

Humean Reductionism About Laws Of Nature

Ned Hall

§0 Introduction

By far the most central and important question about laws of nature is this: Are they mere patterns in the phenomena (patterns that are in some way salient, to be sure—but still, nothing more than patterns)? Or are they something more, something that somehow governs or constrains those phenomena? Disagreement over this issue constitutes the Schism in contemporary philosophical work on laws of nature. Here, I will investigate the prospects for an important position that falls under the “mere patterns” approach: what, for reasons that will emerge, I call “Humean reductionism” about laws of nature. I will review the most prominent arguments against this position from the literature, and add some of my own that, I think, are more effective. All the same, “investigate” remains the operative word. The value of the various anti-Humean arguments lies not so much in their power to refute, as in their power to force a much sharper and more sophisticated articulation of a Humean position capable of resisting them. So, while I immodestly claim that my own anti-Humean arguments are the best out there, I offer them with only the modest aim of improving our understanding of the Schism.

§1 Preliminaries

I’m going to begin by assuming, without argument, a basic distinction between fundamental laws and the so-called “laws” of the special sciences. I will only be concerned with the former. It is the job of fundamental physics, and fundamental physics alone, to uncover the fundamental laws. These laws describe the way in which complete physical states of the world generate successive physical states—or, if you prefer, the way in which complete physical states are followed and preceded by complete physical states.¹ I take it that the fundamental laws of our world fix a distinction between worlds that are nomologically possible, relative to our own, and worlds that are merely metaphysically possible. There is plenty of room for disagreement about how this fixing happens: you might analyze nomological possibility by saying that a world is nomologically possible relative to our own iff it has the same laws as our world; or you might say that it is nomologically possible iff what goes on in it merely conforms to the laws of our world; or you might hold that

¹ What is a complete physical state, in general relativity? It will do to let it be the complete physical state of one space-like hypersurface, treating the laws now as constraining how the rest of space-time can be, compatible with a specific state for some such hypersurface.

it is the relation of nomological possibility that has metaphysical priority, so that the locution “it is a law that P” is to be analyzed as “P is true in every nomologically possible world”. Mostly, these niceties won’t matter as we consider rival philosophical accounts of laws of nature in what follows; all that will matter is that an account somehow or other deliver the needed distinction between nomological possibility and impossibility.

By fixing the structure of the nomological possibilities, the fundamental laws, together with facts about what actually happens, fix the counterfactual structure of our world: that is, they fix truth values for all counterfactuals that (1) have antecedents compatible with the fundamental laws, and (2) are determinate enough to *have* truth values.²

By fixing the counterfactual structure of the world, the laws (again, together with the facts about what actually happens) fix all the facts about what causes what, and what causal dispositions are had by what. It is at this point that the special sciences get to have a say: for what they aim to discover are not laws in anything like the sense at work in fundamental physics, but rather “laws” that are really a certain kind of causal generalization (see for example Hall 2005).

It will be useful to have an example of some fundamental laws, to turn to as a kind of test case for Humean reductionism. Current physics is too messy; so we will rewind about two centuries, and consider Newtonian particle mechanics. That is, we will assume that our world is wholly composed of some finite number of indestructible point-particles, moving in continuous trajectories through space, and each possessed of an unchanging value for mass and charge. As a first pass, we will take the fundamental laws governing these motions to be Newton’s laws of motion and the appropriate force laws. But that is just a rough guide, for we should remember that there are perfectly legitimate formulations of Newtonian point mechanics that make no mention of force—e.g., the Hamiltonian formulation. It is contentious to assume that these are formulations of a strictly different theory than Newton’s (e.g., because Newton’s theory is ontologically committed to forces, and the Hamiltonian formulation is not). What really matters for our purposes is that the laws of our world (as we are imagining it to be) pick out a range of nomologically possible total histories of particle behavior. And for that aim, any formulation will do.

²In other words, we are not expecting the fundamental laws, even with the help of all the relevant historical facts, to settle such difficult issues as what exactly Caesar would have done, had he been in charge in Korea.

A very handy way to map this set of nomologically possible histories is as follows: We will suppose that the fundamental laws delimit a certain set of nomologically possible initial conditions, each of which is given by specifying the number of particles at the “initial time”, along with their masses, charges, positions, and velocities. (Of course, we do not need to be committed to the claim that there is an initial time. It will be enough to assume that the laws delimit a set of possible instantaneous physical states for the world as a whole to be in, at whatever time. But we will stick with the simplifying assumption that there is an initial time.) Second, the laws specify how each such initial condition would evolve forward in time. We can summarize by saying that the fundamental laws “factor” the set of nomologically possible histories of particle behavior into an initial conditions hypothesis (ICH) and a dynamical hypothesis (DH).

Some comments. First, we will assume, in the case of our Newtonian example, that the dynamical hypothesis is deterministic, so that no two nomologically possible histories perfectly agree on particle behavior (i.e., on their distribution of particle positions, masses, and charges) up to some time without agreeing for all time. Later, we will occasionally tweak the example so as to relax this assumption.³

Second, I said that ICH + DH yields the set of nomologically possible histories of particle behavior—not the set of nomologically possible *worlds*. One reason for holding back from the stronger characterization is that, according to most anti-reductionists about laws of nature, more is needed to specify the nomologically possible worlds. Minimally, one needs to say that they are worlds that not only exhibit nomologically possible histories of particle behavior, but that are also governed by the *same laws* as the actual world. And that is a substantive addition—one rejected by some (though not all) versions of Humean reductionism.

Third, it is worth noting that the breadth of the ICH makes for a certain kind of strength. That is, it is, other things equal, a point in favor of a physical theory that it recognizes a wide range of nomologically possible initial conditions. To see why, compare for example Keplerian and Newtonian accounts of the solar system. Granted that the Newtonian account is much more empirically accurate, it is also, from the standpoint of scientific investigation, better in a distinct sense: for it allows us to answer questions not merely about how the elements of the solar system did,

³Note that the assumption of determinism is in fact problematic, although not for reasons that will matter here. Xia 1992, for example, describes a fiendishly clever five-particle Newtonian world in which the five particles, solely as a result of gravitational forces, accelerate to infinity in finite time. Such a world is clearly not backwards deterministic, and time-reversibility of the dynamics shows that it is not forwards deterministic, either. So it is actually a somewhat delicate matter what extra constraints to impose on the ICH in order to get determinism.

do, and will behave, but also about how they would have behaved under alternative physical conditions. We will come back to the significance of this kind of strength later, in §6; we will see that the Humean reductionist has a very difficult time accounting for its importance in physical theorizing.

§2 The Elements of Humean Reductionism

A standard way to set up Humean reductionism is as follows: First, one posits a basic distinction between modal and nonmodal facts about the world (or sometimes, between nomic and non-nomic facts). The core thesis of Humean reductionism is then stated as a supervenience claim: all facts—and in particular, all facts about the fundamental laws—supervene on the totality of nonmodal facts, so that no two worlds different with respect to their fundamental laws without differing somehow with respect to their nonmodal facts. Some authors lose their nerve, and treat this supervenience claim as itself contingent (more on this in a moment), saying only that among “worlds like ours”, no two differ with respect to their laws without differing with respect to their nonmodal facts. (Obviously, some criterion needs to be added for distinguishing which of the worlds are, in the relevant sense, like ours.) Whatever modal status is assigned to the supervenience claim, the aim then becomes to construct a philosophical account of laws that respects it. Given the widespread popularity of the supervenience claim, then, it comes as no surprise that the idea that laws are mere patterns in the phenomena should prove so attractive.

I think, however, that this way of setting things up invites some confusion. First, supervenience is too weak a relation: just notice that it is not asymmetric, whereas reduction—which is clearly what is intended—is. Second, one might reasonably think that, if facts about laws reduce to nonmodal facts, then they must *themselves* be nonmodal facts. If so, the claim that facts about the laws reduce to the totality of nonmodal facts (which, being a totality, includes them) becomes trivial. Or, more carefully, to make the claim *nontrivial* requires a substantive, nontrivial characterization of the modal/nonmodal distinction.

A third and quite distinct worry is this: Suppose you think that such magnitudes as mass and charge are ineluctably bound up with the laws, in the sense that nothing could count as mass unless it were governed by the law of universal gravitation, and likewise nothing could count as charge unless it were governed by Coulomb’s law. (We will henceforth label this position “essentialism”.) Then you will almost certainly hold that there could be no difference between two worlds with respect to the laws without some difference with respect to what fundamental physical magnitudes are instantiated in those worlds. Now, a Humean will quickly point out that your adherence to *that* supervenience thesis does not put you in his camp; for the

subvening facts in question—facts about the pattern of instantiation of fundamental physical magnitudes—are not, on your view, *nonmodal* facts. Maybe so, but again, it would be nice to have a characterization of the modal/nonmodal distinction substantive and clear enough to explain *why* it is so.

So let me suggest a different way of proceeding. I am going to start with the assumption, widespread among reductionists and anti-reductionists alike, that there is such a thing as the fundamental ontological structure of the world: facts about the world to which all other facts reduce.⁴ What the reductionist ought to do is to put forth a specific doctrine about the nature of this fundamental ontological structure. Lewis was admirably clear on this point. Drawing inspiration from Hume, he took the fundamental ontological structure to consist in the pattern of instantiation by space-time points of perfectly natural monadic properties, together with the facts about the spatiotemporal relations among those points.⁵ He could then say that these facts are nonmodal in the precise sense that the metaphysical possibilities involving the perfectly natural properties and relations respect a principle of recombination, according to which, for any way of distributing perfectly natural properties and relations among particulars, there is a possible world in which they are so distributed.

⁴ If you are at a loss as to what “reduce to” means, and find insufficiently illuminating the observation that it is a logically stronger relation than supervenience, then there is probably little that I can do to help you. But what little I *can* do is contained in the following observation: There is a characteristic kind of question we ask when doing applied metaphysics. “What is it for such-and-such a fact to obtain?” For example, we might ask, “What is it for something to persist through time?” A four-dimensionalist will give one answer (it is for it to be composed of momentary time-slices, such that for each of the times in question, one of the constituent time-slices exists at that time); a three-dimensionalist will give a different answer (it is for it to be “wholly present” at each moment at which it exists, whatever exactly that means). Or we might ask, “What is it for someone to act freely?” A typical compatibilist will say that it is for her actions to be caused in the right sort of way by internal psychological states (never mind whether those psychological states themselves have fully sufficient causes in the remote past); a typical incompatibilist will say that it is for her to exercise a kind of metaphysically primitive causal agency. I take it that a correct answer to a question of the form, “What is it for such-and-such a fact to obtain?” must display other facts to which the fact in question *reduces*. One can then repeat the question, targeting these other facts. (What is it for a time-slice to exist at a time? Answer: it is for the contemporaneous space-time points that constitute it to exist.) The fundamental ontological structure of the world is what one arrives at when these questions stop. To say that all other facts reduce to facts about this structure is to say that it is facts about this structure, and only about this structure, that figure in the ultimate answers to “what is it for” questions.

⁵ Lewis was agnostic about whether some further story could be told about what it is for a space-time point to instantiate a perfectly natural monadic property (or for a pair of them to instantiate a perfectly natural relation, etc.). That is, even deeper ontological bedrock might be reached if this instantiation consists in the participation of the space-time point in an Armstrongian universal. (See Lewis 1983.) This issue will not matter to us here. He also granted—quite sensibly, since quantum mechanics teaches us otherwise—that it is a contingent matter whether the only perfectly natural *relations* are spatiotemporal relations. He had other, more regrettable, reasons for viewing his cherished thesis of Humean supervenience as contingent; more on this below.

(Such a principle is intended to capture the Humean idea that there are no necessary connections between distinct existences.) So an improved Humean reductionism about laws would be this: facts about the laws reduce to facts about the distribution of perfectly natural properties and relations. That's much better. First, notice that, since the reductionist is no longer claiming merely that facts about laws reduce to the totality of nonmodal facts—but rather, reduce to some specified subset of those facts—his position avoids the charge of triviality. Second, because a substantive characterization of “nonmodal” is in play (in terms of the recombination thesis), it is perfectly clear why the essentialist's position contradicts our reductionist's position. (For example, our reductionist's metaphysics says that there is a possible world containing just two massive particles, where these particles are accelerating away from each other; according to the essentialist this is strictly metaphysically impossible.)

Still, there is need for one more amendment. For Lewis (and many others) view the fundamental ontological structure of the world through the lens of properties and relations. This is a mistake. Consider mass. Are we to suppose that having mass 5 g and having mass 7 g are distinct, perfectly natural properties? If so, then there is serious trouble with the recombination thesis, for it seems a logical truth that if an object has mass 5 g, then it does not have mass 7 g. I think that what is wanted is a formulation of Lewis's position not in terms of a theory of properties and relations, but in terms of a theory of *magnitudes* (which could be monadic, dyadic, etc.). One could then build modal constraints into the structure of the magnitudes—so that any given particular can have at most one value of a given magnitude—while respecting a more limited inter-magnitude principle of recombination, so that the value a particular has for one magnitude places no metaphysical constraints on its value for any distinct magnitude, or on any other particular's value for any magnitude. Such a limited combinatorial principle would more than suffice to rule out the essentialist position. In sum, the Humean reductionist position ought to be the following: The fundamental ontological structure of the world is given by the distribution of perfectly natural magnitudes in it, where these magnitudes respect an inter-magnitude principle of recombination. All other facts, including facts about the laws, reduce to these facts. Applied to our Newtonian particle world, the thesis is that the facts about its laws reduce to the facts, over all time, about the positions, masses, and charges of all the particles.

There is one more tricky issue to wrestle with, which is the modal status of this reductionist claim, and its relation to the more typically cited supervenience thesis. Now, given the foregoing doctrines about what “fundamental ontological structure” amounts to, once you have specified the fundamental ontological structure of a

world, you get for free that all other facts about the world reduce to facts about it. (Otherwise it wouldn't count as the fundamental ontological structure.) So the crucial issue is really whether the Humean reductionist should hold that it is a merely contingent feature of our world that its fundamental ontological structure is given by the pattern of instantiation of a set of perfectly natural magnitudes that respect a limited principle of recombination. Lewis himself seemed to waffle on this point, for some bad reasons (having to do with a misguided response to Armstrong's spinning sphere example; we will pass over this) that have led to a good deal of mischief in the literature (having to do with a certain challenge posed to Humean reductionist accounts of objective chance; we will touch on this later, in §5.5). Here, I am simply going to assert without much argument that the best position for the Humean to take is an uncompromising one, according to which it is metaphysically necessary that the fundamental ontological structure of the world meet the foregoing characterization. At the very least, the following concession would, it seems to me, be disastrous: "In our world, the fundamental ontological structure is given by the positions, masses, and charges of all particles at all times. So—again, in our world—it is these facts that somehow make it the case that the fundamental laws of nature are the Newtonian laws. But it could have been otherwise. That is, there is a possible world that exactly matches ours with respect to particle motions, masses, and charges, but in which it is a further, irreducible fact that the laws are whatever they are. (In fact, there are many such worlds. In some of them, the irreducible laws are Newtonian laws; in others they are something else altogether.)" It seems to me that the only possible motivation one could have for adopting that position is a completely misconceived desire to partially accommodate the anti-reductionist intuitions of one's opponent. For example, she might insist, as evidence for her anti-reductionism, that it is intuitively obvious (1) that there are possible worlds containing a single particle moving with constant velocity, and in which the laws are the Newtonian laws, and (2) that there are still other possible worlds containing a single particle moving with constant velocity, and in which the sole law is that all particles move with constant velocity. (That example comes up a lot, in conversation.) You might think you could mollify her by adopting the position just quoted. But of course, you won't succeed, for she will turn right around and ask what epistemic right you have to assume that ours is not a world equipped with metaphysically primitive laws. All in all, a bad move: it seems to me that the vastly more sensible policy for the Humean is to forgo all attempts at conciliation. At any rate, this war of intuitions involves some further subtleties; we will come back to it in §5.1.

So now we have the core thesis of Humean reductionism: it is metaphysically necessary that facts about laws reduce to facts about the pattern of instantiation of perfectly natural magnitudes, magnitudes that respect a limited principle of recombination. There are two more theses to add to get a fuller characterization of the key elements of Humean reductionism. First is a commitment to the objectivity of laws. It is notoriously difficult to say, in general, what “objectivity” amounts to, but here I think we can make do with the following thought: According to the reductionist positions we will be considering, what it is for the fundamental laws to be such and such is not to be explained by appeal to psychological facts about humans. For example, what it is for the laws of our particle world to be the Newtonian laws is not for them to be treated in a certain way by practicing scientists. It is, rather, for it to be the case that the particle motions/masses/charges meet some (yet to be specified) condition, statable in nonpsychological terms. (By the end of the paper, we will see that the most interesting and defensible version of Humean reductionism allows for some erosion in this commitment to objectivity; see in particular §5.6.) Finally, we assume that the way that facts about laws reduce to the fundamental facts is nontrivial. More specifically, a successful reductionist account must make it the case that, for at least many possible worlds, the set of possible worlds that count as nomologically possible relative to them is neither so expansive as to include all possible worlds, nor so narrow as to include only the given world itself. We will, in particular, take it for granted that this constraint applies to our Newtonian particle world.

So much for the basic framework. Let’s now work our way up to the most interesting reductionist accounts with a brief discussion of naïve regularity theories as a warm-up exercise.

§3 The Naïve Regularity Theory

The naïve regularity theory has been refuted so frequently and so thoroughly that we don’t need to spend much time on it. Still, I think we can learn some useful lessons by taking a bit of care to see why exactly it goes wrong. So let’s apply the account to our Newtonian particle world.

The usual guiding idea is that a law is, or corresponds to, a true law-like sentence, where what makes a sentence law-like is that it is universal in form, and contains only appropriately qualitative non-logical vocabulary. Here I will add what seems clearly to be a friendly amendment, in the form of an additional constraint on the non-logical vocabulary. So let’s consider a language (a first-order language, I suppose) whose non-logical vocabulary consists solely of expressions for describing positions, masses, and charges of particles at times. That restriction guarantees that the non-

logical vocabulary is appropriately qualitative, and also avoids the immediate embarrassment of letting true claims about tables, chairs, etc. count as fundamental laws of nature. Then, where Φ is a sentence of this language, we will say that it is a law that Φ iff Φ is true, and the prenex form of Φ begins with a universal quantifier (so that Φ expresses a regularity).

Now for the problems. First, a minor one. Observe that if a sentence Φ is of the needed kind, then it is true not only in our world, but also in a world that contains novel physical magnitudes, not instantiated in our world. That is, our language seems to lack the resources to say that mass and charge are the *only* fundamental physical magnitudes there are. No matter; let us suppose that there is a straightforward way to augment the language so that it can express this claim.

Now for the more serious problems. Let us try to figure out which possible worlds, individuated up to purely qualitative identity, are compatible with the laws of our world, as this theory would have them. First, on the reasonable assumption that a universally quantified sentence is automatically true in a world in which nothing exists⁶, it follows that that world will count as nomologically possible. We will come back to the significance of that observation in a moment. Next, it is quite easy to see that among the remaining nonempty worlds, a world w will count as nomologically possible iff it contains exactly the same number of particles as our world, and there is a one-one mapping between the particles of our world and the particles of w that preserves masses, charges, and positions at all times. This result is probably obvious, but just in case, here is how we can proceed:

Observe that there is a sentence of our language, beginning with a universal quantifier, that says that for anything that exists, it is one of exactly N particles. The true one of these provides the “law” that guarantees that every nonempty nomologically possible world has the same number of particles as the actual world. Next, and letting N now be the actual number of particles, our language has the resources to define, for any time t , a predicate F that is satisfied by a particle at a time iff it is one of a set of exactly N particles that has exactly the configuration of masses, charges, and positions exhibited by all the particles at time t .⁷ Letting F and

⁶ Reasonable, even though it goes against the assumption usual in logic textbooks, that a first-order language must have a nonempty domain. I have no problem disagreeing with the textbooks here, since it seems to me silly to suppose that the claim that something exists is a logical truth.

⁷ I have tacitly assumed here that our language has names for every real number (and so, unlike the first-order languages we are used to, has an uncountable non-logical vocabulary). I do not consider that assumption problematic. But in case you’re worried, it is not necessary: if we give our language only the ability to name every rational number, we will still get the result in the text, that the only nonempty nomologically possible world (up to qualitative identity) is the actual world. I leave it as a fairly easy exercise to demonstrate this result.

G be two predicates of this type, the set of true sentences of the form “for all x , if x is F at t , then x is G at t ” will pin down the trajectories of our N particles exactly. So, if we ignore the metaphysical view that holds that two possible worlds can differ with respect to what is true at them without differing in any qualitative respect, the upshot is that the only two possible worlds compatible with the laws of our world are the empty world, and our world itself.

That is embarrassing for two distinct reasons. The first, obviously, is that the set of nomologically possible worlds is much too small—it should have uncountably many more members than it does. But the second, interestingly, is that it is also in one respect too big: for we might want it to *fail* to include the empty world. Suppose, for example, that we thought that Newton was right in holding that absolute space exists. Presumably, it was part of Newton’s intent in postulating absolute space that its existence is a non-accidental feature of our world. (It is not as if Newton thought that, as it happens, there is absolute space; but that it could, compatible with the fundamental physical laws, easily have been otherwise.) If so, then every nomologically possible world ought to include as one part absolute space. Now, as it stands, we didn’t give our language the expressive resources to talk about absolute space. But that was a minor oversight, and the point is that rectifying it would not have changed the verdict that the empty world—which lacks, among other things, absolute space—is one compatible with the laws of our world, and so nomologically possible.

So there are really two lessons to learn here from the failure of the naïve regularity theory. The first is that it falls apart not for the reasons usually cited—its inability to deal appropriately with vacuous laws, etc.—but simply because it fails to draw a nontrivial distinction between what is nomologically possible and what is nomologically impossible. The second is that it puts an illegitimate emphasis on regularities. As the example of absolute space shows, there is no particular reason to think that existential claims are intrinsically ill-suited to characterizing what is in common to the set of nomologically possible worlds. That observation will lead naturally to a friendly amendment to the vastly more interesting and sophisticated reductionist account we will take up next.

§4 The Best System Account: An Overview

Lewis, citing inspiration from Mill and Ramsey, famously put forward what he took to be an improved regularity account. For the reasons noted at the end of the last section, I think the label “regularity account” is inapt. Still, he was unquestionably correct that the account vastly improves on its predecessors. There

are, I think, two distinct guiding ideas behind this account: the one that typically gets official mention, and a second one working behind the scenes.

The official idea is that the laws are, or correspond to, true statements that collectively encode, in a highly efficient manner, a large amount of information about the world. The unofficial position is a kind of “ideal observer” view, according to which the fundamental laws are whatever a suitably placed observer, implementing the best scientific standards for judging what laws are, would take them to be. I will consider these two ideas in turn, and then highlight an obvious way in which we might think to reconcile them. (By the end of the paper, we will see the prospects for a reconciliation dim substantially.)

§4.1. *The official guiding idea*

First the official idea. Here is a classic statement from Lewis:

Certainly not just any regularity is a law of nature. Some are accidental. So an adequate regularity analysis must be selective. Also, an adequate analysis must be collective. It must treat regularities not one at a time, but rather as candidates to enter into integrated systems. For a given regularity might hold either as a law or accidentally, depending on whether other regularities obtain that can fit together with it in a suitable system. (Thus I reject the idea that lawhood consists of ‘lawlikeness’ plus truth.) Following Mill and Ramsey, I take a suitable system to be one that has the virtues we aspire to in our own theory-building, and that has them to the greatest extent possible given the way the world is. It must be entirely true; it must be closed under strict implication; it must be as simple in axiomatisation as it can be without sacrificing too much information content; and it must have as much information content as it can have without sacrificing too much simplicity. A law is any regularity that earns inclusion in the ideal system. (Or, in case of ties, in every ideal system.) The ideal system need not consist entirely of regularities; particular facts may gain entry if they contribute enough to collective simplicity and strength. (For instance, certain particular facts about the Big Bang might be strong candidates.) But only the regularities of the system are to count as laws.

We face an obvious problem. Different ways to express the same content, using different vocabulary, will differ in simplicity. The problem can be put in two ways, depending on whether we take our systems as consisting of propositions (classes of worlds) or as consisting of interpreted sentences. In the first case, the problem is that a single system has different degrees of simplicity relative to different linguistic formulations. In the second case, the problem is that equivalent systems, strictly implying the very same regularities, may differ in their simplicity. In fact, the content of any system whatever may be formulated very simply indeed. Given system S , let F be a predicate that applies to all and only things at worlds where S holds. Take F as primitive, and axiomatise S (or an equivalent thereof) by the single axiom $\forall xFx$. If utter simplicity is so easily attained, the ideal theory may as well be as strong as possible. Simplicity and strength needn’t be traded off. Then the ideal theory will include (its simple axiom will strictly imply) all truths, and *a fortiori* all regularities. Then, after all, every regularity will be a law. That must be wrong.

The remedy, of course, is not to tolerate such a perverse choice of primitive vocabulary. We should ask how candidate systems compare in simplicity when each is formulated in the simplest eligible way; or, if we count different formulations as different systems, we should

dismiss the ineligible ones from candidacy. An appropriate standard of eligibility is not far to seek: let the primitive vocabulary that appears in the axioms refer only to perfectly natural properties. (Lewis 1983, pp. 41-2)

Lewis's formulation here of what we will henceforth call the Best System Account (BSA) needs some amendments. But first, I want to make sure that we don't lose sight of the overarching idea. Lewis takes it that there is some *canonical scheme* for representing facts about the world. Then any correct representation that makes use of this scheme will have two features: First, it will have a degree of informativeness, determined purely by which possible worlds the representation rules out. So it automatically follows that if one correct representation rules out more possible worlds than a second (i.e., every world in which the second true is one in which the first is true, but not vice versa), then the first is more informative. There are thus maximally informative representations, made so by being true only in the actual world. Second, it will have a degree of simplicity, determined by broadly syntactic features of the representation ("broadly", because we want to leave it open whether the canonical scheme of representation is linguistic; more on this shortly). These two factors of simplicity and informativeness then determine an ordering—presumably, partial—among all the correct representations there are, in terms of how well each one balances simplicity and informativeness. Lewis's hope is that the nature of our world will yield a clear winner:

The worst problem about the best-system analysis is that when we ask where the standards of simplicity and strength and balance come from, the answer may seem to be that they come from us. Now, some ratbag idealist might say that if we don't like the misfortunes that the laws of nature visit upon us, we can change the laws—in fact, we can make them always have been different—just by changing the way we think! (Talk about the power of positive thinking.) It would be very bad if my analysis endorsed such lunacy. I used to think rigidification came to the rescue: in talking about what the laws would be if we changed our thinking, we use not our hypothetical new standards of simplicity and strength and balance, but rather our actual and present standards. But now I think that is a cosmetic remedy only. It doesn't make the problem go away, it only makes it harder to state.

The real answer lies elsewhere: if nature is kind to us, the problem needn't arise. I suppose our standards of simplicity and strength and balance are only partly a matter of psychology. It's not because of how we happen to think that a linear function is simpler than a quartic or a step function; it's not because of how we happen to think that a shorter alternation of prenex quantifiers is simpler than a longer one; and so on. Maybe some of the exchange rates between aspects of simplicity, etc., are a psychological matter, but not just anything goes. If nature is kind, the best system will be *robustly* best—so far ahead of its rivals that it will come out first under any standards of simplicity and strength and balance. We have no guarantee that nature is kind in this way, but no evidence that it isn't. It's a reasonable hope. Perhaps we presuppose it in our thinking about law. (Lewis 1994, p. 479)

This passage raises an issue—that the BSA might fail to guarantee that facts about the fundamental laws are objective—that we will have to come back to later. For now, just take note of the big picture: According to the BSA, laws are regularities entailed by that candidate system (read: correct representation of the world that makes use of the canonical scheme of representation) that achieves the best balance of simplicity and informativeness.

§4.2. *Some amendments*

Now for the amendments (some of them: §6 will introduce a final and important one). First, for reasons already noted, I think it was a mistake for Lewis to insist that laws must be regularities. I would therefore simplify the account by omitting this requirement, and presenting the account as, in the first instance, a theory of nomological possibility: the nomologically possible worlds are exactly those in which the best system is true. So, if the best system includes the claim that there is absolute space, then it is nomologically necessary that there is absolute space. “It is a law that A” will then be true iff A is true in every nomologically possible world—never mind if A is not explicitly included in the best system for our world.

Two comments. First, in refashioning the BSA as a theory of nomological possibility—as opposed to a theory of laws, which can subsequently be used to define nomological possibility—we do lose a bit of grain. That is, one might hold the view that even once we have specified which worlds are nomologically possible, it is a further question what the laws are that underlie these nomological possibilities. For example, one might think that there is a real question which formulation of Newtonian mechanics accurately expresses its laws—even when rival formulations generate exactly the same set of nomological possibilities. Or, one might find it embarrassing that if it is a law that P, then it is also a law that (P or Q), for any proposition Q, no matter how irrelevant.

This strikes me as a non-issue. As far as I can tell, the purposes for which we *need* a notion of fundamental law—most notably, to fix a counterfactual structure for the world, and thereby fix facts about causation, etc.—require no such finer resolution; a mere distinction between what is nomologically possible and what is not will suffice.

Second, the first quoted passage from Lewis did seem to suggest that he had a specific reason for wanting to restrict laws to regularities, presumably because he did not want the hinted at facts (existential facts?) about the Big Bang to count as laws. I confess, though, that I really have no idea what was worrying him. It is certainly true, as we will see in §5.5, that some modern defenders of the BSA (Barry Loewer and David Albert in particular) think the best system should include fairly substantive

claims about the initial conditions of our world. But they are quite happy to give these claims the status of laws, even though they don't take the form of regularities.

Next, Lewis's conception of the canonical scheme of representation is probably wrongheaded. He seems to think that it should be a first-order language, whose nonlogical vocabulary expresses natural properties and relations. But physics doesn't talk the language of properties and relations; as already noted, it talks the language of magnitudes—magnitudes most naturally represented by *mathematical* and not *linguistic* entities. An adequate scheme of representation should pay heed to this practice. Sticking to first-order logic, we could of course talk about masses and charges by means of such predicates as “— has mass —”. But then we will need to add extra, nonlogical axioms, stating that nothing can have more than one mass, that mass values must be non-negative real numbers, etc. These axioms will inevitably add to the syntactic complexity of the candidate system that employs them; but it seems quite unfair to charge a system for such added complexity. (Remember that the virtues exhibited by the best system are supposed to mirror the virtues we see in scientific theorizing. Are we really to suppose that a scientific theory automatically loses points for describing the world in terms of mathematically structured magnitudes, rather than unstructured properties and relations?)

Second, the standard way that contemporary physical theories delimit the class of nomological possibilities is by means of mathematical models, typically defined in part by differential equations. So why not let the canonical scheme of representation simply be such models, or rather the equations that define them? There would still, unquestionably, be a role for the distinction between natural physical magnitudes and gerrymandered ones. For example, we could require that the variables that appear in a set of defining equations correspond only to genuine (i.e. perfectly natural) physical magnitudes.

At any rate, I think we can leave for another day the question of how best to work out this alternative account of the canonical scheme of representation. It is quite clear that pursuing it would do nothing to threaten the spirit of Lewis's proposal (and would probably go a long way towards improving its lettering).

Let's briefly consider how the BSA (as I have partially reconceived it) would apply to our Newtonian particle world. The basic idea is that we want to capture as much information as we can about the motions, masses, and charges of particles by exhibiting equations that particle behavior must conform to. If we were allowed to invent any old magnitude we wished, allowing a parameter representing it into our equations, then this would be easy: define a “magnitude” M that is stipulated to have value 1 for a particle exactly if that particle belongs to a world of particles whose behaviors are just what they actually are. But if we require that the parameters that

appear in the equations correspond to natural magnitudes—i.e., position, mass, and charge—then presumably things aren't so easy. It is not all implausible that the equations that appear in Newtonian mechanics (in, say, its Hamiltonian formulation) achieve, among all such candidate equations, far and away the best balance of simplicity and informativeness.

§4.3. *The unofficial guiding idea*

So much for the first guiding idea. The second is, in fact, already hinted at in the first of the passages quoted from Lewis. Here it is again: "...I take a suitable system to be one that has the virtues we aspire to in our own theory-building, and that has them to the greatest extent possible given the way the world is." Elsewhere, he is even more explicit: "The standards of simplicity, of strength, and of balance between them are to be those that guide us in assessing the credibility of rival hypotheses as to what the laws are." (Lewis 1986, p. 123)

To help flesh this idea out more fully, let us imagine someone whom I will call a Limited Oracular Perfect Physicist. What makes our LOPP a perfect physicist is that, given as evidence any information about the world, she is perfectly able to judge what hypotheses about the fundamental physical laws are most strongly supported by that evidence. What makes her oracular is that she has, as evidence, quite a lot of information about the world. But not, of course, all information: else her job would be too easy. (For example, she doesn't directly receive as evidence information about what the fundamental laws are.) Specifically, we will suppose that what she has available to her as evidence are all the facts about the distribution of perfectly natural magnitudes. In the case of our Newtonian particle world, her evidence consists in perfect information about the motions, masses, and charges of every particle, together with the further information that the world contains nothing else. The second guiding idea, then, is roughly that the laws are whatever she says they are.

Some comments. First, and most obviously, the device of introducing a fictional LOPP is dispensable. What the second guiding idea really needs to assume is that there are, implicit in the practice of physics, evidential standards for determining what the fundamental physical laws are that induce a mapping from possible worlds to something like a probability distribution (or perhaps a family of such distributions) over propositions about the fundamental laws of nature. It is nothing more than a useful heuristic to imagine a creature who holds that mapping in her head. Next—and this is quite important—it is no part of anti-reductionism per se that it must reject the existence of such a mapping. That is, an anti-reductionist can—and should—perfectly agree that there is a fact of the matter about how one's credence over alternative hypotheses about the fundamental laws of nature ought to be

distributed, were one fully apprised of the positions, masses, and charges of every particle at every time. (In the case of our well-behaved Newtonian world, the anti-reductionist presumably will say that one ought to be all but certain that the laws are the Newtonian laws.) But the anti-reductionist will view this as a purely *epistemic* fact of the matter. So we can summarize the second guiding idea this way: the Humean reductionist is taking standards that *both* sides endorse—but that his anti-reductionist opponent views as solely epistemic standards—and elevating them to the status of standards that are constitutive of laws of nature.

There are two subtleties that we need to take note of. They have the same source, which is that it is not legitimate for our Humean reductionist to assume that the epistemic standards implicit in the best practice of physics have a built-in bias toward his metaphysical position about laws. In fact, what makes the most dialectical sense is for him to assume the opposite (and, for reasons that we will get to later, this assumption is independently plausible). He will then simply need to compensate—with luck, only slightly—for this anti-reductionist bias. Let me explain.

Imagine that we are reductionists, and that we are interviewing our LOPP. She has been informed of the complete history of particle motions, masses, and charges, and has formed in response an opinion about what the laws are, hence which worlds are nomologically possible. We ask her to tell us which worlds these are. She tells us that since, of course, facts about laws do not reduce to nonmodal facts, in order to tell us which worlds are nomologically possible she must tell us two things: she must tell us, for each nomologically possible world, what its history of particle behavior is, and she must also tell us what its laws are. Never mind that, according to her, the second part is easy, since it is nomologically impossible for the laws to be other than they *actually* are. We are, all the same, distressed. For she is not giving us exactly the information we want—indeed, she is giving us “information” (namely, about metaphysically irreducible laws) that we think does not exist. So we rewind, and ask our question more carefully: “Please, oh LOPP, tell us which worlds, described up to but not beyond their nonmodal features, are nomologically possible.” We need not worry that her answer to *that* question will presuppose anything objectionable. The moral—and admittedly, it is a fairly minor one—is that our slogan formulation of the second idea went slightly awry. We should have said that the Humean reductionist takes epistemic standards *for determining the nomological possibilities for the nonmodal facts*—standards that both he and his opponent endorse as such—and promotes them to the status of standards that are constitutive of the nomological possibilities for the nonmodal facts. (The reductionist then adds, of course, that these are *all* the facts.)

The second way in which the reductionist must compensate for an anti-reductionist element in the LOPP’s standards is a bit more subtle. I will postpone

discussion of it until §4.5. Let us first take advantage of the second guiding idea in providing a fairly user-friendly explanation of the Humean's treatment of objective chance.

§4.4. *The treatment of objective chance*

We need a new example. I will borrow one that I've used before, for a different purpose (see Hall 2004).

Imagine a particle world that is almost like our Newtonian world, with one significant exception: The behavior of its particles conforms exactly to Newton's laws of motion and the relevant force laws, except that whenever two particles collide, they either rebound in a perfectly elastic collision, or stick together in a perfectly inelastic collision (forming, let us say, a particle whose mass and charge are the sums of the masses and charges of the colliding particles). We would like to add to this description that there is, for each collision, some objective chance of each outcome; but the reductionist can say no such thing at the outset. For, according to him, such facts do not belong to that elite set of facts to which all other facts about a world reduce.⁸ So we must proceed with our description of this world in a more roundabout way. Let us therefore stipulate, first, that an absolutely enormous (but finite) number of collisions takes place, and second, that the frequencies of the two types of outcome (elastic vs. inelastic collision) show no dependency on any of the physical parameters that characterize the collisions: no dependency on the intrinsic physical parameters of the colliding particles, and no dependency on their relative velocities, accelerations, etc.

What do I mean, "no dependency"? We need to be careful here. Of course there will be *some* function that takes complete specifications of these parameters as input and outputs a real number that correctly describes the frequency of elastic outcomes among collisions with those values of the parameters. (Given that many combinations of values are never realized in this world, then there will be infinitely many such functions.) The point is that by the reductionist's lights, this function will be far too *complicated* to deserve to be included as part of the fundamental laws of this world.

We can get a firmer purchase on how the reductionist introduces facts about chances to this world by considering what our LOPP will conclude, when given complete information about how particles in this world behave. To begin, she will surely give the lion's share of her credence to a certain kind of probabilistic

⁸ So I am setting aside here the possibility of a Humean reductionist view that treats objective chances as among the perfectly natural magnitudes. For good reasons to set this possibility aside, see for example Ismael 1996.

hypothesis, according to which each interaction has the same single-case chance of yielding a given outcome, where these objective probabilities are independent of one another. But it is just as clear that she will not settle on a unique value for this single-case chance. Rather, she will have a credence distribution over a range of values from 0 to 1 (exclusive), sharply peaked around the value that is equal to the actual relative frequency. Presumably, what the reductionist would like to do in this case is to “round off”, and insist that the proper metaphysical conclusion to draw is that the single-case chance *just is* the relative frequency, since this is the value on which her credence distribution peaks.

We will come back, in the next section, to a lingering problem about this “rounding off” maneuver. For now, I wish to draw attention to something else. Our LOPP takes as evidence facts about particle motions, masses, and charges—but *not about chances*. She outputs a choice of a “best” (i.e., most probable) hypothesis, which *inter alia* is a hypothesis *about chances*. Now, we are assuming that she is the kind of anti-reductionist who invests in the notion of “chance” a content that no reductionist could give it (since, most obviously, perfect knowledge of the nonmodal facts leaves her with residual uncertainty about what the chances are). No matter: since we are assuming that she perfectly implements the very best scientific criteria for theory choice, we can read off of her choice for “best” theory for a given world a way to fix, in a well-motivated and reductionist-friendly fashion, facts about chances for that world. In general, the claim “ $ch_t(A) = x$ ” (read: the chance, at time t , that the proposition A will come true, is x) is true at world w iff the theory that the LOPP would choose as best if given as evidence the totality of fundamental, nonmodal facts about w , when applied to the time t state of w , yields the conclusion that $ch_t(A) = x$.⁹ More simply: “ $ch_t(A) = x$ ” is true iff a suitably informed LOPP *says* it’s true (really: most likely to be true). Less simply, but more illuminatingly: *What it is* for a claim of the form “ $ch_t(A) = x$ ” to be true is for the fundamental, nonmodal facts of the world to be such that they would lead a LOPP informed of them to conclude that this claim was (most likely to be) true.

What remains is for the BSA to supply a specific hypothesis about what the standards are upon which our LOPP relies. Lewis adds, to the usual standards of simplicity and informativeness, a standard of statistical *fit*, measured by how likely a probabilistic candidate system judges the actual history of the world to be. That may not be the best choice; see for example Elga (2004). But worries about it are, I think, comparatively superficial: for notice that they do not threaten the unofficial guiding

⁹ I’m assuming as a necessary truth that chances behave in a Markovian fashion. We can relax this assumption by letting the best theory be applied to the entire history of w up to t .

idea, but only one specific way of fleshing it out (i.e., one specific proposal about what the relevant and implicit standards for theory choice are).

Time to return to our unfinished business: what should a reductionist say, when the evidence available to our LOPP leaves her undecided as to what the laws are?

§4.5. *A problem with “rounding off”*

Observe that, even when our LOPP is given complete information about our Newtonian particle histories, it is plausible that she will reserve some share of credence, however tiny, for hypotheses about the fundamental laws according to which they are not the Newtonian laws, but something else—something that, presumably by some remarkable coincidence, happens to conspire with the initial conditions of our world to make it *look* as if it is a Newtonian world. So the reductionist will need to perform a “rounding off” maneuver—in this case, taking the nomological possibilities to be given by that hypothesis to which the LOPP accords the highest share of credence. Let’s explore this issue further, for it does, alas, contain a potentially serious problem for the reductionist, which I have only the beginnings of an idea how to solve.

It obviously won’t do for our reductionist to say that, when the LOPP fails to assign a credence of 1 to some law-hypothesis, the proper metaphysical conclusion to draw is that it is indeterminate what the laws are, with the range of indeterminacy running over those hypotheses to which she assigns credence greater than 0. At the same time it will, by the reductionist’s lights, *sometimes* be indeterminate what the laws are. So perhaps he should say that, provided the LOPP assigns a credence sufficiently above 0.5 to some hypothesis, it is a determinate fact that the nomological possibilities are specified by that hypothesis.

That might work, as far as it goes. But it doesn’t go far enough, as our discussion of the chancy Newtonian world already demonstrated. There, we saw that our LOPP would end up assigning a credence *distribution* over a continuous range of hypotheses (and so: at most infinitesimal credence, to any specific one of them). And what the reductionist would nevertheless like to claim, about that case, is that there is a determinate fact about the probabilistic laws, given by that hypothesis on which the LOPP’s credence peaks.

But there will be other, structurally similar cases in which the reductionist will want to claim that there is some *objective indeterminacy* in the laws. To see this, let us complicate our chancy Newtonian example, and imagine that the phenomena support a functional version of a probabilistic law, as follows: This time, the frequencies of the two types of outcome *do* show a simple dependency on the

physical parameters that characterize the collisions. Specifically, they conform very closely to the following equation:

$$\text{frequency of elastic outcomes} = e^{-kM},$$

where M is equal to the sum of the masses of the colliding particles and k is some constant.

So far, there is no new problem—though it will presumably still be the case that our LOPP distributes her credence over a range of possible values for k . Now for a devious twist. Suppose that there is some other way in which to parameterize the collisions—say, by the combined charges of the incoming particles C —that gives rise to frequency data that is just as simple, but that yields a different and incompatible probabilistic law—say, this one:

$$\text{frequency of elastic outcomes} = e^{-k^*C},$$

where k^* is some different constant. And by “incompatible”, I mean *really* incompatible: for many and perhaps all collisions, the two laws assign substantially different single-case chances. However much we might indulge in pious Lewisian hopes that our world is not so ill-behaved, it is, by the reductionist’s lights, unquestionably metaphysically possible for one to behave in this way.

What a reductionist should say about such a world is that it is indeterminate what its probabilistic laws are. The question—which I am going to have to leave mostly open—is how to construct a detailed theory that says that we have indeterminacy in *this* case, but not in the first case examined above (the case of our simple chancy Newtonian world). For what it’s worth, my guess is that the best strategy for the reductionist is to say that in a case like this, the implicit standards that both he and his opponent endorse would *themselves* recognize a case of indeterminacy. That is, our LOPP would respond to this data not by dividing her credence more or less equally between the two probabilistic laws, but rather by insisting that, as an epistemic matter, there is no determinate fact about what one’s credences ought to be in this case. That verdict seems to me intuitively very plausible—after all, the data are *utterly bizarre*—but the matter deserves more consideration than I can give it here.

§4.6. *Reconciling the two guiding ideas*

So we have two, apparently distinct, guiding ideas. I will close this section by highlighting the obvious way to try to knit them together. Take the second guiding idea—that the implicit standards for judging lawhood are in fact constitutive of

lawhood—to be the central, nonnegotiable idea. Take the first idea—that lawhood is a matter of balancing simplicity and informativeness (and perhaps fit, in the case of probabilistic systems)—to be a substantive proposal about the nature of these implicit standards. Put another way, the “candidate systems” that the first guiding idea talks about can be thought of as constituting the range of hypotheses that our LOPP considers. She employs some criteria for ranking them in terms of plausibility, and what the first guiding idea tries to do is to say explicitly what those criteria are.

Later on I will cast doubt on whether this reconciliation will really work. But for now it will be very useful to keep in mind when reviewing challenges to the BSA, for it will be important to distinguish those challenges that target the reductionist’s specific proposal to analyze laws in terms of simplicity and informativeness, from those that target the underlying idea that the practices implicit in physics for judging lawhood can be used to tell us what laws are.

The next section takes up a long menu of challenges to the BSA. Most of them can, I will argue, be met rather handily; but the last of them will require a substantial revision to that account, presented in §6. This revision leads in turn, I think, to a very serious dilemma for the Humean reductionist. I will close in §7 with some speculative remarks about how the Humean might best overcome this dilemma.

§5 A menu of challenges to the BSA

§5.1. *The central anti-Humean intuition*

The simplest and most direct way to argue against reductionism is by way of the following kind of thought experiment: Imagine a world with a single particle, moving with constant velocity. There at least two possibilities—each, so this argument would have it, a genuine metaphysical possibility. The first is that this particle’s motion is governed by a law that says that all particles move with constant velocity. The second is that the motion is governed by the Newtonian laws. A possible world of the first type perfectly duplicates a possible world of the second type with respect to its nonmodal facts, even though the laws of these worlds are different. So facts about laws do not reduce to nonmodal facts.

For some philosophers, the game is over at this point. For, they will claim, there is an a priori connection between what is conceivable and what is metaphysically possible, and thanks to this connection the thought experiment *establishes* that there are two possible worlds differing in their laws, but not in their nonmodal facts. But there is no such a priori connection—at least, not one nearly so direct enough to do the work needed of it here. My reasons for saying this are not the usual Kripkean ones (which, arguably, would not apply in a case like this anyway). Rather, I’m

impressed by the existence of many debates in fundamental ontology, each of which would lead to paradox, if this modal rationalist thesis of an a priori connection between the conceivable and the possible were correct. For example, a presentist holds that the only things that exist (tenselessly speaking) are things that exist now; an eternalist denies this claim. All parties agree that they are disputing a necessary truth: it's not that there are some genuinely metaphysically possible worlds in which presentism is true, and others in which eternalism is true. At the same time, each position is conceivable. (Never mind that a typical eternalist will all too frequently deny that presentism is conceivable, and vice versa: this is just the kind of rhetorical posturing that one comes to expect from philosophers, when knock-down arguments aren't available.) We could multiply examples: think of debates about the existence of God, or about the nature of persistence through time, or about unrestricted composition, etc. The proper verdict is that there is no quick route from judgments of conceivability to conclusions about the nature of fundamental ontology.¹⁰

All the same, while it would be appropriate for the reductionist to deny that the foregoing thought experiment refutes his position, it would not be appropriate for him to wholly ignore the intuitions that drive it. Now, I think that there is a perfectly good response that a reductionist can make to this kind of thought experiment. But first, three bad responses:

To begin, a reductionist might grant that these are genuine possible worlds, but insist that in our world, facts about fundamental laws *do* reduce to the nonmodal facts about how perfectly natural magnitudes are distributed. We have already considered and rejected this option (in §2).

Second, some reductionists (e.g. Loewer 1996 and Beebe 2000) claim that much or all of the appeal of such thought experiments derives from the lingering influence of a wholly misconceived view of laws as the dictates of a divine lawmaker. This is a curious claim. Never mind that hard-headed anti-reductionists introspect no such influence. What is unclear is how the diagnosis could possibly be relevant, even if such influence were present. To see this, it is worth noticing that there is nothing in a divine command account of laws that would prevent a Humean reductionist from endorsing it. Here, for example, is how a BSA-friendly divine command account

¹⁰ Chalmers—probably the most visible contemporary defender of modal rationalism—would, I take it, insist that the conceivability-possibility connection holds, provided that the judgments of conceivability are suitably *ideal*; see for example his 2002. I doubt that he's right about this, but at any rate it hardly matters: there is enough controversy about the status of the anti-reductionist conceivability judgments at work in the thought experiment that no one could plausibly claim that they are "ideal", in the relevant sense.

would go: Among the fundamental, nonmodal facts about the world are the fact that God exists, and has various intentions regarding the behavior of the physical parts of the world. Among the rest of the fundamental, nonmodal facts are facts that make it the case that the physical world happens to conform to these intentions. So, plausibly, a maximally simple and informative description of the world will include a claim that the natural portions of it conform in their behavior to God's intentions. So this fact will, according to the BSA, be nomologically necessary. Its nomological necessity will underwrite such counterfactuals as this: If God had had such-and-such different intentions, the natural world would have behaved in a correspondingly different way. And it is in virtue of the truth of such counterfactuals that it is true to say that things in the natural world behave the way they do *because* God so intends it. In short, you could be a thoroughgoing Humean reductionist and still think that God pushes things in the natural world around.

Presumably, a typical divine command theorist does not have *this* conception of her position, but rather has a conception according to which it *conflicts* with Humean reductionism. But that is because she has, as part of her *antecedently available conceptual repertoire*, the resources to characterize her divine command theory in the intended (i.e., anti-reductionist) manner. Specifically, she has a conception of *power* according to which whether something (in this case, God) has it is not a fact that reduces to the totality of nonmodal facts. Equipped with such a conception, she can formulate her preferred divine command theory. And she could just as well have formulated any of a number of distinct, and decidedly secular, anti-reductionist theories (e.g., by attributing powers to objects).

The third bad reason to reject the anti-reductionist modal intuition purports to rest on a demonstration that there is no coherent philosophical account of laws compatible with that intuition. Here I have in mind Lewis, who quite rightly argues that Armstrong's account of laws as higher-order relations between universals is inadequate, but then concludes, on the basis of an equally inadequate survey of the alternatives, that reductionism is the only viable position that remains. Here is an example of where the crucial bit of overlooking happens:

To begin, we may be certain a priori that any contingent truth whatever is made true, somehow, by the pattern of instantiation of fundamental properties and relations by particular things. In Bigelow's phrase, truth is supervenient on being... If two possible worlds are discernible in any way at all, it must be because they differ in what things are in them, or in how those things are. And "how things are" is fully given by the fundamental, perfectly natural, properties and relations that those things instantiate. (1994, pp. 473-4)

Lewis takes it that Armstrong's account of laws respects the allegedly a priori constraint he lays down here¹¹, because Armstrong's posited higher-order universal of "necessitation" counts as a perfectly natural relation. It's just that this account is deeply objectionable on other grounds (notably, Armstrong never makes it adequately clear why the "necessitation" universal deserves that name).

But now consider an anti-reductionist view like Maudlin's (2003), which holds that laws are metaphysically primitive features of the world, that determine how complete physical states generate successive physical states. Are we to suppose that this primitivism somehow violates the a priori principle that truth supervenes on being? Lewis seemed to think so.¹² That verdict does not withstand scrutiny. Notice, in the passage just quoted, how Lewis slides from the perfectly benign characterization of the key constraint in the penultimate sentence, to the tendentious and indeed question begging formulation in the last sentence. The benign characterization is reminiscent of and perhaps equivalent to the thesis I endorsed in §2, that there is such a thing as the fundamental ontological structure of the world. But it may well be that, for Maudlin, characterizing the complete fundamental ontological structure of the world requires us to break out of preconceived metaphysical categories—in particular, facts about laws do not consist in facts about the instantiation by particulars of properties and relations. That just shows that the claim that "truth supervenes on being" loses its status as uncontroversially a priori once glossed in the way Lewis prefers. If, on the other hand, this claim simply amounts to the thesis that there is some fundamental ontological structure to the world, and all facts about the world supervene on facts about this structure, then by all means we should uphold it as a nonnegotiable a priori truth. But, *pace* Lewis, we should not expect this claim to be of any use in settling the debate between reductionists and anti-reductionists.

What, then, is the best reductionist response to the thought experiments? Simply to claim that constructing a theory that respects them comes at too great a cost in clarity and coherence. Yes, the reductionist position comes at a clear intuitive cost. Yes, other things equal, it would be better not to have to pay that cost. But other things are not equal. Why not? That is a vexed question. It is not, as Maudlin rightly emphasizes (2006), that the reductionist position can be properly motivated by appeal to some empiricist condition on concept formation, since we no longer think

¹¹ I can't resist drawing attention to Lewis's deft and ingenious use of rhetoric here: he doesn't state that we *are* certain a priori of his key claim (such a bold pronouncement would have set off alarm bells in the attentive reader), but rather that we *may be* certain a priori of this claim. That makes things go down more smoothly. But then he turns right around and treats this claim as nonnegotiable!

¹² This impression was confirmed in conversation with him.

we have in hand a plausible candidate for such a condition. Nor can it be properly motivated simply by complaining that facts about the anti-reductionist's laws are underdetermined by the empirical data, since general anti-realism about the unobservable is—the mischief wrought by van Fraassen and others notwithstanding—a ridiculous epistemological position.

But we have to be careful here, for these observations hardly show that the world is safe for anti-reductionism. Agreed: there are certain bad empiricist theories of concept formation (the impressions-based theory of Hume; various verificationist theories) according to which the anti-realist position is conceptually incoherent. Agreed: there is a certain bad epistemological theory (generalized skepticism about the unobservable) according to which the anti-realist position yields the unacceptable conclusion that we can't possibly know what the laws are. But the existence of bad philosophical theories proves nothing. In particular, that these theories are so bad does very little (*pave* Maudlin) to show that there are no epistemic or conceptual grounds for doubt about anti-reductionism. And it is clearly such doubts, however loosely formulated, that form the primary motivation for Humean reductionism. Here is a nicely representative expression of these doubts, from Loewer:

Carroll and Maudlin drop the metaphors of directing and guiding and simply maintain that laws fail to supervene on the Humean facts. So far as I can see, there is no incoherence in their position. There are possible worlds in which some regularities instantiate a non-Humean property X and in which these regularities satisfy all the conditions on laws with the exception of [the requirement that laws govern regularities]. However, there are still metaphysical puzzles about A-laws ["A" for Armstrong, a leading anti-reductionist]. It is the fact that a generalization instantiates property X that is supposed to empower it to explain its instances, support counterfactuals, etc.; i.e., it is that fact which makes it a law. The metaphors of directing and guiding or Armstrong's invocation of necessitation are supposed to provide some sort of an account of how A-laws explain their instances, support counterfactuals, etc. But once these metaphors are rejected [as, Loewer argues elsewhere, they should be] it is unclear why or how the satisfaction of X enables a generalization to perform these feats. Carroll and Maudlin simply accept that it is a basic fact that A-laws explain, etc., without providing any account of what it is about them that enables them to do so. Their attitude is hardly different from Armstrong's recommendation of natural piety. Our reasons for believing that there are A-laws have to be very strong to justify such devotion. (Loewer 1996, p. 119)

So here, I think, is where we stand. The reductionist should recognize that much of our ordinary conception of law of nature, and indeed, much of our scientifically informed conception, has a distinctly anti-reductionist cast to it. So he should be forthright that he is advocating an at least modestly revisionist account of laws of nature. He needs to make clear the benefits of going revisionist, and make clear that the costs are not exorbitant. But by the same token, there is unfinished business on

his opponent's part in laying out a transparently coherent anti-reductionist theory, one that shows—without relying on unhelpful metaphors of “directing”, “guiding”, or “governing”—how laws, as she conceives them, can do the various jobs required of them.

The honest verdict, then, is that as far as this part of the debate is concerned, we have an unsatisfactory stalemate. Let's move on.

§5.2. *Essentialism about physical magnitudes*

A second, and quite different anti-reductionist intuition defends the essentialist position mentioned earlier by targeting the reductionist commitment to a limited principle of recombination. A reductionist will, for example, recognize as possible a world containing just two particles, oppositely charged and both possessing mass, but moving in such a way that they are accelerating *away* from each other. Granted that this world is not nomologically possible; it is still metaphysically possible. Or so the reductionist will and *must* claim, given that he holds that mass, charge, and position are perfectly natural magnitudes such that the way any one of them is instantiated places no metaphysical constraints on the ways the others are instantiated. The anti-reductionist we are now considering denies that there is such a possible world—or, more carefully, denies that it has been correctly described as one in which *mass* and *charge* are instantiated. (She can perfectly grant that there are possible worlds in which two particles instantiating other, “alien” physical magnitudes accelerate away from each other in the given matter.) She might hold, with Shoemaker (1980), that nothing could count as mass or charge unless it were governed by exactly the same fundamental laws as actually govern it; or she might hold a more moderate position (which may be Cartwright's: see her 1999), according to which something counts as having mass only if it has broadly the same causal powers that actual massive objects do (where “broadly” is intended to be more permissive than “exactly”). Either way, the reductionist who thinks of mass and charge as among the fundamental perfectly natural magnitudes cannot agree with her.

The essentialist position is sufficiently controversial, and the main arguments for it sufficiently weak, that I think it is quite acceptable for the reductionist to make no attempt at accommodating the essentialist's intuition. But if he really wants to, there is a strategy he can pursue that will go some way towards mollifying this particular opponent. Since the strategy is independently interesting—and since there is also, as we will see, a much better motivation for it—we'll take a look.

The second, unofficial guiding idea behind Humean reductionism about laws provides the key. To see how, consider again the reductionist treatment of objective

chances. According to the reductionist, facts about the objective chances are not among the fundamental nonmodal facts of any world. All the same, there certainly *are* facts about objective chances. They are, very roughly, whatever a suitably informed LOPP would say they are. Take the case of our chancy Newtonian world. Here, our LOPP will have as evidence all the facts about how particles move, and what their masses and charges are. She will then consider hypotheses that make explicit use of the notion of chance. One of these hypotheses will, by her lights, be best. What makes it the case that the chances are what they are is simply *that* this hypothesis about them is best. We thus get, according to the reductionist, somewhat roundabout truth conditions for claims about objective chance: for example, the claim that at time t , the chance that particles A and B will collide elastically is x is true just in case the suitably informed LOPP would say it is true (never mind that, being an anti-reductionist, she means something different by this claim).

A reductionist could perform exactly the same maneuver with respect to facts about mass and charge. That is, he could say that mass and charge are not fundamental, perfectly natural magnitudes, but rather are introduced in a manner exactly analogous to the way chance-facts are introduced. Consider our Newtonian particle world. The reductionist we are now considering holds that the *only* fundamental nonmodal facts there are about this world are facts about the *locations* of particles at all times.¹³ But a candidate system is now allowed to hypothesize, as it were, that particles are also characterized by additional magnitudes, and introduce equations connecting the values of these magnitudes to particle positions. These systems will thus start out as partially uninterpreted in a way exactly analogous to the way in which probabilistic candidate systems start out as partially uninterpreted (remaining so unless and until they are declared “best”). What would make it the case that there *are* masses and charges is just that there is a candidate system that says so, and that, partly *by* saying so, manages to achieve an optimal combination of simplicity and informativeness (informativeness, remember, only with respect to particle positions). So in our particle world, all they really are are facts about positions of particles at all times; but if we pretend that in addition, each particle is characterized by an unchanging value of two magnitudes (one with real values, the other with nonnegative real values), then we can write down very simple equations that encapsulate quite a lot of information about the particle motions. The final step is to let the facts about particle motions that make it the case that these equations achieve such an optimal balance of simplicity and informativeness constitute the

¹³ Actually, he need not be quite so dogmatic: he could remain agnostic about the existence of other perfectly natural magnitudes, insisting only that their distribution makes absolutely no difference to what the fundamental physical laws are.

truth-makers for claims about particle mass and charge, so that those claims can now be understood as literally correct. What results, again, is a philosophical position about mass and charge that is exactly analogous to the position the BSA already takes about objective chance. Mass, charge, and chance are all, in a certain specific sense, *manufactured* magnitudes.¹⁴

As such, they violate certain of our intuitive preconceptions. For example, just as we tend to think of the facts about the objective chance of decay of an unstable atom as being facts wholly intrinsic to the atom, so too we tend to think that the facts about the mass and charge of a particle are facts wholly intrinsic to that particle. In both cases, this intuition of intrinsicness has to go, if the reductionist is right. Still, this maneuver results in a position that unquestionably respects the letter (if perhaps not the spirit) of the essentialist's intuition. For consider a world containing just two particles, accelerating away from each other. Could it be that these particles both have a value for mass, and have opposite charges? No. For on the present account, *what it is* for a particle to have a value for mass or charge is, *inter alia*, for it to belong to a world whose best system is the very same as the best system for our Newtonian particle world. And that system is not even *true* of the two-particle world we are considering. More generally, while we can still expect a principle of recombination to govern configurational magnitudes such as position—so that the position of one particle at one time places no metaphysical constraints whatsoever on the position of any other particle at any other time—only a highly constrained version of such a principle governs mass and charge. Specifically, if we limit our attention to the nomologically possible worlds, then we can say that for any given time, it is possible

¹⁴ Barry Loewer has suggested, both in conversation and in an unpublished paper, that there might be a way to develop a best-system account that *wholly* dispenses with what he sees as an unfortunate commitment on Lewis's part to a doctrine of natural properties and relations. The hope is that candidate systems can somehow specify *both* the laws *and* the natural properties as a "package deal", so that which candidate system is best fixes which properties are natural. The idea in this section for a modified reductionist position first took shape as an attempt on my part to provide a partial fulfillment of Loewer's aims: for candidate systems can now be seen as specifying what at least *some* of the basic physical magnitudes are. But Loewer's more ambitious goal seems to me unreachable, and that is because I do not see how to accomplish even the more limited goal without firmly adhering to a distinction between basic physical magnitudes that are *perfectly natural* (even if these turn out to be *only* configurational magnitudes) and merely "manufactured" basic physical magnitudes. Put in terms of our LOPP, the worry is this: She needs to be given *some* information about the world—information that is specifiable *independently of what the best system is*—in order to have a basis on which to judge the various law-hypotheses. Granted: she need not be given any information about chances, nor even about masses and charges. (So facts about these quantities can be seen to obtain in virtue of which system is best.) But if, in addition, we deprive her of information about *positions*, what does she have left to go on?

for the particles existing at that time to have any distribution of masses and charges.¹⁵ But that's it.

I mentioned above that there was another motivation for the “manufactured quantities” approach to mass and charge than the generous but probably misguided desire to appease the essentialist. It derives from the thesis that the primary aim of physics—its first order of business, as it were—is to account for *motions*, or more generally for change of spatial configurations of things over time. Put another way, there is one Fundamental Why-Question for physics: Why are things located where they are, when they are? In trying to answer this question, physics can of course introduce *new* physical magnitudes—and when it does, new why-questions will inevitably come with them. (So it is no part of the thesis we are considering that physics is *only* concerned with explaining motions; it is just that the other explanatory demands on it are, in a certain sense, derivative on this one.) Now, what exactly is physics doing, when it introduces these new magnitudes?

One answer is that it is simply postulating new magnitudes, end of story. Perhaps these are fundamental, perfectly natural magnitudes; perhaps they are magnitudes that further investigation will show reduce to more fundamental ones. Regardless, it is part of this realist view that the success of a physical theory that postulates some magnitude at explaining motions—even perfect success, down to the most microscopic level—provides no guarantee that the theory is true in what it says about the given magnitude. For example, a Newtonian theory might perfectly capture the motions of particles in our world, but for all that be false: for it might be

¹⁵ Matters are in fact a bit more complicated, for two reasons. First, in some of these nomologically possible worlds, the best system will not be the best system of our Newtonian world; so nothing in these worlds will count as having mass or charge. Accordingly, it would be more accurate to state the recombination thesis this way: First, define an “initial condition” as a specification of a finite number of particles, together with a distribution of masses, charges, positions, and velocities among them. Then for every mathematically possible initial condition, there is a nomologically possible world whose particle motions are given exactly by the equations of the best system for our world, as applied to this initial condition. In other words, the various possible ways of assigning values to the uninterpreted magnitudes of mass and charge, together with values to the (interpreted, obviously) values of position and velocity, produce, given the equations of our best system, a class of possible motions; these are exactly the nomologically possible motions. It does not matter that some of these motions will be exhibited in worlds in which there is no such thing as mass and charge. The second reason matters are more complicated than suggested in the text is that the text implicitly supposed that the best system for our world supplies an ICH that is maximally permissive. It is far from obvious that it should. How do we know that much greater informativeness cannot be achieved, at minor cost in complexity, by narrowing down the range of possible initial conditions in some manner? We don't, as it turns out. But since this is precisely the issue over which §6 obsesses, we can afford to ignore it for now.

that the particles of our world merely behave *as if* they have masses and charges, when in fact they uniformly lack these magnitudes.

A second, very different answer goes instrumentalist: “mass” and “charge” are merely introduced as calculational devices, in a way that utterly deprives claims about their distribution of truth-values. Rival claims about the mass or charge of some particle might lead to better or worse predictions about the motions of it and other particles; but while these *predictions* can certainly be evaluated for truth, the rival claims that partly underlie them cannot.

The version of reductionism sketched in this section can be seen as trying to chart a course between these two answers. With realism, it holds that claims about masses and charges have truth-values. But in an instrumentalist spirit, it holds that the truth conditions for these claims concern, in effect, the predictive success of theories that employ them. In this way, it seeks to avoid the highly implausible view—a view that instrumentalism seems inevitably committed to—that a significant fragment of scientific language does not have the function of describing what the world is like, while at the same time denying the possibility of a certain kind of underdetermination—a kind that can easily seem absurd, given the central explanatory aims of physics. (We will come back to this underdetermination worry in §5.7, below.)

Verdict: There is no serious threat from the essentialist position, but a natural way for the reductionist to accommodate it turns out to have independent interest.

§5.3. *Counterfactuals, explanation, induction*

The next set of challenges to reductionism can, I think, be dispensed with rather quickly. To set the stage, let us put our anti-reductionist glasses on, and consider how they will color our views of counterfactuals, explanation, and induction. The crucial idea is this: From the standpoint of an anti-reductionist, there is all the difference in the world between a possible world in which there are genuine laws governing the phenomena and a possible world that lacks such laws. For an anti-reductionist, the latter sort of world quite literally exhibits no constraints on how things in it can behave. For example, suppose there is some other particle world, with particle histories that exactly match those of our world. But, whereas our world is governed, we will stipulate, by Newtonian laws, this other world is governed by no laws whatsoever. For the anti-reductionist, that means that all of the following claims are true of this “Hume world”:

- Counterfactuals that have perfectly determinate truth values in our world lack them in the Hume world—e.g., counterfactuals that say what would have happened if such-and-such a particle had been moving in such-and-such a way.

- Any regularities in the Hume world, however seemingly systematic, are literally inexplicable. Asked why they occur, the best one can say is that it is sheer coincidence.

- Rational agents in the Hume world (assuming it is possible for there to be any) ought, if apprised of their sorry situation—that is, informed that they reside in a world governed by no laws—to make no plans for the future. For they ought to view the evidence of past regularities as providing no rational basis for expecting these regularities to persist.

That such claims are true of the Hume world is a perfectly plausible consequence of the anti-reductionist position. But it would be question begging for the anti-reductionist to judge that, since a reductionist claims that the world we live in is the kind of world that *she* would describe as a Hume world, it follows that the reductionist must *himself* be committed to the view that counterfactuals typically lack truth values, that regularities cannot be explained, and that induction is irrational. What we have here, in short, is a difference in outlook on one metaphysical issue—what are laws?—that is so fundamental that it automatically infects the respective parties' outlook on a range of other metaphysical and epistemological issues.

A useful comparison might be the following: A dualist thinks that there is a fundamental metaphysical difference between the conscious and the non-conscious. Presumably, this difference has all sorts of ethical implications. For example, by the dualist's lights, it might be perfectly ethically permissible to treat a genuine zombie in ways that it would be unimaginably horrific to treat a genuinely conscious human being—and this, quite regardless of whether the zombie displays the same overt physical behavior that a genuinely conscious human being does (*appears* to be writhing in agony, and so on). Suppose that is right, and imagine a dualist who argues against her physicalist opponent as follows: “According to you, we are all just zombies. But if so, then it doesn't matter how we treat each other. So your view has unacceptable ethical consequences, and therefore stands refuted.” That argument is Just. Plain. Silly.

Verdict: The off-cited complaints about reductionism that it cannot give a proper account of counterfactuals, explanation, and induction are, likewise, silly.

§5.4. *Nomological possibility*

The next complaint—that the reductionist position yields a mistaken account of nomological possibility—is silly for a slightly different reason. The idea is this: an anti-reductionist will presumably hold that facts about what the laws are are themselves nomologically necessary, so that any world *w* and any world *w'* that is nomologically possible relative to *w* must share the same laws. That does not follow

from a reductionist position. Indeed, the BSA straightforwardly contradicts this claim. For example, our Newtonian particle world will count as nomologically possible a world with a single particle moving with constant velocity. But the best system for such a pared down world will likely include the claim that all particles move with constant velocity—and so will not count our own world, rich as it is in complicated particle motions, as nomologically possible relative to *it*.

Now, it is important to note here something that is very often overlooked in the literature, which is that the intuition that is being violated is quite distinct from the central anti-reductionist intuition discussed in §5.1. One way to see this is to notice that it is quite easy to construct a reductionist position that respects the given intuition. For example, we could begin with the BSA, and then define a nomologically possible world as a world in which the best system for our world is not only *true*, but is also the *best system*: that immediately secures the desired result that nomological possibility consists in the perfect sharing of laws.

Still, all we really accomplish by this maneuver is moving the bump in the intuition-carpet around a bit. For we will have preserved the intuition about nomological possibility at the price of saying strange things about which specific worlds are nomologically possible. In our Newtonian example, we get the conclusion that any nomologically possible world is a world with Newton's laws, but at the cost of saying that it is nomologically impossible that there be just a single particle.

How bad an intuitive cost do we have here? Not that bad, as far as I can see. The best way to pay it, I think, is for the reductionist to opt for the less revisionist account of which worlds are nomologically possible (so that the single-particle world counts as nomologically possible), admit that some of the nomologically possible worlds will have different laws—and then go on to challenge his anti-reductionist opponent to demonstrate how this slightly unusual position about nomological possibility will make a difference to the way in which the notion of law is used in actual scientific practice.

I doubt she will be able to succeed. For example, she might exhibit certain counterfactuals about which the reductionist is forced to say odd things, e.g. this one: “If the world had contained just one particle, then it would have been such that, if it had contained a second particle, then both particles would have moved with constant velocity.” According to the reductionist, a nomologically possible world with just one particle is, plausibly, one in which it is a law that all particles move with

constant velocity; so it is a world in which the embedded counterfactual is true. That appears to make the counterfactual as a whole true, which is, admittedly, odd.¹⁶

On the other hand, such nested counterfactuals do not appear to have any significant role to play in the central scientific practices of prediction and explanation. What's more, the reductionist could propose a special semantical rule governing the evaluation of such counterfactuals, according to which one must, if possible, stay within the sphere of nomological possibility—*actual* nomological possibility—when assigning them truth values. So the given counterfactual would be evaluated as follows: First, we find the closest nomologically possible world that contains just one particle. Then, we find the closest world to *it* that (1) is nomologically possible relative to the *actual* world; and (2) contains two particles. Since this is a world in which the embedded consequent is false, the counterfactual as a whole is false.

Verdict: The anti-reductionist intuition about nomological possibility poses no serious threat to reductionism.

§5.5. *Objective chance and subjective probability*

A closely related problem concerns the reductionist treatment of objective chance, and an alleged conflict with a certain plausible principle relating chance to *credence*. To set the stage, recall the reductionist account of the chance-facts that obtain in the simple version of a chancy Newtonian world: Particles sometimes collide; when they do, they sometimes rebound elastically, and sometimes stick together; the frequency data show no (simple) dependency of outcome type on the parameters that characterize collisions; so the best system, plausibly, will describe each collision as having a fixed single-case chance for each outcome, where these chances are independent of one another. The *values* for these chances, finally, will simply equal the corresponding *frequencies* of each outcome type. In particular, let f be the frequency of elastic collisions; then for any time immediately before a collision, the chance, at that time, that that collision will be elastic is f .

Focus on time zero for this world. The laws, we are supposing, are an amalgam of the Newtonian laws and the probabilistic collision law. Together, they will fix a probability distribution for various possible future histories: the histories that have, at time zero, some chance of coming about. (And this probability distribution is really all we need: chances at *later* times derive from it by conditionalizing on intervening history. See my 2004.) Among these histories will be ones in which the

¹⁶ And it is no help to shift to the version of reductionism that makes nomological possibility symmetric; that just turns the given counterfactual into a counterlegal, which is even more odd.

frequency of elastic collisions *differs* from *f*. But if one of *these* histories comes about, then the best system for the world that results will *not* be the *actual* best system; instead, it will be a system that includes a probabilistic collision law that *contradicts* the actual collision law, by specifying different single-case chances. In short, there are worlds *w* that have laws different from the actual laws, and more specifically have at time zero chance distributions over possible futures different from the actual chance distribution—but that count as not merely nomologically possible (i.e., in conformity with the actual laws, in their nonmodal respects), but possible in the stronger sense that the actual laws assign them some non-zero chance of coming about. For short, the Humean’s probabilistic laws say that there is some chance that they *won’t* turn out to be the laws.

As before, the reductionist can avoid this outcome by a crude but effective maneuver: stipulate that a world counts as nomologically possible iff it has the same best system as the actual world, and then modify the chances by conditionalizing the time-zero distribution on the proposition true in all and only nomologically possible worlds. For example, in our chancy Newtonian world, the resulting distribution will say that only those future histories in which the frequency of elastic collisions is *f* have any chance of coming about. But as before, the crude maneuver creates new problems: in this case, it yields objective chances that behave in exceedingly funny ways. (E.g., the collision chances won’t remain constant, from trial to trial. See Arntzenius & Hall 2003 for more details.) Better, I think, just to live with the peculiar consequence that the laws assign a non-zero chance to some “undermining” futures.

Now, just how peculiar a consequence *is* that? First, some perspective: remember that a reductionist is *already* committed to denying a very firmly entrenched intuition about chances, that we all share: namely, that it is metaphysically possible for *distinct* chances to give rise to exactly *the same* phenomena. (Our LOPP thinks so; that is why, when fully informed about particle behavior in our chancy Newtonian world, she nevertheless assigns a credence *distribution* over a continuous range of hypotheses about the collision chances.) He owes us some good reasons for abandoning this intuition; we’ll come back to what those might be in a moment. But for now it’s enough to point out that if he can provide them, they will surely compensate, in addition, for the somewhat less blatantly anti-intuitive result that there can be undermining futures.

Unless, that is, this result leads to a contradiction. Lewis, and many others, once thought so. For it seemed a plausible principle relating subjective degrees of belief to objective chances that, roughly, if one knew what the objective chance was of some given outcome, and didn’t have any additional “inadmissible” information about that

outcome (for example, news from the future as to what that outcome *would be*), then one's subjective degree of belief in that outcome should equal this chance value, on pain of irrationality (see Lewis 1980 for the classic discussion). But now transport yourself to time zero in our chancy Newtonian world, and imagine that you are apprised of the time-zero chances, and of nothing else. Then according to our principle, you should assign some non-zero credence to the prospect that future history exhibits a frequency of elastic collisions *different* from f . But wait—you're a Humean. So you think—indeed, are *certain*—that the chance-facts of which you have been informed are made true by the total pattern of nonmodal facts that is about to unfold. More specifically, you are certain that if that pattern includes a frequency of elastic collisions different from f , then those chance-facts could not possibly be what you have been told they are. So you should assign *no* credence to the prospect that future history exhibits a frequency of elastic collisions different from f . Contradiction! We can conclude, it would seem, that the reductionist's position is, even if true, quite literally unbelievable. (Lewis famously called this problem the “big bad bug”.)

Not so fast. The argument errs, although in a way that is easy to miss.¹⁷ Since this part of the story has been well rehearsed in the recent literature (see my 2004), I will be brief. The credence-chance principle just alluded to contains a mistake: it fails to properly incorporate the way in which so-called “inadmissible” information ought to affect credence, when chances are known. Properly amended, the principle no longer says that the presence of inadmissible information *breaks* the connection between credence and chance; rather, it gives a precise formula for how that information *changes* the connection. The crucial additional observation is that for the reductionist,

¹⁷ There is one error, all too often attributed to it in the literature, that it does *not* make, which is the “error” of failing to take note of the contingent nature of the thesis of Humean supervenience. (See for example Vranas 2002; this is an example of the “mischief” alluded to in §2.) It is true that Lewis himself held his thesis to be contingent. He did so in part for good but irrelevant reasons: the official thesis says, among other things, that the only perfectly natural relations there are are spatiotemporal relations, and Lewis granted the possibility of others. And in part for bad but irrelevant reasons: he thought that Armstrong's “spinning sphere” thought experiment shows that there are worlds in which some things persist through time by being wholly present at each moment of their existence (and so not by being composed of momentary time-slices), and he thought that two such worlds could differ with respect to what is true at them (in one, the sphere is spinning; in the other it is stationary) without differing with respect to their Humean facts. Never mind why this is a bad reason; neither it nor the good reason provide the slightest grounds for thinking that Lewis thought of the Humean's account of *chance* as contingently true. And even if he *did* think this, he *shouldn't* have: it is a ridiculous mistake for a reductionist to concede that an anti-reductionist account of chance is metaphysically possible (essentially for reasons we've already rehearsed: how could he then claim to know that it wasn't *actual*?). And anyway, conceding this possibility wouldn't really help in avoiding the contradiction presented in the text (I'll leave it as an exercise to show why).

information about chances is *itself* “inadmissible”, since it is in part information about the future history of outcomes. (So the original problem wasn’t really that the reductionist account of chance *contradicts* a plausible credence-chance principle, but rather that it renders that principle useless.) That is no problem, since the revised credence-chance principle can smoothly handle this inadmissible information: it says, roughly, that if one knows the chances, then one’s credence should equal those chances, *conditional on* one’s total evidence. In knowing the chances, one’s total evidence includes (according to the reductionist) information about the future: in our example, it includes the information that the frequency of elastic outcomes will be *f*. The *unconditional* chance of some different frequency is non-zero; but the chance, *conditional* on this information, is of course zero. So too, our agent’s credence. There is no contradiction, after all. The icing on the cake, finally, is that the revised credence-chance principle can be motivated—motivated extremely well, in my view—on grounds *entirely independent* of the metaphysical debate over whether chance-facts reduce to nonmodal facts. So there is not even room for the complaint that the reductionist avoids the “bug” only at the cost of adopting an unattractive credence-chance principle.

Questions involving the epistemology of chance prove to be a red herring. But we would still like to know what advantages of the reductionist position might compensate for its unintuitive *metaphysical* consequences. Here, one finds the usual complaints that the anti-reductionist alternative is too obscure, supplemented by a *new* complaint that it is in fact the *anti-reductionist* who runs into trouble with credence-chance principles. Thus Lewis:

Be my guest—posit all the primitive unHumean whatnots you like. (I only ask that your alleged truths should supervene on being.) But play fair in naming your whatnots. Don’t call any alleged feature of reality “chance” unless you’ve already shown that you have something, knowledge of which could constrain rational credence. I think I see, dimly but well enough, how knowledge of frequencies and symmetries and best systems could constrain rational credence. I don’t begin to see, for instance, how knowledge that two universals stand in a certain special relation *N** could constrain rational credence about the future coinstantiation of those universals. (Lewis 1994, p. 484)

That there is almost nothing here by way of argument has not, regrettably, prevented Lewis’s opinion from gaining widespread acceptance (for a typical example, see Schaffer 2003). That reaction is long overdue for a correction. Having discussed this issue before (in Hall 2004), I will again be brief. Lewis overlooks three key points. First, any sane epistemology must accept that there are some constraints on rational credence that cannot be explained in terms of other such constraints; it is open to the anti-reductionist to treat the appropriate credence-chance principle as

one of these. Second, there are versions of anti-reductionism about chance that *can* explain the normative force of the credence-chance principle in terms of other normative constraints on credence—for example, the “primitivist hypothetical frequentism” discussed in Hall 2004. Third, and most significantly, a close look at worked-out examples strongly suggests that *reductionism* lacks the resources to explain credence-chance principles in terms of other principles; “dimly” turns out to fall far short of “well enough”. So there is really no point in pretending any longer that reductionism holds some special advantage over its rivals when it comes to accounting for the relationship between credence and chance.

It does, however, appear to hold an advantage in another arena, which is in the treatment of statistical mechanical probabilities. Suppose you put an ice cube in water; there is an astronomically small but nonzero probability that it won’t melt. Now, what sort of claim is that? One unhappy alternative: the interaction between ice cube and water is a genuinely indeterministic process, and the claim simply reports the objective chance. Another, slightly happier but still pretty unhappy alternative: it is a claim about the degree of belief it is reasonable to have that the ice cube won’t melt. And a much better, truly contented alternative: it is a probability that derives from a probability distribution over microstates compatible with the initial macrostate of the universe, conditionalized on facts about the present macrostate of the universe (including, in particular, the macro-facts about the ice cube and the water; see Albert 2000 for an explication of this treatment of statistical mechanical probabilities). What is needed next is an account of the status of this ur-distribution. At this point, the anti-reductionist seems at a loss, for it seems essential to her position that objective probabilities can only characterize irreducibly stochastic processes, and by definition no such process can result in a setting for the initial microstate of the universe. By contrast, there is at least room for the reductionist to explore the possibility that a candidate system for a world—even a candidate system that posits deterministic dynamical laws—might significantly increase its informativeness by positing, in a compact form, a probability distribution over initial states. Remember, furthermore, that the truth conditions for a claim that there is such-and-such an ur-distribution will be given by the way in which the nonmodal facts single out the candidate system making this claim as best; so the reductionist is not at all committed to a way of conceiving of chances that makes it unintelligible how there could be probabilities over initial states. This point, I think, displays a genuine and significant advantage for the reductionist. Loewer, who has done more than any other philosopher to develop the ideas sketched here (see for example his 2001), likes to say that the key advantage of the BSA is that it can make room for

laws about initial chances. But I don't think that is quite right; rather, the key advantage is that it can make room for *intelligible claims* about initial chances.

Verdict: The reductionist runs roughshod over certain entrenched intuitions about the metaphysics of objective chance; but he has something definite to show for it.

§5.6. *The problem of the ratbag idealist*

The next challenge to consider was already raised in one of the passages quoted above from Lewis; we will call it the “problem of ratbag idealism”. Here, I think, Lewis and other reductionists have in fact missed a chance to perform a nifty judo move on their opponents. Recall the challenge: standards for judging candidate systems—in particular, the standards of simplicity—are supposed to be “up to us” in a way that exposes a deep and offensive subjectivity in the BSA. It is at this point that the second guiding idea—which, remember, can be phrased as the claim that the standards for judging candidate systems should be the very standards that an anti-reductionist endorses as appropriate *epistemic* standards for figuring out what the laws are—comes into its own.

Here is the idea. Perhaps it was a mistake for the BSA, as set forth by Lewis and others, to implicitly assume that the simplicity of a theory is in fact one of the things we focus on as a guide to its truth. If so, the problem is not with the core idea behind the BSA but with its particular implementation. So suppose otherwise. Suppose, that is, that *everyone* should agree that one of the considerations our LOPP should use in deciding which of the various law-hypotheses is correct is its simplicity (when expressed in accordance with some appropriate scheme of representation). The anti-realist insistently reminds us that these standards are merely *epistemic* standards, standards such that we ought, as a normative epistemic sense matter, to believe that hypothesis that scores highest with respect to them. But then the problem of the ratbag idealist is very much *her* problem. In fact, it is even *more* her problem. For consider: If it is somehow “up to us”, in an objectionable sense, what counts as simple, then by the anti-reductionist's lights, central facts of normative epistemology are *also* up to us: facts about what one ought to believe about the laws, in the face of empirical evidence. To parrot Lewis: “Now, some ratbag idealist might say that if we don't like the misfortunes that our epistemological principles have visited upon us, we can change these principles—in fact, we can make them always have been different—just by changing the way we think!”

That, I think, is a much deeper lunacy than the kind Lewis seeks to evade. In effect, the anti-realist who endorses the ratbag idealist challenge to the BSA is committing herself to following position: “I believe there are facts about the world

that logically go beyond the sorts of facts that even in principle could serve as our evidence. Still, we can have good reason to think that we have got these facts right, if we construct theories about them that, among other things, score well on our standards of simplicity. Of course, what counts as ‘simple’ is a matter having to do with highly idiosyncratic features of human psychology, features that we could change if we wanted to. All the same, I am committed to treating simplicity of a theory as an epistemically good guide to its truth.”

That really does strike me as crazy. By contrast, the reductionist position seems perfectly sane, at least by its own metaphysical lights. And here, just to be clear, I am parting company with Lewis. I think he was deeply mistaken to think that there was any danger whatsoever posed by the problem of ratbag idealism. To see why, we need to step back for a moment, and reflect on the implications of the basic metaphysical outlook adopted by the reductionist. It helps to keep firmly in mind that he thinks that for any world, *all there is* to that world is a distribution throughout space and time of various perfectly natural magnitudes. For example, all there is to our Newtonian particle world are some particles moving around, with masses and charges. That’s it. It is emphatically *not* that the facts about these particles serve as clues to something “behind the scenes” that is directing their behavior. That is quite the wrong way to think about it. In fact, a much better way to think about the status of laws, given such a background metaphysics, is *pragmatically*.

Here is an example, that will help clarify what I have in mind: Suppose you are a biologist, and you have long since learned to renounce vitalism. You face the question what, given your anti-vitalist commitments, should count as “life”. The first thing you should do is to scrap that question, and replace it with a slightly different one: What, given your anti-vitalist commitments, is the best—i.e., theoretically most useful—way to draw a distinction between things in the world that might roughly correspond to our ordinary living/nonliving distinction? You might find that there are a number of candidate ways to draw this distinction. You might find that the choice between them should be made entirely on grounds of theoretical utility—indeed, that for some purposes one way of drawing the distinction is useful, for other purposes another. So too, I claim, with laws—at least, given a reductionist metaphysics. We have an array of nonmodal facts about a world—as it might be, facts about how many particles there are, what their masses and charges are, and how they are moving. Science, presumably, is in the business of investigating *that* stuff. It is not in the business of investigating any further stuff that lies behind the scenes, for the scenery constitutes all of reality. So do not ask what, given such a metaphysical outlook, laws are. Ask instead how, given such a metaphysical outlook, one might usefully draw a distinction between certain facts about the world that are in some

sense *distinctively appropriate targets for scientific inquiry*, and other facts that are less interesting and central.

As far as I can see, for a thoroughgoing reductionist about “laws”, the only interesting sense in which that word picks out a subject matter worthy of philosophical attention is if it demarcates a class of claims about the world that are, somehow, *distinctively appropriate targets for scientific inquiry*. It is, of course, a good and hard question what would make some claim a *distinctively appropriate target for scientific inquiry*. We will look into this issue in a bit more depth below, in §6. For now, I simply wish to note that there is evidently plenty of room for the view that it is in part facts about us—*idiosyncratically* about us, and our peculiar human psychologies—that play a role in determining the most appropriate way for us to structure our investigation of the world. How could the details of our peculiar human situation *not* be relevant to this matter?

Here, the contrast with the anti-reductionist could not be sharper. On her view, the laws governing a world demarcate a range of facts about the world profoundly, fundamentally distinct from any other kinds of facts, a range of facts with deep and interesting metaphysical connections to the realm of nonmodal facts. And given that they are empirically accessible at all, only the most remarkably incurious person could fail to think that scientific inquiry ought to be structured in part with the aim of uncovering them. By contrast, a reductionist cannot think of laws as having some such natural epistemic magnetism to them. To put the point pithily, for the anti-reductionist, the nature of laws *automatically* makes them *distinctively appropriate targets for scientific inquiry*; whereas for the reductionist, it is exactly the reverse. On his view, all it could come to for a claim to achieve the status of law is that, for some other and *independent* reason, it is a *distinctively appropriate target for scientific inquiry*. Why should he worry in the slightest if those other reasons make crucial reference to highly idiosyncratic features of human psychology? It would be as if a biologist lost sleep over the thought that, if the distinctive interests of the community of biologists were different, they might draw the distinction between living and nonliving things differently. Yes, they would. And that would be perfectly appropriate. By the same token, if the community of scientists (or perhaps just physicists) sees fit to change its standards for what counts as simple, then according to the reductionist there ought to be a corresponding shift in what counts as a fundamental law—that is, what counts as a *distinctively appropriate target for scientific inquiry*. So what?

Verdict: It would have been smarter for the anti-reductionist to keep quiet about the ratbag idealist. But taking stock of the problem this creature poses allows us to

significantly sharpen our understanding of the differences between reductionism and anti-reductionism. (§6 will take advantage of this improved understanding.)

§5.7. *A new skeptical problem?*

Our next problem comes from van Fraassen's *Laws and Symmetry*. Or at any rate, it is inspired by van Fraassen; his own discussion is sufficiently obscure that I'm not sure I've captured his intent. No matter. The problem I will set forth is sufficiently interesting, whether it is the one van Fraassen had in mind or not.

Imagine two worlds. The first is our own, well-behaved Newtonian particle world. The second perfectly matches our world in certain key respects: in particular, it has the same number of particles, and they follow exactly the same trajectories as the particles of our world. But the masses and charges are completely different—swapped all around, we can suppose, in such a way that there is absolutely no hope of writing down simple equations involving mass and charge that will capture a lot of information about particle motions in this world. Such a world is clearly possible. In fact, countless such worlds are possible: just hold fixed the particle trajectories in our world, and change their other magnitudes. (One could even add brand new magnitudes.)

Why is this a problem? Well, it seems to confront us with a skeptical dilemma. Here we are, as scientists in our world, trying to figure out what its laws are. We have no direct access to the perfectly natural magnitudes that characterize our world at the microscopic level. What we do have access to are motions of large things, which of course are determined in a fairly direct way by motions of small things. But the screwed-up particle world we are imagining reproduces these motions exactly. That is, it is completely, and in principle, observationally indistinguishable from our own world. So what possible reason could we have for thinking that we are in a well-behaved Newtonian particle world, as opposed to a particle world with randomly assigned, and therefore horribly behaved, perfectly natural magnitudes? The Lewisian appeal to perfectly natural magnitudes thus seems to erect an insurmountable epistemic barrier between us and knowledge of, our even reasonable belief about, the fundamental laws.

I see two responses available to the reductionist. The first draws on the quasi-instrumentalist treatment of mass and charge sketched in §5.2. There, it was offered in part as a way of mollifying essentialist intuitions about the connection between magnitudes such as mass and charge and the fundamental laws governing them. But it will, obviously, work equally well to dissolve the skeptical worry we are considering here.

There is, however, a second way to deal with that worry, not by dissolving it, but by exposing it as one we should be happy to live with. It is well to remember that it is not always a bad thing if one's metaphysical commitments engender skeptical possibilities. For example, I think that there is a fact of the matter about what our world is like in its unobservable aspects—and that what it is like is certainly not logically determined by what the observable aspects are like. That opens the door to the skeptical possibility that, in spite of paying close attention to our best observational evidence, we get facts about the unobservable badly wrong. Yes, there obviously is such a skeptical possibility. And it is equally obvious that its mere existence does not establish the bizarre epistemic doctrine of anti-realism about the unobservable. By the same token, the mere existence of worlds that perfectly reproduce the motions of our world, but whose other perfectly natural magnitudes are horribly ill-behaved, does not by itself establish anti-realism about laws as the BSA conceives them.

There is a bit more to say, and saying it helps to show not merely that van Fraassen hasn't yet made his case, but that further effort is unlikely to help. For consider how we ought to describe our messed-up particle world, from the perspective of the BSA. It is, after all, a world that has its own laws. Given the way that masses and charges are swapped around, those laws are certainly not the Newtonian laws. They are something else. There seem to be two possibilities. The first is that there is a unique and highly informative best system for this world. The second, and more likely possibility, is that there is no candidate system that is both clearly superior to its rivals and highly informative. If the first possibility is correct, then the right way to describe this world is as follows: It has powerful laws, very different from our own, that combine with initial conditions so remarkably chosen that the motions that result look just like Newtonian motions. How remarkably chosen? Well, if any aspect of those initial conditions had been even slightly different, then the motions would have been manifestly non-Newtonian. But look, that is a skeptical possibility that we should *all* recognize. *Of course* it is possible to have very different laws, that permit a range of initial conditions, among which there will be at least one that is highly misleading—in the sense that it would, under those laws, evolve in such a way that the resulting history yielded overwhelmingly powerful evidence that the laws were something else. Our *own* laws are like that. For example, there are initial conditions compatible with our laws that would result in locally anti-thermodynamic behavior; e.g., initial conditions that would lead to a history in which every ice cube put in a glass of warm water *grew*. Of course, the measure of such initial conditions is astronomically small. But nevertheless, they exist. And if one of them had obtained, we would have been led badly astray about what the laws of our

world are. By the same token, the scientists in the screwed-up particle world are badly wrong about their laws—but only because of what is, given the laws that *in fact* govern their world, a fantastically unlucky choice of initial conditions.

The second possibility receives a similar treatment. According to it, our screwed-up particle world simply lacks substantive laws, because the best system does not score very well on the standard of informativeness. Most likely, these laws will be non-deterministic (and not probabilistic either). So, given an initial condition, all they will manage to do is to lay down broad constraints on how that initial condition can evolve forward in time. So the right thing to say about this world is that, from its initial conditions, it just so happens—purely by accident—to evolve in a way that looks Newtonian. And again, this is a possibility we ought to recognize, *regardless* of our commitment to the BSA. We *think* that our world comes equipped with very powerful laws, perhaps even deterministic ones. But it *could* be that its laws are extremely loose, permitting a vast range of non-Newtonian physical behaviors that, through sheer luck, never happen to manifest themselves. It would be a foolish overreaction to design our metaphysics so as to rule out this possibility.

Notice that in each case, it is not just that we find a way of describing the skeptical possibility that makes it seem quite mundane. In addition, these redescription reveal the skeptical possibilities to be ones that we have very good epistemic grounds to dismiss, assigning them vanishingly small shares of credence. The reasonable conditional probability, on the assumption that our world contains only very loose laws, that we would see perfectly Newtonian behavior is negligibly different from zero. Likewise, the conditional probability, on the assumption that our laws are quite strict but radically non-Newtonian, that the initial conditions would be so precisely tuned as to reproduce Newtonian phenomena, is negligibly different from zero. If you doubt any of this, remember our example of the ice cubes: Yes, it is possible, even holding fixed what we think our laws are, that the next time you put an ice cube in warm water, it will grow. But the appropriate degree of credence to assign to this prospect is so small as to make it effectively ignorable.

Verdict: van Fraassen thinks that a serious skeptical problem bites the reductionist who endorses a Lewisian metaphysics of perfectly natural magnitudes; but close inspection reveals that this problem is wholly toothless.

§5.8. *The information that counts*

We come now to a problem that, I think, requires a substantial amendment to the BSA. To set things up, let us suppose that the BSA, when applied to our Newtonian particle world, succeeds to this extent: the best system—the one, remember, that optimally combines simplicity and informativeness—includes the

Newtonian dynamical laws. And let us suppose it also includes a broad constraint on initial conditions, so that the combined result is that it counts a world as nomologically possible only if that world contains a finite number of particles, each possessing some value of mass and charge, and moving on trajectories that perfectly conform with the Newtonian dynamical principles. So far, so good. The problem is that there is every reason to think that this “only if” cannot be strengthened to an “if and only if”; that is, that the best system for our Newtonian world will exclude too many worlds that ought to count as nomologically possible.

A quick way to see this is to notice that we get a vast increase in informativeness, at only a small cost in simplicity, if we include in a candidate system a statement that says exactly how many particles there are. Only *small* cost in simplicity? Well, the cost might be modest, if the number of particles is a large number that is not mathematically identifiable in a simple way. But suppose it is. Suppose it turns out that there are exactly 2^{64} particles in our world. Then a statement to that effect buys a lot of informativeness at a negligible cost in simplicity. And it seems quite wrong to think that, *for that reason*, it ought to count as nomologically necessary that there are 2^{64} particles. In the first place, particle number should turn out to be a nomologically contingent feature of the world. In the second place, even if you have somehow convinced yourself that there can be good grounds for considering particle number to be a nomologically necessary feature of the world, surely those grounds cannot be that this number is easily mathematically describable.

At any rate, there are more pressing versions of this problem. Suppose, for example, that there is some moment of time such that the complete physical state of the world at that time happens to have a very simple to describe form: there is some relatively simple, compact way to say exactly what that state is. Let this state be *S*. Then, if a candidate system includes the Newtonian dynamical principles, it buys an *enormous* increase in informativeness by adding a statement to the effect that at some time, the complete physical state of the world is *S*. For doing so will shrink the set of nomological possibilities down to *one*. (Here I am taking advantage of the fact that Newtonian dynamics are two-way deterministic.) But that is a disaster, for remember that it is a nonnegotiable desideratum on an account of laws that it yield a nontrivial distinction between what is nomologically possible and what is not.

Finally, it is a disaster that is unavoidable, essentially because it is guaranteed that our world exhibits a simple to describe state. For considering that the state of the world at any time can be coded up, in a very simple way, by a single real number: just take all the coordinates, masses, and charges of all the particles, expressed in decimal notation, and interleave the digits. Suppose we include this number in a candidate system; then once again, we get an increase in informativeness that shrinks the set of

nomological possibilities down to one. Call this the problem of *the phony fundamental constant*.

“But,” comes the response, “don’t we get this increase in informativeness at the price of a huge increase in complexity? For even though the way in which this real number codes up complete information about a physical state is simple, the number *itself* is very unlikely to be at all simple. That is, there will be no mathematically compact way to specify its value.”

I do not think there is any refuge in this response. Remember that it is part of the practice of physics to include in its theories *fundamental constants*. But it is *no* part of the practice of physics to insist that those constants have mathematically tractable values. Pick any mathematical standard of complexity of number you like, and a physicist will say that for all she knows, the fundamental constants might be *that* complex. (And do not think that we can buy mathematical simplicity by a suitable choice of units; the quickest way to see that this won’t work is to remember that some of the constants appearing in our current physical theories are unitless.) Yet it is nomologically necessary that the constants have the values they do.

More to the point, there is something silly in all of these defenses, for they miss what is really going wrong in these problem cases, which is that a candidate system is scoring points for the *wrong kind of informativeness*. Recall how, in section 1, we factored our Newtonian specification of nomological possibilities into two parts: there is an Initial Conditions Hypothesis, that specifies what initial conditions are nomologically possible; and there is a Dynamical Hypothesis, that specifies how each such initial condition would evolve. The problems we are considering have the common feature that a candidate system wins merely because it narrows down the range of possibilities allowed by the ICH. That, I claim, is an increase in informativeness that ought not to count in favor of a candidate system that exhibits it.

Stronger: it ought to count *against* a candidate system that exhibits it. (This idea is not new: compare Earman’s discussion in his 1984.) Part of the reason involves an appeal to intuition: facts about initial conditions seem, intuitively, to be the sorts of things that a good physical theory will treat as contingent. But a better reason focuses on the central explanatory practices of physics, which make essential use of counterfactuals concerning worlds with different initial conditions, but governed by the same dynamics as our world. What, for example, explains the fact that the planetary orbits deviate from perfectly elliptical orbits? The presence of other planets, and corresponding perturbation by interplanetary gravitational forces: for *if all other planets were absent*, then a given planet would orbit the sun in a nearly perfect ellipse. More generally, I claim, the ability to provide sharp and determinate truth

conditions for a wide range of counterfactuals is precisely what lends a good physical theory its explanatory power. So, by unduly restricting the ICH, a candidate system will hamstring its ability to evaluate counterfactuals in a way that fatally undermines its explanatory potential. So it was a serious mistake for the BSA to incorporate a standard of informativeness according to which a candidate system automatically improves its score by narrowing the range of worlds it counts as nomologically possible. This is not the kind of informativeness that counts; and what's more, once this standard is invoked, there is no way to avoid the devastating conclusion that the best system for our Newtonian particle world will count that world *but no others* as nomologically possible.

Verdict: The official formulation of the BSA needs a substantial revision, in the form of an account of informativeness that is more sensitive to the actual practice of physics. The next section provides just such an account, and produces an elegantly amended BSA. But victory for the reductionist will be short-lived, for he will be hard-pressed to explain why, by his metaphysical lights, the amended BSA should be preferable to the original version.

§6 A solution and a problem

§6.1. *Solving the problem of too much information*

If the diagnosis in §5.8 is correct, then it is easy to see how to produce an improved version of the BSA. Here is a sketch:

Let a candidate system have two parts: an ICH, and a DH. A system is better to the extent that each of these parts can be specified in a simple way. But in the case of the ICH, a system is better to the extent that this hypothesis is *un*informative—i.e., to the extent that it admits *more* initial conditions as nomologically possible. Where we want a candidate system to be informative is, instead, solely in its dynamical hypothesis. And, with respect to the informativeness of the DH, we can say something much more informative about what we mean by “informative” than Lewis did. That is because there is a clear, non-arbitrary standard of informativeness for a dynamical hypothesis to aim for: determinism, in both temporal directions.

Next, there are fairly clear and non-arbitrary ways of falling short of the ideal of determinism. For example, a DH might be deterministic only towards the future. Or, it might be deterministic for all but a very small measure of initial conditions. Or it might be indeterministic only in a “punctuated” sense, so that only when certain specific and localized kinds of interactions take place does the DH fail to describe their outcomes in a deterministic fashion. (Our chancy Newtonian world was like that.) And even where the DH falls short of determinism, it might succeed in laying

down precise probabilities. It is, I think, a very interesting question how to develop a precise catalog of the varieties of less-than-perfect informativeness that a DH might exhibit. I will not pursue that question further here. Still, it is quite striking how, by shifting attention from the candidate system as a whole to just one component of it—its dynamical hypothesis—we have *already* arrived at a vastly more interesting characterization of “informativeness” than the utterly bland characterization routinely trotted out in the literature defending the BSA, that says only that informativeness is somehow a matter of narrowing down possibilities. It’s not; and even if it were, it would remain an embarrassment that the literature has been unable to say anything more illuminating.

Why is two-way determinism the gold standard of informativeness for a DH? It is not *just* that this kind of informativeness is pretty clearly what physicists (many of them, anyway) would say that they aim for in their theories. They have, in addition, an excellent reason for setting their sights on such theories, since meeting this standard directly serves the theoretical purposes emphasized in §5.8: a theory that does so can, as a result, answer with great precision questions both about what would have happened, and about what would have to *have* happened, had conditions in the world been different in some specified respect. That is one extremely important way for a theory to increase its exploratory power. The other way, of course, is for the theory to incorporate a maximally permissive ICH.

These observations suggest an amendment to the official guiding idea behind the BSA: what candidate systems aim for is not an optimal combination of simplicity and informativeness, but rather an optimal combination of simplicity and *explanatory power*—where explanatory power in turn is a function of the uninformativeness of the system’s ICH together with the informativeness of its DH. As before, the nomologically possible worlds will be exactly those that conform to the candidate system that wins out in this competition. (Or: that conform to each of those candidate *systems* that are tied for first. More, in a moment, on why we might need this formulation.) And as before, the desiderata we are seeking to maximize will be in some tension with each other. But—having “resolved” candidate systems into two distinguishable parts—we can say a bit more than defenders of the BSA usually do about the nature of this tension.

To begin, the most straightforward way to achieve a maximally uninformative ICH is to take advantage of the limited principle of recombination governing perfectly natural magnitudes, so that the *nomologically* possible initial conditions will exactly coincide with the *metaphysically* possible ones, or at least the metaphysically possible ones that exhibit only the perfectly natural magnitudes found in our world. In the case of our Newtonian particle world, that will get us an ICH that counts as

possible any initial condition in which there are a finite number of particles, each possessing some value for mass, charge, and position (and in which there is nothing else).¹⁸ But it will presumably get us many more possibilities. For example, it will include initial conditions in which there are an infinite number of particles—even uncountably infinite. And it may include initial conditions featuring entities *other* than particles, but having mass and charge. By doing so, it will make the job of constructing a DH that is both simple and informative much, much harder.

Consider some ways this might happen. First, it might be that our ICH (even when it is restricted so that the conditions it admits exhibit no magnitudes other than mass and charge) allows for possibilities that some candidate DH—as it might be, the DH that incorporates the standard Newtonian dynamical equations—simply doesn't apply to. For instance, one cannot, for straightforward *mathematical* reasons, even apply the equations in the Hamiltonian formulation of classical particle mechanics to an initial condition featuring an uncountable infinity of particles. Of course, a candidate DH could accommodate such an initial condition simply by adding an extra clause, and indeed could do so in quite a simple fashion: for example, it could stipulate that if the world starts out with uncountably many particles, then all of these particles immediately vanish. Both simplicity *and* determinism (at least, toward the future) are retained, and the ICH is allowed to include many more possibilities. But the problem is that there are far too many ways for a DH to be simple and informative, while still “covering” those possibilities featuring uncountably many particles. And that is because the particle behavior exhibited in the *actual* world gives, intuitively, no guidance as to how to treat such bizarre initial conditions.

One fix is to stipulate that what is nomologically possible is only what all such DH's *agree* is nomologically possible. That would have the odd result that there could, compatible with the laws of our Newtonian particle world, have been uncountably many particles; but that if there *had* been, pretty much anything could have happened (provided only that it was simple to describe). A better fix is probably

¹⁸ There is a complication here concerning velocity, which I will mostly pass by. The issue is this: facts about the velocities of particles, at a time, are plausibly *not* facts wholly intrinsic to the state of the world at that time. So if we think of our ICH as specifying a range of possible *momentary* states, then the DH we construct on top of it has, in our Newtonian example, no hope of being deterministic, in the sense that it specifies exactly one future, for each initial condition. (That's because the standard Newtonian equations need information not just about initial positions, but also about initial velocities.) But it comes infinitesimally close: for it is still the case that, for any two nomologically possible worlds, any time t , and any time interval $\Delta t > 0$, if the worlds agree on the positions, masses, and charges of all particles throughout $[t, t + \Delta t)$, then they agree throughout their entire histories.

to say that the ICH should be restricted, if doing so makes it possible to produce, in a *non-arbitrary fashion*, a simple and informative DH.

In our example, that requirement would plausibly lead us to rule out any initial condition containing infinitely many (even countably infinitely many) particles. But problems will likely remain. For example, should we allow an initial condition in which there are just two oppositely charged, massive particles, moving on a straight line toward each other? Alas, the usual equations blow up, in such a case.¹⁹ We could, perhaps, evade this problem by insisting that the initial conditions specify, along with their masses and charges, only the *positions* of particles, and not also their velocities. Then we could say that the simplest, most informative DH (which is the one whose dynamical equations are the standard Newtonian ones) specifies, for each initial condition, a *range* of possible forward evolutions: namely, forward evolutions compatible with each way of *tractably* assigning velocities to the particles. (This way of putting things makes it vivid that our Newtonian DH is not two-way deterministic, in the very strongest sense that it fully fixes the history of the world, given only a single instantaneous state at some time. But we knew that already; cf. footnote 18.) Whether this maneuver will work is beyond my technical competence to assess.²⁰ At any rate, the interaction between the need for a simple and informative DH, coupled with a simple and uninformative ICH, clearly merits further study.

§6.2. *A dilemma for Humean reductionism*

Our new, improved BSA immediately and effectively deals with the challenges raised in §5.8, concerning particle number, the possibility of a simple state, and the phony fundamental constant. What's more, it *seems* to do so in a perfectly well-motivated way: for remember that the intention of the BSA is to incorporate, as the standards for judging candidate systems, standards implicit in the best practice of physics. And those standards place a premium on physical theories that say a lot about what would have happened, under alternative possible conditions.

But now we come to a problem. If, once again, we pay very close attention to the background metaphysical commitments that underlie the BSA, we can see that it faces a curious but difficult dilemma.²¹ Building the ICH/DH distinction into the

¹⁹ And remember that there are other, more devious examples, where the equations apply, but where the particles end up accelerating to infinity; see footnote 2.

²⁰ It will work, just in case there is, for every configuration of masses, charges, and positions for finitely many particles, some way of assigning particle velocities such that the standard Newtonian differential equations yield a solution, when applied to these boundary values.

²¹ There is an additional cost—though it pales by comparison to the one about to be discussed: a reductionist who adopts the amended BSA will not, at least without some fancy footwork, be able to endorse the treatment of statistical mechanical probabilities discussed in §5.5. For it is part of that

BSA seems to be necessary, in order to prevent Humean reductionism from being too violently at odds with the actual practice of physics, and in particular with its aim of constructing theories that are maximally explanatory powerful. For example, if physicists treat particle number in a Newtonian world as a nomologically contingent feature of that world, then, at the very least, someone who disagrees with them on purely philosophical grounds is in a highly uncomfortable position. But, however successful the new version of the BSA is at dodging potential counterexamples—and in particular, in securing agreement with the physicists about which initial conditions should be treated as nomologically possible—it is, by the Humean’s lights, *entirely unmotivated*.

To see why, we should return to the conception of “law” emphasized in §5.6: some claim about the world counts as a law only to the extent that it is, on grounds *independent* of its nomological status, a distinctively appropriate target for scientific inquiry. Now, for the reductionist there is one obvious way to pick out a special set of claims: they will be claims about the kind and distribution of those perfectly natural magnitudes that characterize our world. (The anti-reductionist, too, can view these claims as special.) Such claims have a distinguished metaphysical status, insofar as all *other* facts about the world reduce to them. But, just as obviously, that is not *enough* of a distinction: one wants some *further* way of picking out, among all the true claims about the distribution of perfectly natural magnitudes, an elite subset that are distinctively appropriate targets for scientific inquiry. (Without some such further distinction, the reductionist’s account of laws will effectively collapse into the naïve regularity theory.)

An anti-reductionist can say that certain claims about the fundamental, nonmodal facts deserve special attention *because they hold as a matter of law*. And so it makes perfect sense to structure inquiry in science—and especially in physics—in such a way that it has a chance of uncovering these laws. But for a Humean reductionist, *all there is*, fundamentally, is an assortment of nonmodal facts—facts that are all perfectly on a par, metaphysically speaking. Scientific inquiry is, in the first instance, merely directed at finding out about *them*. It would seem, then, that any further structure to this inquiry can be imposed only by the need to find out about those facts *efficiently*. Helen Beebe has provided an especially eloquent articulation of this idea:

treatment that the best system for an entropically well-behaved world like ours will include a statement to the effect that the world starts out in a very low entropy macrostate. And that is an example of a restriction on ICH that the amended BSA cannot, on the face of it, accommodate.

So the idea is something like this. Suppose God wanted us to learn all the facts there are to be learned. (The Ramsey-Lewis view is not an epistemological thesis but I'm putting it this way for the sake of the story.) He decides to give us a book—God's Big Book of Facts—so that we might come to learn its contents and thereby learn every particular matter of fact there is. As a first draft, God just lists all the particular matters of fact there are. But the first draft turns out to be an impossibly long and unwieldy manuscript, and very hard to make any sense of—it's just a long list of everything that's ever happened and will ever happen. We couldn't even come close to learning a big list of independent facts like that. Luckily, however (or so we hope), God has a way of making the list rather more comprehensible to our feeble, finite minds: he can axiomatize the list. That is, he can write down some universal generalizations with the help of which we can derive some elements of the list from others. This will have the benefit of making God's Big Book of Facts a good deal shorter and also a good deal easier to get our rather limited brains around. (Beebe 2000, p. 547)

It makes perfect sense, given a reductionist metaphysics, that a distinctively appropriate target for scientific inquiry is to figure out exactly which “axioms” God has written down: for doing so directly serves what can be the only ultimate aim of inquiry, which is to figure out as much of the nonmodal truth about the world as possible (there being no other truth to be had). But that seems to point us, alas, directly to Lewis's original formulation of the BSA, with its premium on the kind of informativeness that consists merely in the narrowing down of possibilities. He suggested, in one of the passages we quoted, that the standards of simplicity and informativeness are “virtues that we aspire to in our own theory building”. If I am right about the proper diagnosis of the problems raised in §5.8, this suggestion is mistaken. While simplicity may be such a virtue, it is precisely uninformativeness that we aspire to in the part of our theory that tells us what initial conditions are possible. It is incomprehensible why this should be a virtue if, in the first instance, all that scientific inquiry consists in is an investigation of what the nonmodal facts are about our world.

So the Humean reductionist seems to face the following dilemma: On the one hand, he can avoid conflict with the practice of physics by structuring the BSA around the ICH/DH distinction in the way described above; but it is unintelligible why, by his lights, the resulting standards are the right ones to use in choosing a best system. On the other hand, he can choose standards that make perfect sense, given his metaphysical commitments; but doing so appears to lead straight to the original version of the BSA, a theory that from the standpoint of the actual practice of physics is insane. This choice between a guilty intellectual conscience and insane revisionism is not a happy one.

Consider, by contrast, what metaphysical viewpoint would motivate the importance of the ICH/DH distinction, and in particular an adherence to standards of theory choice that looks for an uninformative ICH, coupled with an informative

DH. It is not far to seek: it is the view that information about the counterfactual structure of the world is *sui generis*, and not merely disguised information about its nonmodal structure. If that is what you think, then *of course* you will want to build into your theorizing about the world techniques for accurately mapping this *sui generis* counterfactual structure. And that is exactly what the distinction between initial conditions and dynamics accomplishes. Notice, finally, that it is precisely a commitment to a *sui generis* counterfactual structure that unites *all* the disparate varieties of anti-reductionism. So the central feature of physical theorizing that we have focused on cleanly divides the Humean reductionist from his philosophical opponents. His position is not only philosophically uncomfortable, but lonely.

§7 Conclusion

As I noted at the outset, I think that the problem I have raised—challenging though I hope it is—is really an occasion for the Humean reductionist to sharpen his position still further. I am not sure what the best way is to do this, but here are some preliminary ideas (prompted mainly by conversations with Barry Loewer, who is one of the most thoughtful reductionists currently writing on the topic).

Focus just on the problem posed by the phony fundamental constant, the one that codes up complete information about the physical state of our particle world at some time. I suggested that the proper way to deal with this problem is to build the ICH/DH distinction into the BSA—but that this suggestion is really a Trojan horse, since the resulting account is one that the reductionist ought to reject, given his metaphysics.

Loewer has suggested (in conversation) a very different response. He observes that one serious defect in a system that incorporates the phony fundamental constant is that it codes up information about the state of the world in a form that is almost completely *useless* to us ordinary humans. To dramatize, suppose an Oracle starts to tell you the value of this constant, a few digits at a time. After she has told you a billion digits, you still have something that is virtually useless to you—even if you have unlimited computational power at your disposal. For this information allows you to make almost *no* new predictions. So perhaps the right lesson for the Humean reductionist is to return to Lewis's original formulation, but insist that what is wanted from a candidate system is not merely information that narrows down possibilities, but information that is, in some suitable sense—and taking into account relevant facts about the condition that we humans find ourselves in—*practically useful*.

One way to spell this idea out focuses on Beebe's remarks about our "rather limited brains". We might reasonably demand that candidate systems pay heed to our

febleness, by representing the world in a way that not only optimally balances simplicity and informativeness, but is such that *approximations to it*—which, given our limitations, may be all that we can ever hope to achieve—are *also* highly informative. Then the real problem with the phony fundamental constant has nothing to do with whether its actual value is mathematically identifiable in a simple way. Even if it *were*, the system that incorporated this value would fail the test we are considering: other, “nearby” systems, that got the value of the phony constant slightly wrong, would do so at the price of losing almost all of the extra informativeness the correct system attains by means of the constant.

What I have just sketched seems to me a promising approach for the reductionist to take, although I have two lingering doubts. First, it seems clear that however this strategy is developed, it will end up making the importance that the practice of physics apparently places on the ICH/DH distinction seem misguided; it remains to be seen whether the reductionist can supply a plausible “error theory” of this practice, or whether instead his position ends up looking too revisionist for comfort. Second, it seems to me easy to come up with examples of statements about the world that ought to be viewed as nomologically contingent, but that nevertheless contain quite a lot of information of great practical utility to us. I imagine, for example, that cosmologists would very much like to know the total mass of the universe—and that even approximate information about this mass would be enormously predictively and explanatorily valuable. But would their ability to squeeze useful predictive and explanatory information out of this knowledge do anything to show that facts about the total mass should be viewed as nomologically necessary? To me, it seems not—although I freely confess that I have nothing but intuition to go on here, and that these intuitions may have been corrupted by evil anti-reductionist influences.

It is clear what the debate needs next: a version of reductionism that gets much more serious than Lewis ever did both about what, *precisely*, the standards are for judging candidate systems, and about why—given the reductionist’s metaphysical commitments—those *ought* to be the standards; and a version or versions of anti-reductionism that takes much more seriously the need to present a transparently clear conception of what the extra feature of the world is that goes beyond its distribution of perfectly natural magnitudes, and of how empirical evidence can be a reliable guide to this feature. In addition, the debate needs a shift in methodological priorities that places much less emphasis on intuitions about hypothetical cases, and much more emphasis on attending to the distinctive sort of work that a concept of “law of nature” performs in actual scientific practice. That is an approach that sensibly conceives of inquiry into metaphysics as continuous with scientific inquiry:

not because it is the scientists who really ought to be settling our metaphysical disputes for us (hah!), but rather because an understanding of the fundamental ontological structure of our world will be achieved, if at all, only by paying close attention to the structure of our scientific theories, and most importantly to the reasons why they should or should not have the structure they do.

§8 References

- Albert, David 2000. *Time and Chance*, Cambridge MA: Harvard University Press.
- Arntzenius, Frank and Ned Hall 2003. "On What We Know About Chance", *British Journal for the Philosophy of Science* 54: 171-179.
- Beebe, Helen 2000. "The Non-Governing Conception of Laws of Nature", *Philosophy and Phenomenological Research* 61: 571-593.
- Bogdan, R. J. ed. 1984. *D. M. Armstrong*, Boston: Reidel.
- Cartwright, Nancy 1999: *The Dappled World*, Oxford: Oxford University Press.
- Chalmers, David 2002. "Does Conceivability Entail Possibility?", in T. Gendler and J. Hawthorne (eds.), *Conceivability and Possibility*: 145-200.
- Earman, John 1984. "Laws of Nature: The Empiricist Challenge", in Bogdan ed. 1984: 191-223.
- Elga, Adam 2004. "Infinitesimal Chances and the Laws of Nature", *Australasian Journal of Philosophy* 82: 67-76.
- Gendler, Tamar and J. Hawthorne (eds.) 2002. *Conceivability and Possibility*, Oxford: Oxford University Press.
- Hall, Ned 2004. "Two Mistakes about Credence and Chance", *Australasian Journal of Philosophy* 82: 93-111.
- Hall, Ned 2005. "Causation and Ceteris Paribus Laws", *Harvard Review of Philosophy* 13: 80-99.
- Ismael, Jenann 1996. "What Chances Could Not Be", *British Journal for the Philosophy of Science* 47: 79-91.
- Lewis, David 1980. "A Subjectivist's Guide to Objective Chance", in Lewis 1986, pp. 83-113.
- Lewis, David 1983. "New Work for a Theory of Universals", *Australasian Journal of Philosophy* 61: 343-377.
- Lewis, David 1986. *Philosophical Papers, Volume II*, Oxford: Oxford University Press.
- Loewer, Barry 1996. "Humean Supervenience", *Philosophical Topics* 24: 101-127.
- Loewer, Barry 2001. "Determinism and Chance", *Studies in History and Philosophy of Modern Physics* 32B: 609-620.

- Maudlin, Tim 2003. "A Modest Proposal Concerning Laws, Counterfactuals, and Explanation", unpublished ms.
- Maudlin, Tim 2006. "Why Be Humean?", unpublished ms.
- Schaffer, Jonathan 2003. "Principled Chances", *British Journal for the Philosophy of Science* 54: 27-41.
- Shoemaker, S. 1980. "Causality and Properties", in van Inwagen 1980: 109-35.
- van Inwagen, Peter (ed) 1980. *Time and Cause*, Dordrecht: Reidel.
- Vranas, Peter 2002. "Who's Afraid of Undermining? Why the Principal Principle Might Not Contradict Humean Supervenience", *Erkenntnis* 57: 151-174.
- Xia, Zhihong 1992. "The Existence of Noncollision Singularities in Newtonian Systems." *Annals of Mathematics* 135: 411-468.