would surely be a surprise if science required a commitment to a metaphysically thick account of laws of nature.

There are several things to say here. The first is to emphasize that Humeanism should be viewed as empiricist in the pejorative sense. This is because it falls prey to a classic empiricist fallacy. What Hume himself compellingly argues for is an epistemic thesis: we cannot know that there are necessary connections between distinct existences.³⁴ This is the problem of induction, which we very happily concede. But there is a significant gap between this epistemic thesis and the metaphysical one that says there are no such necessary connections—the same gap, we would argue, between the observation that our access to the world goes via our perceptual system and the conclusion that all that exists are congeries of ideas. It is true that Humeanism is an especially parsimonious metaphysical view of laws that is apparently compatible with our epistemic situation, but mere parsimony is not compelling here. Neither Hume nor anyone else has ever given a compelling argument that we cannot have good reasons to believe a less parsimonious metaphysical theory over a more parsimonious one if the more parsimonious one leads to bizarre or unacceptable consequences.

How should we respond to the problem of induction, then? We do not have anything very original to say here. Our proposal is just to accept another oft-quoted dictum from Hume and proportion one's beliefs to the evidence. Accepting a broadly Bayesian picture on which we revise our beliefs in response to new evidence, with an appreciation of how scientists actively seek out and assess particular kinds of evidence, suggests an anti-Humean (but still, Humean!) humility about the status of laws of nature. We cannot know that we have discovered a necessary connection, or a relation between universals, or a primitive law of nature, but we equally well cannot know that we have not, particularly in

³⁴Recall note 3. Of course, the arguments we make here do not depend in detail on any particular reading of Hume, and so we will not be refuted by arguments that the metaphysical reading of Hume is apt. Rather, our position is that if that is what Hume meant, so much the worse for Hume!

light of the fact that the nomic relationships posited by our scientific theories have turned out to be projectable over a broad range of cases. In other words, repeated experience with successful prediction and intervention give us increasing confidence that certain experimental practices and reasoning styles are reliable, which in turn provides evidence that the nomic structure of the world is sufficiently robust to support those practices and inferences. This position allows scientists to embrace the empiricist ethos without its excesses.

Some readers will find this attitude dissatisfying, or too dismissive of metaphysics. It certainly denies a privileged, *a priori* status to metaphysics. But we would resist the idea that the position just sketched is anti-metaphysical. Instead, it emphasizes the continuity between metaphysics and scientific practice, including highly applied issues arising in laboratory activities. We see this as an elevation of metaphysics. The central point is that commitment to a metaphysical view should bear on one's actions and expectations just as one's other beliefs do, and the success or failure of those actions and expectations should be weighed when assessing one's metaphysical positions. Some Humeans and other *a priori* metaphysicians will resist this claim, arguing that proper metaphysics occurs at a level that cannot possibly impinge on science. If that is truly their intention, though, we fear it render metaphysics not just irrelevant but meaningless. Consider: if the claim that laws do not govern is intended in such a way that it can have absolutely no bearing on whether one should consistently expect the laws of nature to be accurate with respect to the behavior of natural systems in novel contexts, then the claim means nothing at all.³⁵

This sort of quasi-empirical status for metaphysics has roots in Howard Stein's reading of Newton's metaphysics. For Stein's Newton, the goal of physics is to identify math-

³⁵Here we differ from a line of Humean defense that originates with Loewer (2012, 2019). Loewer argues that Humean laws are able to play the role usually attributed to laws in science, even though they do not provide metaphysical explanations. We think that drawing this distinction makes metaphysics irrelevant in a way we would prefer to resist.

ematical expressions describing uniform dynamical regularities in nature. This should be done, to whatever extent possible, without stipulating a prior metaphysics, lest it artificially constrain the theories one entertains. And so it is that Newton comes to describe gravitation as action at a distance, without prejudice for mechanistic contact-based interactions that many of his contemporaries felt were compelled by *a priori* metaphysical considerations. But then, given a successful mathematical theory, one reintroduces a metaphysical theory, on the grounds that certain structure must be attributed to the world in order for the lawlike relationships expressed by the theory to be coherent. Empirical support for a metaphysics comes via the success of the theory that presupposes it.

Newton's own views on laws are hardly satisfactory: he understood laws as a continuing exertion of God's will in the world. But we take inspiration from his method. The important point is that one should attribute to the world at least as much structure as is needed to make sense of your best science. For Newton, this meant resisting Leibniz's *a priori* arguments for relationism, on the grounds that it was too parsimonious to make sense of rotation in Newtonian dynamics. Likewise, we say, the successful practice of experimental science requires us to posit a certain amount of modal structure. Humeanism is too parsimonious to support the practice, just as relationism is too parsimonious to support Newton's laws.

The forgoing discussion raises an important question. What sort of metaphysics of laws is needed to make sense of scientific practice? We suggest that two conditions must be met. First, the metaphysics must underwrite the claim that the local regularities in nature are exchangeable across regions, in the sense that the regularities exhibited in different regions are all expressions or manifestations of the same system (c.f. Smeenk and Weatherall, Forthcoming, on Aristotle's principle). Call this *modal exchangeability*. If massive bodies tend to approach one another on Earth, or in our solar system, they do so also in other galaxies, in the early universe, and so on. Modal exchangeability is needed

both to infer that locally discovered regularities can be projected to the past and future and to distant regions of space. It is also needed to infer that repeated experiments and observations over time and in different places are probing a single system of regularities. This sort of assumption should be rejected by anyone who claims that laws do not govern, because in that case there is nothing linking the regularities in one region with the regularities elsewhere.³⁶

Second, one needs to assume that whatever regularities are present in nature are robust under small, local perturbations. That is, the true system of regularities concerns not just how the world is, but also how the world would be were some small change to occur in a local region and the consequences of that change were propagated to other regions that the laws require the change to affect. Call this *nommic robustness*. Accepting nomic robustness is to fail Earman's empiricist loyalty test, because it means that there can be two possible universes that are identical with regard to what actually happens, but which differ with regard to what would happen were some small intervention to occur. Nevertheless, we need this principle to make sense of the practice of interventionist experiment, per our arguments above. If we did in fact live in a world that was ambiguous between two sets of laws, we would go out and try to engineer a situation that would tell us which one was right. We would intervene. Whatever our account of laws is, it must make it true that there is a fact of the matter that it makes sense to try to discover.

What accounts of laws fit with this picture? Most metaphysically thick accounts of laws of nature will support these principles. For instance, Armstrong-Tooley-Dretske style accounts, on which laws are necessary relations between universals, would support both principles (insofar as actual candidate laws, such as partial differential equations, can be made to fit into that framework at all), because necessary relations are both projectible

³⁶That is, there is nothing prospectively linking them. For some Humeans, the "best" system may be one that emphasizes similarities across regions (Hicks, 2018; Dorst, 2019; Jaag and Loew, 2020). But that does not license the inference that those same regularities would arise in other, not-yet-observed regions.

and robust (Armstrong, 1983; Tooley, 1977; Dretske, 1977). Causal powers views also apparently support these principles, as long as one takes the powers to be shared among saliently similar entities (Cartwright, 1980; Woodward, 1992). And accounts on which laws of nature are primitive, such as those of Carroll (1994) and Maudlin (2007), would also suffice. Likewise, views on which there are basic modal facts, or modal structure, about what would occur under various non-factual circumstances, such as those of Lange (2009) or Berenstain and Ladyman (2012), could support these principles.

We will not endorse any of those views here. Instead, we will just note that these two principles could themselves be taken to constitute a minimal view, in the broad family of "modal structure" views, according to which there is some objective modal structure in the world, reflecting the possible consequences of local interventions. The idea is that the regularity facts concern not just the actual events, organized or summarized in some way, but also the fact that these are robust regularities that consistently obtain under the same circumstances across time and space. That they are regularities in this sense would support experimental practices, because it would permit an interpretation on which experiments are probing the regularity facts and modal structure by checking how known regularities extend to novel circumstances.

Of course, this view is incompatible with Humean supervenience and it fails Earman's loyalty test. But we suggest that it is nonetheless a broadly empiricist view. The reason is that one can learn, and we suggest we *have* learned, that the same regularities obtain across time and space by prior observation and experiment. In other words, there is an observable metaregularity that our first principle holds in regions we can observe. It is a defensible, though defeasible, inference from there to our first principle. Then, once the first principle is accepted, one can infer that what happens under certain circumstances in one place and time is also what would have occurred under the same circumstances at a different place and time. In other words, one can infer, from principle 1, that by surveying

a heterogeneous sample of actual events across space and time, what the space of possible, but non-actual, events in any given region of space-time would be. Thus, observation of actual events can support inference about the structure of possible events. This reasoning involves assumptions whose validity will rise or fall with the success of the practice—just as Newton’s metaphysics rises or falls with his physics. But it is certainly compatible with the core empiricist commitment that knowledge is grounded in experience.

6. The Final Sequence

Humean metaphysics, like many empiricist-inspired positions, is ultimately motivated by an epistemological worry: our experience of the world consists of only what does happen, and so anything we say about what can happen or what must happen must be inferred from events that actually occur. So far so good. But to resolve this epistemological worry, Humeans ignore the reasoning and practices that scientists adopt, with great success, to draw inferences about nomic structure from those actual events. Instead, they offer a metaphysical view according to which there is nothing to discover about nomic structure aside from how to best summarize or regularize the events that do occur. Put another way, Humeans confuse a difficult problem—but one that scientists have apparently solved—for an impossible one, and so they deflate away structure in the world that in fact we have sound and reliable knowledge about. The result, we have argued, is a fundamental incompatibility between the Humean worldview and successful experimental practice. If Humeanism were true, experimental science would be fruitless and incoherent; experimental science is well-motivated and successful, ergo Humeanism is false.

It is true that we do not merely “observe” nomic or modal structure in the way we observe that the sky is blue or that low atmospheric pressure precedes a storm. But there are many things we know that we do not simply observe, such as that the Sun

is larger than the moon (despite having the same angular size in the sky) or that the half life of Uranium-235 is 700 million years. Instead, we learn these things by complex inferences from what we do observe, along with interventions designed to shape what we observe and facilitate those inferences. Nomic structure is no different in principle than these other complex facts that require theory-mediated observation, intervention, and experimentation. Importantly, though, the practices that have emerged within science for discovering this sort of nomic structure require us to suppose that there is a robust, governing nomic structure out there to discover. For this reason, Humeans have no account of what scientific practice does or why it works.

The basic problem with Humeanism is that it does not have space for scientists as embodied epistemic agents, intervening in the world and influencing the it in a way that both depends on, and is revelatory of, laws that constrain what actions are possible and what the consequences of those actions will be. We do not passively observe or measure a mosaic of occurrent facts or events. We actively participate in those events, and we reason about them in ways that motivate further actions whose outcomes provide justification for yet other inferences and actions.

We note that the absence of intervention from the Humean worldview may cause problems beyond just their account of scientific epistemology. For one thing, the central role of intervention in contemporary accounts of causal inference suggests that Humeans will have significant difficulty in making sense of that practice either. For another, if one follows Hacking (1983) in thinking that intervention and representation are tightly linked, then the Humean has issues with not only the epistemology but the very semantics of science. For Hacking, a key part of using any observational apparatus, such as a microscope, is the manipulation of what is being observed to see how the “image” changes as a result: it is practical intervention (simultaneously with observation) that “creates the ability to distinguish between visible artifacts of the preparation or the instrument, and

the real structure that is seen with the microscope.”³⁷ Without intervention, one cannot establish the connective tissue that links theory to phenomena.

If not Humeanism, then what view of laws should we adopt—especially if one shares the empiricist instincts that motivate the Humean and many working scientists? We have not endorsed an account of laws here, but we do suggest a methodology. We should attribute to the world at least as much structure (and arguably, only that structure) needed to make sense of both our best scientific theories and the practices that have led to their discovery and evidential support. We have suggested two principles that appear to be necessary for any satisfactory account of laws, judged on these grounds: modal exchangeability and nomic robustness. Many accounts of laws seem to underwrite these principles. We also suggested that these principles could themselves be the basis for a minimalist account of laws, following a modal structuralist model. But crucially, Humeanism apparently must either deny these principles, or else deny basic tenets of the view, such as the supervenience claim. And it is for that reason that Humeanism must be fail.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 2419967.

References

- , 1928/1978. Universals of law and of fact, in: Mellor, D.H. (Ed.), *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*. Humanities Press, Atlantic Highlands, NJ, p. 128–133.
- Adlam, E., 2022. Laws of Nature as Constraints. *Foundations of Physics* 52, 28. doi:10.1007/s10701-022-00546-0. <https://doi.org/10.1007/s10701-022-00546-0>.

³⁷(Hacking, 1983, p 191)

- Altmann, S.L., 1992. *Icons and Symmetries*. Oxford University Press, Oxford.
- Armstrong, D., 1983. *What Is a Law of Nature?* Cambridge University Press, Cambridge.
- Bacon, R., 1267. *Opus majus*.
- Beebe, H., 2000. The non-governing conception of laws of nature. *Philosophy and Phenomenological Research* 61, 571–594.
- Beebe, H., 2011. Necessary Connections and the Problem of Induction. *Noûs* 45, 504–527. arXiv: 41330869. <https://www.jstor.org/stable/41330869>.
- Berenstein, N., Ladyman, J., 2012. Ontic structural realism and modality, in: Landry, E., Rickles, D. (Eds.), *Structural Realism: Structure, Object, and Causality*. Springer, Dordrecht, pp. 149–168.
- Bhagal, H., 2021. Induction and the Glue of the World. *Australasian Journal of Philosophy* 99, 319–333. doi:10.1080/00048402.2020.1793788. <https://doi.org/10.1080/00048402.2020.1793788>.
- Bhagal, H., Perry, Z., 2017. What the Humean Should Say About Entanglement. *Noûs* 51, 74–94. doi:10.1111/nous.12095. <https://onlinelibrary.wiley.com/doi/abs/10.1111/nous.12095>.
- Buchwald, J.Z., 1979. The Hall Effect and Maxwellian Electrodynamics in the 1880's. Part I: The Discovery of a New Electric Field. *Centaurus* 23, 51–99. doi:10.1111/j.1600-0498.1979.tb00360.x. <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1600-0498.1979.tb00360.x>.
- Callender, C., 2015. One world, one beable. *Synthese* 192, 3153–3177. doi:10.1007/s11229-014-0582-3. <https://link.springer.com/article/10.1007/s11229-014-0582-3>.
- Carroll, J.W., 1994. *Laws of Nature*. Cambridge University Press, Cambridge.
- Cartwright, N., 1980. Do the laws of physics state the facts? *Pacific Philosophical Quarterly* 61, 75–84.
- Chen, E.K., Goldstein, S., 2022. Governing Without a Fundamental Direction of Time: Minimal Primitivism About Laws of Nature, in: Ben-Menahem, Y. (Ed.), *Rethinking the Concept of Law of Nature*. Springer International Publishing, Cham, pp. 21–64. doi:10.1007/978-3-030-96775-8_2. https://link.springer.com/10.1007/978-3-030-96775-8_2.
- Cohen, J., Callender, C., 2009. A better best system account of lawhood. *Philosophical Studies* 145, 1 – 34.
- Currie, A., 2018. *Rock, Bone, and Ruin: An Optimist's Guide to the Historical Sciences*. MIT Press, Cambridge, MA.
- Dewey, J., 1929. *The Quest for Certainty: A Study of the Relation of Knowledge and Action*. Minton, Balch and Company, New York.
- Dorst, C., 2019. Towards a Best Predictive System Account of Laws of Nature. *The British Journal for the Philosophy of Science* 70, 877–900. doi:10.1093/bjps/axy016.

- <https://www.journals.uchicago.edu/doi/10.1093/bjps/axy016>.
- Dretske, F., 1977. *Laws of Nature*. volume 44.
- Earman, J., 1984. *Laws of Nature: The Empiricist Challenge*, in: Bogdan, R.J. (Ed.), D.M. Armstrong. Springer Netherlands, Dordrecht. Profiles, pp. 191–223.
- Earman, J., Roberts, J.T., 2005a. Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature Part I: Humean Supervenience. *Philosophy and Phenomenological Research* 71, 1–22. doi:10.1111/j.1933-1592.2005.tb00428.x. <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1933-1592.2005.tb00428.x>.
- Earman, J., Roberts, J.T., 2005b. Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature Part II: The Epistemological Argument for Humean Supervenience. *Philosophy and Phenomenological Research* 71, 253–286. doi:10.1111/j.1933-1592.2005.tb00449.x. <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1933-1592.2005.tb00449.x>.
- Emery, N., 2023. The governing conception of laws. *Ergo* 9, 16.
- Esfeld, M., Hubert, M., Lazarovici, D., Dürr, D., 2014. The Ontology of Bohmian Mechanics. *The British Journal for the Philosophy of Science* 65, 773–796. doi:10.1093/bjps/axt019. <http://bjps.oxfordjournals.org/content/65/4/773>.
- Fernandes, A., 2023. Naturalism, functionalism and chance: Not a best fit for the humean, in: Loew, C., Jaag, S., Hicks, M.T. (Eds.), *Humean Laws for Human Agents*. Oxford UP, pp. 86–107.
- Gilbert, 1820. Entdeckung ausgezeichneter Wirkungen des geschlossenen galvanisch-electrischen Kreises auf die Magnetnadel, und der Kraft der galvanischen Electricität zu magnetisiren: Vorbericht. *Annalen der Physik* 6, 291–294.
- Guo, B., 2022. Can humeans be scientific realist? URL: <https://philsci-archive.pitt.edu/22717/>. unpublished manuscript.
- Hacking, I., 1983. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge University Press, Cambridge.
- Hall, N., 2015. Humean Reductionism about Laws of Nature, in: Loewer, B., Schaffer, J. (Eds.), *A Companion to David Lewis*. Wiley-Blackwell, Hoboken, pp. 262–277.
- Hicks, M.T., 2018. Dynamic Humeanism. *The British Journal for the Philosophy of Science* 69, 983–1007. doi:10.1093/bjps/axx006. <https://www.journals.uchicago.edu/doi/full/10.1093/bjps/axx006>.
- Hodges, W., 1993. *Model Theory*. Number 42 in *Encyclopedia of Mathematics and Its Applications*, Cambridge University Press, Cambridge.

- Huggett, N., 2006. The Regularity Account of Relational Spacetime. *Mind* 115, 41–73. doi:10.1093/mind/fzl041. <https://academic.oup.com/mind/article-abstract/115/457/41/1028776/The-Regularity>
- Jaag, S., Loew, C., 2020. Making best systems best for us. *Synthese* 197, 2525–2550. doi:10.1007/s11229-018-1829-1. <https://doi.org/10.1007/s11229-018-1829-1>.
- Lange, M., 2009. *Laws and Lawmakers: Science, Metaphysics, and the Laws of Nature*. Oxford University Press. doi:10.1093/acprof:oso/9780195328134.001.0001. <https://academic.oup.com/book/7001>.
- Lange, M., 2022. Can Pragmatic Humeanism Account for the Counterfactual Invariance of Natural Laws? *The British Journal for the Philosophy of Science* doi:10.1086/723621. <https://www.journals.uchicago.edu/doi/10.1086/723621>.
- Langton, R., 2009. Speaker's Freedom and Maker's Knowledge, in: *Sexual Solipsism: Philosophical Essays on Pornography and Objectification*. Oxford University Press, Oxford, pp. 289–310.
- Lewis, D., 1973. *Counterfactuals*. Harvard University Press, Cambridge, MA.
- Lewis, D., 1979. Counterfactual Dependence and Time's Arrow. *Noûs* 13, 455–476. doi:10.2307/2215339. <https://www.jstor.org/stable/2215339>.
- Lewis, D., 1981. Ordering semantics and premise semantics for counterfactuals. *Journal of Philosophical Logic* 10, 217–234.
- Lewis, D., 1994. Humean Supervenience Debugged. *Mind* 103, 473–490. doi:10.1093/mind/103.412.473, arXiv:2254396. <http://www.jstor.org/stable/2254396>.
- Loew, C., Jaag, S., 2020. Humean Laws and (Nested) Counterfactuals. *The Philosophical Quarterly* 70, 93–113. doi:10.1093/pq/pqz037. <https://doi.org/10.1093/pq/pqz037>.
- Loewer, B., 2012. Two accounts of laws and time. *Philosophical Studies* 160, 115–137.
- Loewer, B., 2019. Humean laws and explanation. *Principia: An International Journal of Epistemology* 23, 373–385.
- Loewer, B., 2021. The package deal account of laws and properties (PDA). *Synthese* 199, 1065–1089. doi:10.1007/s11229-020-02765-2. <https://doi.org/10.1007/s11229-020-02765-2>.
- Maimonides, M., 1946. *The Guide for the Perplexed*. Pardes Publishing House, New York.
- Martins, R.D.A., 2003. Resistance to the Discovery of Electromagnetism: Ørsted and the Symmetry of the Magnetic Field, in: *Volta and the History of Electricity*. Hoepli, Milan, pp. 245–265.
- Maudlin, T., 2007. *The Metaphysics within Physics*. Oxford University Press, Oxford.
- Mill, J.S., 1843. *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*. John W. Parker, London.

- Miller, E., 2014. Quantum Entanglement, Bohmian Mechanics, and Humean Supervenience. *Australasian Journal of Philosophy* 92, 567–583. doi:10.1080/00048402.2013.832786. <http://dx.doi.org/10.1080/00048402.2013.832786>.
- Ørsted, H.C., 1830. Thermoelectricity, in: Brewster, D. (Ed.), *The Edinburgh Encyclopædia*. Blackwood, Edinburgh. volume 18, pp. 573–589.
- Sebens, C.T., 2024. How do Laws Produce the Future? <http://arxiv.org/abs/2407.00210>. arXiv:2407.00210.
- Smeenk, C., Weatherall, J.O., Forthcoming. *The Aims and Structure of Cosmological Theory*. Oxford University Press, Oxford.
- Starr, W., 2022. Counterfactuals, in: Zalta, E.N., Nodelman, U. (Eds.), *The Stanford Encyclopedia of Philosophy*. winter 2022 ed.. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2022/entries/counterfactuals/>.
- Tooley, M., 1977. The Nature of Laws. *Canadian Journal of Philosophy* 7, 667–698. arXiv:40230714. <https://www.jstor.org/stable/40230714>.
- Woodward, J., 1992. Realism about laws. *Erkenntnis* 36, 181–218.