10

Humean Supervenience

BARRY LOEWER

Over the last couple of decades David Lewis has been elaborating and defending a metaphysical doctrine he calls "Humean Supervenience" (HS). Here is how he introduces it.

> Humean supervenience is named in honor of the great denier of necessary connections. It is the doctrine that all there is to the world is a vast mosaic of local matters of particular fact, just one little thing and then another. . . . We have geometry: a system of external relations of spatiotemporal distances between points. . . . And at those points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. For short: we have an arrangement of qualities. And that is all. There is no difference without difference in the arrangement of qualities. All else supervenes on that.[1]

In this paper I explore and to an extent defend HS. The main philosophical challenges to HS come from philosophical views that say that nomic concepts—*laws*, *chance*, and *causation*—denote features of the world that fail to supervene on non-nomic features. Lewis rejects these views and has labored mightily to construct HS accounts of nomic concepts. His account of laws is fundamental to his program, since his accounts of the other nomic notions rely on it. Recently, a number of philosophers have criticized Lewis's account, and Humean accounts of laws generally, for delivering, at best, a pale imitation of the genuine item.[2] These philosophers think that the notion of law needed by science requires laws—if there are any—to be fundamental features of our world that are completely distinct from and not supervenient on the particular facts that they explain. I side with Lewis against these philosophers. Here I will argue that although Lewis-laws don't fulfill *all* our philosophical expectations, they do play the roles

From *Philosophical Topics* 24 (1996): 101–27.

that science needs laws to play. The metaphysics and epistemology of Humean laws, and more specifically, Lewis-laws, are in much better shape than the metaphysics and epistemology of the main anti-Humean alternatives. However, I do have misgivings about Lewis's account. Both he and his critics assume that the basic properties are so individuated so that the laws are not metaphysically necessary. If this assumption is rejected, then the question of Humean supervenience lapses. I conclude with a brief discussion of this position.

I. Formulating and Fixing HS

Call a property "Humean" if its instantiation requires no more than a spatiotemporal point and its instantiation at that point has no *metaphysical* implications concerning the instantiations of fundamental properties elsewhere and elsewhen. Lewis's examples of Humean properties are the values of electromagnetic and gravitational fields and the presence or absence of a material particle at a point.[3] HS says that every contingent property instantiation at our world holds *in virtue* of the instantiation of Humean properties. If M is a contingent property, then an instantiation of M holds *in virtue of* instantiations of $P_1, P_2, \ldots, P_n$ only if in every metaphysically possible world at which the P instantiations hold the M instantiation also holds.[4] Lewis illustrates the *in virtue of* relation with the example of a grid of pixels. The grid exemplifies a particular picture, say, a depiction of a cube, in virtue of the firing of the grid's pixels. Note that HS doesn't require that if an individual instantiates a property F, it does so in virtue of Humean properties instantiated only at points where that individual is located. The instantiation of F may supervene on a larger pattern of Humean property instantiations and even on their totality.

Lewis says that HS is contingent. There are un-Humean possible worlds that contain facts that fail to supervene on the mosaic of Humean property instantiations at those worlds. At an un-Humean world, for example, consciousness might be instantiated by nothing smaller than a complex organism and the totality of Humean property instantiations at that world may not be metaphysically sufficient for its instantiation. At such a world, consciousness is an *emergent* un-Humean property. HS says that the actual world contains no properties like that.

Why believe Humean Supervenience? Hume can be interpreted as advocating HS but in a very different form and for very different reasons than

Lewis. For Hume the fundamental properties are kinds of impressions instantiated in the sensorium. All true judgments supervene on the distribution of these properties. So judgments that one impression or kind of impression is nomically connected with another either are strictly false or must be construed as supervening on the distribution of fundamental properties of impressions. If this interpretation of Hume is accurate, then his version of HS is not at all plausible. But Lewis's reasons for defending HS are not Hume's. He says that he defends HS "to resist philosophical arguments that there are more things in heaven and earth than physics has dreamt of."[5] In other words, his motivation is to support physicalism. He characterizes physicalism this way:

> [M]aterialist supervenience means that for anything mental there are physical conditions that would be sufficient for its presence, and physical conditions that would be sufficient for its absence.[6]

Physicalism says that whatever happens in our world happens in virtue of physical happenings. Lewis thinks that it is the job of physics to provide an inventory of the fundamental physical properties and optimistically suggests that "present day physics goes a long way toward a correct and complete inventory."[7] Among the properties he mentions are mass and charge; properties that he also takes to be Humean. Despite this guidance he doesn't tell us what makes a property a fundamental physical one, and without such an account physicalism is threatened with vacuity.[8] A proposal that I think captures what many have on their minds when they speak of fundamental physical properties is that they are the properties expressed by simple predicates of the true comprehensive fundamental physical theory. The true comprehensive fundamental physical theory is the minimal theory that accounts for changes of the locations and motions of macroscopic spatiotemporal entities and also for changes in the properties that account for locations and motions and so on.[9] If current physics is on the right track, then charge and mass may be *fundamental* physical properties but mental properties are not. Although mental properties are invoked to account for the motions of macroscopic entities (e.g., Smith's believing that his friend is across the street is invoked to account for his crossing the street), they are not expressed by predicates of the minimal comprehensive theory that can in principle account for the motions of macroscopic entities.[10] If physicalism is true, then mental properties and all other contin-

gent properties are instantiated in virtue of the instantiations of fundamental physical properties.[11]

HS and physicalism are distinct doctrines. HS doesn't entail physicalism since it is compatible with there being Humean properties that are not physical.[12] Physicalism doesn't entail HS since there is no guarantee that the fundamental properties posited by physics are intrinsic properties of spatiotemporal locations. In fact, it seems pretty clear that contemporary physics does dream of non-Humean properties. I have in mind so called "entangled states" that are responsible for quantum nonlocality, i.e., for quantum theory's violations of Bell inequalities.[13] The entangled state of a pair of particles fails to supervene on the intrinsic properties of the separate particles. That is, the local properties of each particle separately do not determine the full quantum state and, specifically, do not determine how the evolutions of the particles are linked.[14] Since we have reason to believe that quantum theory is true, we have reason to think that HS is false.

Lewis is aware of the objection. He initially responded by pointing out that quantum theory is not in very good philosophical shape.

> I am not ready to take lessons in ontology from quantum physics as it now is. First I must see how it looks when it is purified of instrumentalist frivolity . . . of doublethinking deviant logic . . . and—most of all—when it is purified of supernatural tales about the power of the observant mind to make things jump.[15]

However, there are versions of quantum mechanics—David Bohm's version for one—that are so purified.[16] Bohm's theory has a realist interpretation, conforms to standard logic, has no jumps at all, and doesn't figure the observant mind in its fundamental laws. The defender of HS cannot hide behind the hope that quantum theory is incoherent or merely an instrument for predicting experimental outcomes.

More recently, Lewis has accepted that HS needs to be reformulated to accommodate quantum nonlocality.[17] Here is a suggestion for how to do so in the context of Bohm's theory. The quantum state of an n-particle system is a field in $3n$ dimension configuration space where the value of the field at a point in configuration space is the amplitude of the quantum state at that point.[18] These field values can be thought of as intrinsic properties of points.[19] The ontology of Bohm's theory also includes a "world particle" whose location and motion in configuration space determines the loca-

tions and motions of ordinary material particles in three-dimensional space, and the locations and motions of these particles determine the manifest world.[20] If Bohm's theory is the correct and complete physical theory, and if physicalism is true, then everything would supervene on the quantum state and the location of the world particle. We can think of the manifest world—the world of macroscopic objects and their motions—as shadows cast by the quantum state and the world particle as they evolve in configuration space.[21]

The lesson for a defender of HS to take from quantum mechanics is to count a property as Humean in a world iff it is an intrinsic quality of points in the fundamental space of that world. If Bohm's theory (or any other version of quantum mechanics that construes the wave function realistically) is correct, then that space is configuration space.[22] Given this account of Humean properties, quantum nonlocality poses no threat to HS.

I am not sure whether my reformulation of HS takes care of all incompatibilities between HS and physics. But since Lewis's aim is to defend HS from philosophy and not from physics, let us turn to philosophical challenges to HS. The most important challenge is from philosophical views concerning nomic concepts; that is, from views about laws, causation, and chance.[23] Nomic features are not intrinsic to space-time points, so HS requires that they supervene on Humean properties. But according to the non-Humean tradition, they don't. The failure of supervenience is expressed in metaphors associated with laws and causation. Some advocates of the non-Humean tradition say that laws *govern* or *guide* the evolution of events and that causation *provides the cement* of the universe. What determines and cements *Es* can't supervene on *Es*. According to non-Humeans, the nomic facts, rather than being determined by the Humean facts, determine them! Since Lewis says that he defends HS to resist philosophical arguments that there is more than physics tells us, he must think that physics does not tell us that there are non-Humean laws or causation. But physics does not speak unambiguously. Certain regularities—e.g., Schrödinger's equation—are said to express laws, some laws posit probabilities, and physicists often claim that one event causes another (e.g., the absorption of a photon causes a change in the energy of an electron). The question is whether the "laws," "probability," and "causation" that physics speaks of are non-Humean or can be accommodated by HS. Of course, one way of defending HS is to deny that there are nomic facts. Perhaps they are projections of the mind, as Hume is reputed to have thought, or the inven-

tions of philosophical misinterpretations of science, as van Fraassen suggests.[24] But defending nomic nonfactualism would be a Herculean undertaking. Laws and chances obviously play important roles in the sciences, and many of our concepts, for example, functional concepts, have nomic commitments.[25] So if nomic concepts are not completely factual, then the thought that a certain functional concept is instantiated is also not completely factual (or is false). Defending HS by denying nomic facts is too costly. The other way of defending HS is to show that, appearances to the contrary, nomic facts do supervene on Humean facts. More specifically, the nomic notions employed by physics and the other sciences are compatible with HS. Of course, this approach must deflate the governing and cementing metaphors that are associated with nomic concepts. But that may not be too high a price to pay if the resulting notions can do the work required of them in the sciences.

II. Lewis's Accounts of the Nomic

Lewis defends HS by constructing reductive Humean accounts of laws, chances, and causation. I will mainly be concerned with his account of laws, but it will be useful to quickly sketch his accounts of all three kinds of nomic facts. Lewis accounts for laws and chances together by building on a suggestion of Ramsey's.

> Take all deductive systems whose theorems are true. Some are simpler, better systematized than others. Some are stronger, more informative than others. These virtues compete: An uninformative system can be very simple, an unsystematized compendium of miscellaneous information can be very informative. The best system is the one that strikes as good a balance as truth will allow between simplicity and strength. How good a balance that is will depend on how kind nature is. A regularity is a law iff it is a theorem of the best system.[26]

Chances enter the picture by letting deductive systems include sentences that specify the chances of events.

> Consider deductive systems that pertain not only to what happens in history, but also to what the chances are of various outcomes in various situations—for instance the decay probabilities for atoms of various isotopes. Require these systems to be true in what they say about history. . . . Require also that these systems aren't in the business of guessing the outcomes of

> what, by their own lights, are chance events; they never say that A without also saying that A never had any chance of not coming about.[27]

As Lewis says, axiom systems more or less fit the facts, are more or less strong, and are more or less simple. Strength is measured in terms of the informativeness of the implications of the axioms, fit in terms of the chance of the actual history, and simplicity in terms of syntactical and mathematical complexity and the number of independent assumptions. These features of deductive systems trade off. Strength and fit can often be improved at the cost of simplicity and vice versa. By assigning probabilities to types of events, systems sacrifice strength for fit, but they may also make great gains in simplicity. The best system is the one that gets the best balance of the three, while not both implying that q and that the chance that q is less than 1. The laws of a world are the generalizations that are entailed by the best system for that world. Among the laws may be regularities that mention chances; e.g., for any tritium atom the chance of its decaying in time interval t is x. The totality of chance laws entails what Lewis calls "history-to-chance conditionals"; i.e., statements of the form $H_t \rightarrow P_t(E) = x$. These specify the chances of future courses of events after t if the history up through t is H_t. The chance of E at t is derived from the history up to t and the history-to-chance conditionals. So as not to prejudice our discussion I will call the laws and chances delivered by this account the L-laws and L-chances. Of course, Lewis thinks that the L-laws and L-chances are the laws and chances.

A proposition is L-physically necessary just in case it is true in every world compatible with the L-laws. L-physical necessity thus defined is less than metaphysical necessity, more than mere actuality (not every truth is physically necessary), but thoroughly grounded in actuality. An interesting consequence of Lewis's account is that there are physically possible propositions that are incompatible with the laws being the laws and incompatible with their chances. We will return to this point later.

Lewis's account of causation is in terms of counterfactuals. The counterfactual $A \rightarrow B$ is true just in case there are worlds at which A and B are true that are more similar to the actual world than is any world at which A is true and B is not true. For the case in which the laws are deterministic (the indeterministic case is a bit more complicated), similarity is evaluated in terms of the extent to which worlds match the actual world in particular fact and the extent to which the spatiotemporal regions of those worlds are

compatible with the laws of the actual world. Similarity in these two respects trades off. Generally, perfect match can be improved at the expense of more extensive violation of law and vice versa. According to Lewis, his account makes the counterfactual "if Nixon had pushed the button, there would have been a nuclear holocaust" come out as true. There is a world whose events conform to the actual laws and match the events of the actual world up until time *t*, when events in a small spatiotemporal region (in Nixon's brain) violate the actual laws and lead, in conforming with the actual laws, to Nixon's pushing the button a few moments later and then to the nuclear holocaust. (Of course this assumes that the button is connected, the missiles prepared, and so forth. If these conditions were not present, then the counterfactual would be false.) Lewis thinks that this world is more similar to the actual world in match and conformity to the laws than is any world at which Nixon pushes the button and there is no nuclear holocaust. Match with the actual world after the button is pushed can be restored, but only by eradicating all traces of Nixon's button pushing. Lewis thinks that this would require widespread and big violations of the actual laws.[28]

To a first approximation, Lewis's account of L-causation is: Event *e* L-causally depends on event *c* just in case *c* and *e* are distinct occurring events and if *c* had not occurred, *e* would not have occurred (or the chance of *e* would have been smaller). Event *c* L-causes event *e* just in case there is a chain of events *c* . . . *e* related by causal dependence. Of course, Lewis claims that L-causation is causation.[29]

III. Some Clarifications

Lewis's reductions of laws, chance, and causation to Humean concepts are a philosophical tour de force. If correct, they show that the nomic features of the world are compatible with HS, and that goes a long way toward demystifying them. But are his reductions correct? Like any reductions they should be evaluated in terms of how well they ground and illuminate the practices involving the concepts. These practices are reflected in and are to an extent codified by our beliefs involving them. So we need to examine whether Lewis's reductions preserve our central and supportable nomic beliefs and how well they fit in with our other well-supported views. For example, it is generally believed that laws play a central role in explanations. If this is so, then it counts in favor of the reduction of laws to L-laws

if L-laws play that role, and it counts against the reduction if they don't. If L-laws (or any other nomic concepts) satisfy a sufficient number of our central and well-supported nomic beliefs and nothing else satisfies them equally well or better, then the reduction of laws to L-laws is successful. Exactly how many or which of our nomic beliefs must be respected is not clear cut. What one philosopher sees as a reduction, another may see as an elimination.[30] But if it can be shown that L-laws satisfy enough of our central beliefs concerning laws (and other nomic concepts) to play the roles that laws are supposed to play in the sciences and that nothing else plays these roles any better, then we will have good reasons to call L-laws "laws."

Since Lewis's accounts of chance, counterfactuals, and causation all involve laws, if the HS account of laws is not defensible, then even if the other accounts are correct, they would not establish that these nomic features are compatible with HS. For this reason, I will focus on Lewis's proposal that laws are L-laws.

There are some aspects of Lewis's account of laws that I want to clarify prior to seeing whether L-laws can play the role that laws are supposed to play.

Philosophers have understood "is a law" as applying to a number of different kinds of entities: sentences, propositions, or certain nonrepresentational features of reality, i.e., whatever it is that makes a particular sentence or proposition express a law. I will understand the L-laws as being propositions. They are the propositions expressed by the generalizations that are implied by the best axiom system.

"It is an L-law that p" is true at a world iff there is a unique best axiom system Φ for that world and among the theorems of Φ is a sentence that expresses the same proposition as "p." These truth conditions have some important consequences: First, "it is a law that p" implies "p." Second, "it is a law that" creates intensional contexts. So it may be a law that Fs are followed by Gs and it may be that F and F^* are coextensional, while it is not a law that F^*s are followed by Gs. Third, what makes a proposition an L-law at a world w is the "vast mosaic of local matters of particular fact" at w. No part of that reality that can be isolated makes a general proposition lawful or accidental.

Each of the notions "simple," "informativeness," and "best" needs clarification. Lewis thinks of simplicity as an objective property of expressions in a language (e.g., a conjunction is less simple than its conjunct) or of the proposition expressed by a sentence. Some mathematical propositions

are objectively simpler than others. The informativeness of a sentence is measured in terms of the number of possibilities it excludes. Lewis seems to think that the informativeness of a system is the informativeness of the conjunction of its axioms. I make a different suggestion below. Lewis doesn't say what "best" is, but it is reasonable to think of its content as being determined by scientific practice. He readily admits that all these notions are vague. But he thinks that it is not implausible that, given the way our world is, all the ways of clarifying them will count the same generalizations as laws.

There is a problem concerning the languages in which best systems are formulated. Simplicity, being partly syntactical, is sensitive to the language in which a theory is formulated, and so different choices of simple predicates can lead to different verdicts concerning simplicity. A language that contains "grue" and "bleen" as *simple* predicates but not "green" will count "All emeralds are green" as more complex than will a language that contains "green" as a simple predicate. More worrying: Let S be a system that entails all the truths at our world, and let F be a predicate that applies to all and only things at worlds where S holds. Then $(x)Fx$ is maximally strong and very simple. It is the best system for our world. The trouble is that it entails all true regularities, and so all regularities are L-laws.

Lewis's remedy is to insist that the simple predicates of the language in which systems are formulated (and in which their simplicity is evaluated) must express *natural* properties or universals. But which are the natural properties? One suggestion for picking out natural properties is not appropriate in the present context. It is that they are the properties that appear in the laws or that possess causal powers. This doesn't work, since it would make the analysis of laws and causation circular. Lewis's view seems to be that, since it does so much useful work, we should accept the notion of a natural property as a primitive.[31] He does say that it is plausible that the simple predicates of current physics are good candidates for expressing natural properties. But how does he know that? Perhaps Lewis's account should not be faulted for relying on the notion of a natural property since every other account of laws—both Humean and non-Humean—helps itself to a distinction between properties that are fit and those that are unfit for laws. But one worries that if the notion of a natural property is simply taken as a primitive, then we will have no epistemic access to which propositions are laws.[32] The problem isn't merely that all possible *evidence* may underdetermine which propositions are laws but that even if we know all the true

are objectively simpler than others. The informativeness of a sentence is measured in terms of the number of possibilities it excludes. Lewis seems to think that the informativeness of a system is the informativeness of the conjunction of its axioms. I make a different suggestion below. Lewis doesn't say what "best" is, but it is reasonable to think of its content as being determined by scientific practice. He readily admits that all these notions are vague. But he thinks that it is not implausible that, given the way our world is, all the ways of clarifying them will count the same generalizations as laws.

There is a problem concerning the languages in which best systems are formulated. Simplicity, being partly syntactical, is sensitive to the language in which a theory is formulated, and so different choices of simple predicates can lead to different verdicts concerning simplicity. A language that contains "grue" and "bleen" as *simple* predicates but not "green" will count "All emeralds are green" as more complex than will a language that contains "green" as a simple predicate. More worrying: Let S be a system that entails all the truths at our world, and let F be a predicate that applies to all and only things at worlds where S holds. Then $(x)Fx$ is maximally strong and very simple. It is the best system for our world. The trouble is that it entails all true regularities, and so all regularities are L-laws.

Lewis's remedy is to insist that the simple predicates of the language in which systems are formulated (and in which their simplicity is evaluated) must express *natural* properties or universals. But which are the natural properties? One suggestion for picking out natural properties is not appropriate in the present context. It is that they are the properties that appear in the laws or that possess causal powers. This doesn't work, since it would make the analysis of laws and causation circular. Lewis's view seems to be that, since it does so much useful work, we should accept the notion of a natural property as a primitive.[31] He does say that it is plausible that the simple predicates of current physics are good candidates for expressing natural properties. But how does he know that? Perhaps Lewis's account should not be faulted for relying on the notion of a natural property since every other account of laws—both Humean and non-Humean—helps itself to a distinction between properties that are fit and those that are unfit for laws. But one worries that if the notion of a natural property is simply taken as a primitive, then we will have no epistemic access to which propositions are laws.[32] The problem isn't merely that all possible *evidence* may underdetermine which propositions are laws but that even if we know all the true

sentences (except sentences that say which are the natural properties) of every possible language, we still don't know which express laws until we know which predicates express natural properties.

Here is a different suggestion for specifying the language in which the axiom systems are formulated that doesn't rely on the notion of a natural property. I assume that it is the job of physics to account for the positions and motions of paradigm physical objects (planets, projectiles, particles, etc.). This being so, the proposal is that we measure the informativeness of an axiom system so that a premium is put on its informativeness concerning the positions and motions of paradigm physical objects. And further, we measure the informativeness of a system not in terms of its content (i.e., set of possible worlds excluded) but in terms of the number and variety of its theorems. Systems have infinitely many theorems, so we just can't compare systems by counting theorems. One way to deal with this difficulty is to discount the contribution of a theorem to the informativeness of the system which implies it by the length of its proof in some regimented proof system. So in evaluating the "informativeness" of a system, we enumerate its proofs by their length, award points for the informativeness of a theorem, with extra points awarded if it is about the motions of ordinary objects, and then divide by the length of the proof.

If the above account of informativeness can be worked out, then it will immediately take care of the trivialization problem. The system $(x)Fx$ would not be counted as "informative" since, although its theorem $(x)Fx$ is very informative, it has no theorems that mention the positions and motions of ordinary objects. The other worry was that systems formulated in "gruesome" languages may vie with systems formulated in our language for simplicity and strength but may entail different generalizations. But if the systems agree with respect to the number and variety of theorems which mention positions and motions, etc., then we have no reason to believe that this will be the case, and we have some reason to believe that it won't be the case. It seems likely that the gruesome system will have to be a bit more complicated to equal an ungruesome system in informativeness. And if there are gruesome and ungruesome systems that agree in both simplicity and informativeness, they still may imply exactly the same generalizations. If this is right, then we can dispense with natural properties. But there is still an oddity. If the best system formulated in our language entails that "all emeralds are green" and "all rubies are red," then the best system formulated in a language containing the simple predicates "gred" and

"emerubies" will entail that "all emerubies are gred." But maybe that's not so bad since this generalization is nomologically necessary. Perhaps our being disinclined to count it as a law just reflects the bias of the language which we actually use.

IV. Are the L-Laws the Laws?

I now want to examine to what extent L-laws satisfy our central beliefs about laws. Here is a list of the most important features that laws are supposed to have:[33]

(*i*) If it's a law that *F*s are followed by *G*s, then it is true that *F*s are followed by *G*s.

(*ii*) Being a law is a mind-independent property.

(*iii*) The laws are important features of our world worth knowing.

(*iv*) It is a goal of scientific theorizing to discover laws, and we have reason to believe that some of the propositions that the fundamental sciences classify as laws are laws.

(*v*) There is a distinction between lawful generalizations and accidental generalizations.

(*vi*) There are vacuous laws.

(*vii*) Laws are contingent but ground necessities.

(*viii*) Laws support counterfactuals.

(*ix*) Laws explain.

(*x*) Laws are confirmed by their instances.

(*xi*) The success of induction depends on the existence of laws.

(*xii*) The laws govern (direct, constrain, or probabilistically guide) the evolution of events.

(*xiii*) If it is a law that *p*, and *q* is any proposition expressing boundary conditions or initial conditions relevant to the law that are co-possible with *p*, then it is possible that it is a law that *p* and *q*.

Some of these conditions come from scientific practice and others from philosophical reflection (not confined to philosophers). Some are more important than others. Any alleged account of laws that failed to ground a distinction between lawful and accidental regularities is obviously mistaken. On the other hand, an account of laws that didn't endorse the metaphor that laws govern events shouldn't be rejected on that account. The metaphor is obscure and not obviously connected with actual scientific practice.

L-laws clearly satisfy (*i*), (*v*), and (*vi*). With respect to (*v*) and (*vi*), L-laws are a big improvement on traditional regularity accounts. According to regularity accounts, a proposition is a law iff it is expressed by a true generalization whose predicates are nonpositional and projectible.[34] Vacuous generalizations are true, so all vacuous regularities composed of projectible predicates are counted as laws by the regularity account. This can be avoided by requiring that laws have instances, but that would exclude all vacuous generalizations some of which seem to be laws; e.g., the ideal gas.

Reichenbach gives the following example to illustrate the distinction between lawful and accidental generalizations.

(U) There are no solid one ton spheres of uranium.
(G) There are no solid one ton spheres of gold.

Reichenbach observes that (U) is a law but that (G) isn't.[35] Both of these generalizations are true and contain only nonpositional and projectible predicates, so the regularity theory can't distinguish them. But Lewis's account can. It is plausible that quantum theory together with propositions describing the nature of uranium entail (U) but not (G). So if quantum theory is part of the best theory of our world, then (U) will be a law. In fact, the reason we think that (G) is not a law is that we think that the best theory of our world is compatible with (G)'s being false. Adding (G) to fundamental physical theory would produce a stronger system but at a great cost in simplicity.

L-laws also seem to satisfy (*iii*) and (*iv*). If one knows the L-laws, then one would know a lot about the world and have that knowledge in the form of simple compact axioms. Further, it is not implausible that, at least in physics, the goal of theory construction is to find true, strong, well-fitting, and simple theories. The fundamental theories of physics—e.g., quantum theory, general relativity—exhibit these virtues. Propositions that scientists call "laws" are consequences of the fundamental theories (e.g., Schrödinger's law) or of these laws together with sentences connecting higher-level descriptions with quantum mechanical descriptions (e.g., laws of chemical bonding). If we were to learn that a certain system was best for our world, we would have reason to believe that its general consequences are laws.

Whether or not L-laws satisfy (*vii*), (*viii*), and (*ix*) is controversial. L-laws are related to the other L-nomic concepts in more or less the way endorsed by philosophical tradition. L-laws are contingent and the regularities they

entail are L-necessary; L-laws can be premises in deductive arguments that have the form of deductive nomological explanations; and it is generally the case that if it is an L-law that *F*s are followed by *G*s, then the counterfactual "if an *F* occurred, it would be followed by a *G*" will generally be true.[36] Of course, to anti-Humeans, L-laws are sham laws that are capable only of supporting sham counterfactuals, etc.[37] But these complaints should not be taken seriously unless backed up by arguments that show that L-counterfactuals, L-necessity, and L-explanation are not the genuine items. If, for example, genuine counterfactuals do not supervene on Humean facts, then they can't be supported by L-laws. But, although specifics of Lewis's account have been criticized, I know of no argument that shows that the counterfactuals laws are supposed to support express non-HS facts.

Armstrong does argue that Humean regularities cannot really explain.

> Suppose, however, that laws are mere regularities. We are then trying to explain the fact that all observed *F*s are *G*s by appealing to the hypothesis that all *F*s are *G*s. Could this hypothesis serve as an explanation? It does not seem that it could. That all *F*s are *G*s is a complex state of affairs which is in part *constituted* by the fact that all observed *F*s are *G*s. "All *F*s are *G*s" can even be rewritten as "All observed *F*s are *G*s and all unobserved *F*s are *G*s." As a result, trying to explain why all observed *F*s are *G*s by postulating that all *F*s are *G*s is a case of trying to explain something by appealing to a state of affairs part of which is the thing to be explained.[38]

It is likely that he would similarly complain that L-laws don't really explain since the fact that a regularity is an L-law is a complex state of affairs constituted in part by the regularity. But the argument isn't any good. If laws explain by logically implying an explanandum—as the DN model claims—then the state of affairs expressed by the law will in part be constituted by the state of affairs expressed by the explanandum. How else could the logical implication obtain? In any case, L-laws do explain. They explain by unifying. To say that a regularity is an L-law is to say that it *can* be derived from the best system of the world. But this entails that it can be unified by connecting it to the other regularities implied by the best system. I suspect that Armstrong thinks that L-laws don't explain because he thinks that laws explain in some way other than by unifying. I will return to this point later when we discuss his own view of laws.

Can L-laws play the roles in induction that laws are supposed to play? One of these roles is that laws are confirmed by their instances. Let's say

that a generalization "*F*s are followed by *G*s" is confirmed by its instances iff an instance of the generalization increases its credibility and also the credibility that unexamined *F*s are followed by *G*s. Dretske suggests that if laws are mere Humean uniformities, then they are not confirmed by their instances.[39] He seems to think that all Humean uniformities are like "all the coins in Smith's pocket are dimes," in that one instance lends no credibility to another. But, of course, there is a difference between a uniformity which is an L-law and one which is accidental. The question is whether this difference permits confirmation of the former but not the latter. On a Bayesian account of confirmation, the answer is affirmative. There are probability distributions on which Newton's gravitational law (construed as an L-law) is confirmed by its instances but "all the coins in Smith's pocket are dimes" is not confirmed by its instances. Perhaps Dretske thinks that it should follow from the nature of laws that they are confirmed by their instances. It is true that this doesn't follow on the Bayesian account of confirmation. There are probability distributions on which gruesome generalizations rather than L-laws are confirmable. But I don't consider this to be a very strong objection to L-laws since I don't see how any plausible account of laws can *guarantee* that they are confirmed by their instances.[40]

Armstrong claims that "if laws of nature are nothing but Humean uniformities, then inductive scepticism is inevitable."[41] His argument is that if the laws were Humean uniformities, then we could not explain why induction is rational (or necessarily rational), and without such an explanation inductive skepticism follows. I don't want to examine his argument in detail.[42] Suffice it to say that it depends on his claim that non-Humean laws can explain their instances while Humean uniformities cannot. We have already seen that this assumption is question begging.

If "inductive skepticism" means that it is impossible to provide a non-question-begging justification of a system of inductive inference, then I agree with Armstrong's claim that Humeanism makes inductive skepticism inevitable. That is because it is *inevitable period,* whatever laws may be. Hume conclusively showed the impossibility of a non-question-begging justification of any universal system of inductive inference. But if Armstrong means that someone who believed that laws are Humean uniformities (or that there are no non-Humean laws) is irrational in making inductive inferences, then Armstrong is pretty clearly wrong. Suppose that *D* is a scientist who assigns a probability of 1 to HS and also allocates substantial initial probability to simple and strong theories, including the true one. As

she accumulates evidence she will probably (relative to her probability assignment) come to assign a high probability to the true system and to the L-laws.[43] That her decisions are based on assigning high probabilities to many true propositions is likely to make those decisions successful. It is hard to see what reason we could have for thinking that *D* is irrational.

One of the conditions on our list that L-laws seem to violate is mind independence [i.e., condition (*ii*)]. The property of being an L-law is defined in terms of standards of simplicity, strength, fit, and best combination. These standards seem to be relative to us; i.e., to our psychology and interests. My proposal for making informativeness concerning position and motion especially important also may seem to make lawhood relative to our interests. We can imagine cognitive beings whose standards and interests differ greatly from ours. So it is apparently a consequence of Lewis's account that which propositions are laws depends on mental facts about us. This smells, at least a little, like nomic idealism.

But it is not clear that being an L-law is mind dependent in any way that is troubling for the jobs that laws are required to perform in science. First, it should be noted that the mind independence of the lawful regularities themselves is completely compatible with Lewis's account. What is at issue is whether the lawfulness of those regularities is mind dependent.[44] Second, Lewis's account is compatible with the view that scientists are now mistaken concerning which generalizations are L-laws and even with the view that in the Peircian ideal scientists will be mistaken.[45] So being an L-law is compatible with fairly robust realism. Third, the extension of "is a law" at a world *w* is determined not by the standards of simplicity, etc., of the scientists (if there are any) at *w* but by the scientists at our world. This rigidifies the standards and so falsifies the counterfactual that had our standards been different so would have the laws. Fourth, Lewis observes that simplicity, strength, fit, and balance are only partly relative to us. Independently of our psychology or opinions, a linear function is simpler than a quartic function, a second-order differential equation is simpler than a third-order one, etc. So he suggests that the mosaic of Humean facts of our world may be such that the best system is robustly the best. Varying the subjective aspects of simplicity, etc., within the space left by objective criteria may leave the best system unaltered. The upshot is that although the property of being an L-law is partly constituted by psychological factors, which generalizations are the laws is mind independent. So far as I can see, the fact that the concept of laws is partly constituted by concepts in-

volving scientists' standards does not prevent them from explaining, supporting counterfactuals, etc.

Scientists and others often talk of laws *governing* or *guiding* events; i.e., they invoke condition (*xii*). The Laplacian creation myth embodies this way of thinking. God creates the universe by creating the laws and setting the initial conditions and then lets the history evolve under the direction of the laws. Physicists do something similar, at least in thought, when they take dynamical laws, set initial conditions, and then see what consequences ensue. But what do these metaphors of governing and guiding come to? No one thinks that the laws literally govern events.[46] Nor do the laws cause the events. But whatever these metaphors come to it is clear that L-laws don't govern the evolution of events. It is more apt to say that L-laws *summarize* events.

Condition (*xiii*) is closely connected to the idea that laws govern events. If dynamical laws govern events, then any initial conditions that are compatible with the generalizations entailed by the laws can be governed by the laws. It is not surprising, then, that L-laws don't satisfy (*xiii*). John Earman, who is sympathetic to HS, provides a simple example.[47] Consider a world w that contains only a single particle moving at a uniform velocity. The events of this world are compatible with Newton's laws, and it further seems possible that Newton's laws are the laws that obtain at w. But Newton's laws are not the L-laws at w since they are far more complicated and no more informative than the single generalization that all particles move at a uniform velocity.

The failure of L-laws to satisfy (*xiii*) is prima facie a serious matter. Given a set of dynamical laws, physicists consider the consequences of those laws for various initial conditions. No restriction is placed on these conditions other than that they be compatible with the generalizations expressed by the laws. L-laws can be used in this way, but there may be some initial conditions which, while consistent with the generalizations, are incompatible with their being laws.

The feeling that an adequate account of laws should satisfy (*xiii*) runs deep. Michael Tooley and John Carroll describe thought experiments which evoke intuitions based on (*xiii*) and use these thought experiments to argue against HS. Here is a variant of one of Carroll's examples: Consider worlds u and v as follows. Both u and v contain x particles and y particles and Newton's laws of motion obtain in both. The difference is that in u it is a law that when x and y particles interact they exchange the value of some

property—say spin—while in *v* it is a law that they don't exchange spins. The initial conditions of *u* and *v* make for many such interactions. These worlds differ in their Humean facts, so there is, so far, no problem for HS. But relative to each world it is possible—i.e., compatible with its laws—for there to have been initial conditions such that, had they obtained, there would have been no interactions between *x* and *y* particles. At such worlds, do the *u* law or the *v* law concerning *x* and *y* particles hold? We can conceive of both possibilities, so it seems that there are both kinds of worlds. At *u'* it's a law that *x* and *y* particles exchange spins when they interact, and at *v'* it's a law that they don't exchange spins. At *u'* but not at *v'* it's true that if an *x* and *y* particle were to interact, they would exchange spins. Since *u'* and *v'* are identical with respect to their Humean property instantiations, HS is false. Notice that (*xiii*) is invoked in the thought experiment when it is claimed that *u'* and *v'* are possibilities.

Carroll and Tooley seem to think that this kind of thought experiment is sufficient to conclusively refute HS accounts of laws. But the argument falls short of a refutation. The intuitions involved in the thought experiment are doubly suspicious. They involve possible situations that are enormously different from the actual world, and they involve scientific concepts. The assumption that such intuitions are accurate is, at best, questionable and in some cases has been outright discredited. For example, most people have the intuitions that continued application of force is required to keep a body in motion and that the heavier the object, the faster it falls. Obviously these intuitions are misguided. Why should intuitions concerning laws be more reliable?[48]

Pointing out that intuitions are not infallible is enough to show that the thought experiments aren't conclusive refutations. But, unless they can be explained away, they do count against Lewis's reduction. That is, they count against his reduction unless it can be explained why we have such intuitions even though laws fail to satisfy (*xiii*). Although any such explanation is speculative, there is a story that strikes me as plausible for how we could come to believe, mistakenly, that L-laws should satisfy (*xiii*). According to HS, nomic facts supervene on the totality of Humean facts. It will generally be the case that in regions of space-time that are small compared to the whole spatiotemporal region of the world, events that don't violate the laws also don't violate the fact that they are the laws. That is, the following condition can be satisfied by L-laws:

(*xiii*)* Given a set of laws {*L*} similar to the actual laws and any spatio-temporal region *S* that doesn't violate {*L*}, there is a Humean possible world *u* containing a region *S** that matches *S* and such that {*L*} is the set of the L-laws of *u*.

Physicists usually consider small systems whose initial conditions are compatible with what they take to be the lawful generalizations. Because the systems are small parts of the actual world, such systems will invariably also be compatible with these generalizations being L-laws. The practice of applying the laws to small systems (compared to the totality of facts) might lead to the belief that any system—no matter how large—whose initial conditions are compatible with the lawful generalizations is also compatible with these generalizations being laws; i.e., to (*xiii*), the condition that underlies the Carroll-Tooley intuitions. But if the laws are L-laws, then this belief is mistaken. Giving it up may be giving up something that we are used to but it wouldn't have much of an effect on scientific practice.

Let's take stock. L-laws clearly satisfy conditions (*i*), (*iii*) through (*vi*), (*x*), and (*xi*). They also satisfy (*vii*), (*viii*), and (*ix*), if the relevant nomic notions are construed as the corresponding L-nomic notions. It is arguable that L-laws satisfy (*ii*). The only conditions clearly violated by L-laws are (*xii*) and (*xiii*). Condition (*xiii*) is almost satisfied, and to the extent that it is not, we can explain why it's not though we think it should be. Condition (*xii*) is obscure. If there is nothing more to it than what is expressed by (*xiii*), then L-laws satisfy it to the extent they satisfy (*xiii*). If (*xiii*) requires something more, that more has not been expressed without metaphor and has not been shown to be anything required by science. Still, it will strike many philosophers that L-laws are eviscerated versions of laws. If there existed some entity that fully satisfied (*xii*) and (*xiii*) as well as all the other conditions on laws, then these philosophers would prefer to call these items "laws." Of course, these philosophers would thereby reject HS. If they could provide good reasons to believe that there are such robust laws, then they would provide good reasons to reject HS.

V. Non-Humean Accounts of Laws

Anti-Humeans think that Humean accounts at best deliver pale imitations of real laws. They say that real laws are distinct from the facts that they explain and don't supervene on them. I will call these hypothesized entities

"A-laws" after Armstrong who is, perhaps, their most prominent and persistent advocate.[49] A-laws are claimed to satisfy all of our conditions on laws including (*xii*) and (*xiii*). If this is so and there are A-laws, then Lewis's proposed reduction of laws to L-laws should, by his own lights, be rejected, since the A-laws, by satisfying more of our beliefs concerning laws, would better deserve the title "laws." And if there are A-laws, then HS is false, since, as we have seen, satisfaction of (*xii*) entails the failure of HS.

Let's say that "*F*s are followed by *G*s" is an A-law at a world *w* just in case the generalization "*F*s are followed by *G*s" instantiates the non-Humean property *X* at *w*. The property *X* is that property which makes "*F*s are followed by *G*s" an A-law. Armstrong, Dretske, and Tooley (ADT) all think that the property of being a law can be analyzed in terms of one property necessitating another.[50] Carroll and Maudlin propose views on which the concept of lawhood is primitive.[51] They say nothing about the property *X* that makes a generalization a law. It may be simple or complex. By offering an analysis of the law-making property, the ADT account sticks out its neck and is open to some objections that do not seem applicable to the primitivist account.[52] But the problems with A-laws that I will discuss apply to both approaches.

It is in virtue of the *X* property being non-Humean that A-laws satisfy (*xiii*), since as long as the mosaic of Humean property instantiations is logically compatible with a generalization, that generalization may satisfy *X;* i.e., it may be an A-law. It is satisfaction of *X* that empowers a generalization to explain its instances, support counterfactuals, direct or guide the evolution of events, and so forth. According to the anti-Humean, A-laws can support genuine counterfactuals; i.e., counterfactuals that don't supervene on the mosaic of Humean facts. Of course, L-counterfactuals can't do that.

There are worlds in which the A-laws and the L-laws coincide. But, of course, what makes a generalization an A-law is quite different from what makes it an L-law. There are also worlds in which the A-laws and L-laws are quite different. Two worlds can be exactly alike in their Humean facts (and therefore in their L-laws) but differ radically in their A-laws. There are worlds in which none of the generalizations entailed by the best axiom systems for those worlds are A-laws but in which other complicated and isolated generalizations are A-laws. Some worlds may have no L-laws since there are no best axiom systems for those worlds, but they may have many A-laws, etc.

Some Humeans think that the metaphysics of A-laws is incoherent. I partially agree. I don't think that there is a satisfactory way of cashing out

the idea that A-laws guide or direct the evolution of events. These metaphors are supposed to provide a way of understanding how it is that laws ground necessary connections, support counterfactuals, explain their instances, and so on. For example, if we think of a law as literally directing or guiding the course of events, then it may seem that the law together with initial conditions can account for the evolution of events. For laws to operate in this way there must be a law-making feature *M* distinct from the generalization that *F*s are followed by *G*s that brings it about that *F*s are followed by *G*s. What could this bringing about be? One suggestion is that *M causes F*s to be followed by *G*s. But this is unsatisfactory. Not only do we have no idea of what this *M* is and how it causes the regularity, but the suggestion seems to involve an infinite regress. The causal relation between *M* and the regularity is presumably backed by a law that brings it about that *M* causes *F*s to cause *G*s, etc. According to Armstrong, when "*F*s are followed by *G*s" is an A-law, then the universal *F* "brings along" the universal *G* and this bringing-along relation cannot be further explained (though it is a kind of causal relation). He says that "we must admit it in the spirit of natural piety."[53]

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 of the conditions on laws with the exception of (*xii*). However, there are still metaphysical puzzles about A-laws. 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 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.

So what are the reasons for believing that there are A-laws? One way of arguing for A-laws is to argue that there are laws and that L-laws (or other

Humean laws) can't do what laws are supposed to do; e.g., provide explanations, support counterfactuals, ground induction, etc. We have already discussed these arguments and found them to be question begging. There is another line of reasoning suggested by Armstrong to the effect that the existence of A-laws best explains certain regularities, and so, by inference to the best explanation, we have good reason to believe that there are A-laws.

> Laws, however, explain regularities. Even if we take the Humean uniformity itself, that all *F*s are *G*s, it seems to be an explanation of this uniformity that it is a law that *F*s are *G*s. But, given the Regularity theory, this would involve using the law to explain itself. We need to put some "distance" between the law and its manifestation if the law is to explain the manifestation.[54]

This suggests the following argument for A-laws: It is a regularity that *F*s are *G*s. There being an A-law that *F*s are *G*s explains this regularity better than its being an L-law that *F*s are *G*s. So it is reasonable to believe there is an A-law that *F*s are *G*s. There is much wrong with this argument. First, even if the existence of an A-law explained the regularity better than any competing explanation, it wouldn't follow that it is reasonable to believe that *F*s are *G*s is an A-law. At best that would make it prima facie reasonable to believe that it is an A-law. Countervailing reasons might make it unreasonable to believe that the A-law exists.[55] Second, it is not even clear that the fact that a regularity is an A-law is the best explanation of the regularity. As I previously pointed out, A-laws are simply declared to explain by postulation. In contrast, it is clear how L-laws explain. They explain by unifying. If it is an L-law that *F*s are *G*s, then the best system implies that *F*s are *G*s. Deriving *F*s are *G*s from the best system explains this regularity by unifying it.

Sometimes Armstrong suggests that our reasons for believing in A-laws are like our reasons for believing in theoretical entities. For example, we believe that electrons exist because their existence is a component of causal explanations of various phenomena; e.g., chemical bonding. It is reasonable to believe they exist because the theory that posits them unifies phenomena and provides causal explanations. But positing A-laws provides no such explanatory advantages. The hypothesis that there are A-laws which back certain regularities doesn't provide any additional unification. If anything, it is disunifying. Unlike electrons, A-laws, presumably, don't figure in causal explanations. Positing that certain regularities instantiate *X* as a theoretical explanation is doing science not philosophy. And it is doing science very badly, since it adds nothing to our scientific understanding. I

conclude that these arguments that A-laws provide better or even good explanations are ineffective.

John Carroll gives an argument that can be understood as an argument for the existence of A-laws.

> [I]n order for believing or reasoning to be instantiated at least some nomic concepts must also be instantiated. So, granting that the instantiation of any nomic concept requires there to be at least one law, for the error theorist or anyone else to believe any proposition at all, there must be at least one law. Thus, like anyone else, the error theorist cannot correctly believe that our universe is lawless.[56]

I agree that this argument establishes that if anyone believes that there are no laws, then that belief must be mistaken, since belief is a nomic property. But it would be a mistake to think that it establishes that believing that there are no A-laws is pragmatically inconsistent. That follows only if the laws that are required for beliefs are A-laws.[57] But as far as I can see, being a belief can be characterized in terms of L-nomic concepts.

Here is a third argument for A-laws.

> We have reason to believe that there are laws. Furthermore, we find ourselves believing or intuiting that laws satisfy all the conditions on the list. So we have reason to believe that there are laws that satisfy all the conditions on the list. But L-laws fail to satisfy conditions (*xii*) and (*xiii*), while A-laws satisfy these conditions. So we have reason to believe that there are A-laws.

This argument involves an inference from the fact that we have certain intuitions concerning a concept *C* to the conclusion that these intuitions are satisfied by *C*'s reference. There is a long tradition in philosophy of evoking intuitions that are associated with concepts in order to discover the nature of the concepts' reference. This method seems more appropriate for some concepts than for others. But, as was mentioned earlier when discussing conditions (*xii*) and (*xiii*), when *C*'s subject matter is scientific and when the intuitions concern modality, the argument is very weak and easily defeated by alternative explanations of why we have the intuitions we do.

The arguments just canvassed provide very little reason to believe that there are A-laws. Of course, it doesn't follow that there are no A-laws, but the epistemological position of the believer in A-laws is not very attractive. The anti-Humean claims that there is a property *X* that is instantiated by certain generalizations and that it is that property which makes those gen-

eralizations genuine laws and so capable of explaining their instances. But she has no account of how *X* accomplishes this. The Humean also thinks that there is a property—being entailed by the best system—that makes a generalization a law, and she does have an account of how that property makes the generalization explanatory. Chalk one up for the Humean. Further, a traditional epistemological principle—one which is part and parcel of scientific method—is that one should not believe that a certain kind of entity exists unless that entity is required by the best explanation of accepted facts. The only "evidence" that the anti-Humean can point to that would, without begging the question, count in favor of the existence of A-laws is our intuitions of nonsupervenience. That "evidence" is very weak and can be accounted for by the Humean. Since the Humean can account for all the evidence that the anti-Humean can account for and can do so without positing non-Humean properties or anything else that the anti-Humean doesn't already accept, the epistemological principle delivers a verdict in favor of the Humean view. I think that it may be this line of reasoning that Lewis has in mind when he says that he defends HS "to resist philosophical arguments that there is more in heaven and earth than physics has dreamt of." There is no scientific reason for believing in A-laws. Of course, physics tells us that there are laws (e.g., Schrödinger's law), but it doesn't tell us whether or not laws supervene on facts. The philosophical arguments that they don't supervene depend on taking intuitions about laws much more seriously than they deserve to be taken.

VI. Conclusion

It appears, then, that L-laws are pretty good candidates for laws and that on balance we don't have reason to think that there are any better competitors. They deserve the title "laws." Does this mean that HS has been saved from philosophy? Not by a long shot. There is still chance and causation to deal with. I will just register my opinion here that a good case (very similar to the case made for laws) can be made that L-chance can play the role of chance in science.[58] I am much less sanguine about the reduction of causation to L-causation. Lewis's account of causation is bedeviled by problems involving preemption and runs into difficulties when extended to indeterministic worlds. Of course, even if L-causation isn't causation, some other HS account may work. But if no HS account of causation is correct, the situation would be very dicey. On the one hand, causality seems to be so inter-

twined with so many of our concepts (indeed, with the concept *concept*) that if it fails to refer, then most of our thoughts would also fail to refer. On the other hand, look as hard as one might, we just don't find causal relations among the fundamental properties of physics or in the dynamical laws of physics. So we would be in a dilemma of either rejecting aspects of our conceptual scheme or rejecting physicalism, at least in its Humean formulation.

The other philosophical threat to HS, in my view a very serious one, involves an assumption that Lewis makes concerning the relation between natural properties and laws. Interestingly, it is an assumption also made by Armstrong. The assumption is that the fundamental laws are contingent. In other words, it is metaphysically possible for a property to be involved in a law in one world but not in another. This means that a fundamental property, e.g., gravitational mass, may conform to quite different laws, or no laws at all, in different possible worlds. At first, this assumption seems plausible since laws are knowable only a posteriori. But, on second thought, the assumption that properties are individuated independently of laws is quite perplexing. It would mean, for example, that the properties of gravitational mass and positive electromagnetic charge could, in another world, exchange places in the laws of that world, or that the property of gravitational mass appears in the law $F = m_1m_2/r^5$, etc. But this seems absurd. It amounts to supposing that fundamental properties possess a kind of haeccity that makes them the properties they are independently of the laws they figure in.

The alternative, necessitarian account of laws has been around for a while.[59] Some objections to it are easy to deflect. For example, even though laws may be metaphysically necessary it doesn't follow that they are a priori or that they are necessarily instantiated. A possibly more serious objection is that if some properties are dispositional, i.e., necessarily involve laws, then others must be categorical.[60] I don't want to evaluate the viability of this view here. But it is interesting to note that if the fundamental properties are individuated by the laws in which they figure, then the debate between Lewis and Armstrong cannot get off the ground, since the issue of whether nomic facts supervene on non-nomic facts requires that we can make a distinction between the two kinds of facts. Of course, if properties are nomically individuated, then HS is false, since the instantiation of a fundamental property has metaphysical implications for the instantia-

tions of properties elsewhere and elsewhen. The necessitarian account of laws is also at odds with the ADT account. If the fundamental properties are nomically individuated, then the laws are not, as they are on the ADT account, facts over and above occurrent events that govern or guide their evolution. Obviously, the issue of whether properties are nomically individuated needs to be settled before HS can be evaluated. But that is an issue for another paper.[61]

Notes

1. David Lewis, *Philosophical Papers,* vol. 2 (Oxford: Oxford University Press, 1986), ix.

2. E.g., Fred Dretske, "Laws of Nature," *Philosophy of Science* 44 (1977): 248–68; David Armstrong, *What Is a Law of Nature?* (Cambridge: Cambridge University Press, 1983); John Carroll, *Laws of Nature* (Cambridge: Cambridge University Press, 1994).

3. A property is intrinsic to a region R (or point x) if its instantiation at R doesn't metaphysically entail anything concerning contingent property instantiations at other regions. For example, being a planet is not intrinsic. Lewis's examples of Humean properties may be a bit surprising since it is natural to think that, e.g., electromagnetic field values by their very nature conform to certain laws and so have consequences for property instantiations elsewhere and elsewhen. Obviously, Lewis is not thinking of them like that but as categorical properties whose nomic commitments are contingent. I briefly discuss the plausibility of this view at the conclusion of the paper.

4. Perhaps more than this is required for one property to be instantiated in virtue of the instantiations of others. The relation expresses the idea that the first property instantiation is completely constituted by the other property instantiations. For a discussion of attempts to clarify *in virtue of,* see S. Webb, G. Witmer, and J. Yoo, "In Virtue Of" (manuscript, Rutgers University, 1996).

5. David Lewis, "Humean Supervenience Debugged," *Mind* 103 (1994): 473–89; the quotation appears on 474.

6. David Lewis, "Reduction of Mind," in Samuel Guttenplan, ed., A *Companion to the Philosophy of Mind* (Oxford: Basil Blackwell, 1994), 414.

7. Ibid., 412.

8. See the interchange between Tim Crane and D. H. Mellor ("There Is No Question of Physicalism," *Mind* 99 [1990]: 185–206) and Philip Petit ("A Definition of Physicalism," *Analysis* 53 [1993]: 213–33) for arguments that physicalism is vacuous and physicalist rejoinders.

9. Of course a fundamental physical theory—e.g., classical mechanics or quantum theory—does not by itself have any implications concerning the positions and motions of macroscopic objects. Propositions connecting the positions and motions of macroscopic objects with fundamental physical states are needed. In the case of classical mechanics, this connection is generally established by connections between macroscopic objects and the microparticles of which they are composed. The connections in quantum

theory are more complicated and controversial. (See David Albert and Barry Loewer, "Tails of Schrödinger's Cat," in Rob Clifton, ed., *Perspectives on Quantum Reality* [Boston: Kluwer, 1995].)

10. Current views might be wrong. It may be that to account for the motions of some macroscopic entities a hitherto unknown property or entity—perhaps some M-particle?—that exemplifies fundamental mental properties needs to be invoked. In that case we would say that M-particles are physical.

11. For arguments that mental properties supervene on physical properties, see Barry Loewer, "An Argument for Strong Supervenience," in Elias E. Savellos and Umit D. Yalcin, eds., *Supervenience: New Essays* (Cambridge: Cambridge University Press, 1995) and David Papineau, "Arguments for Supervenience and Physical Realization," in *Supervenience: New Essays*. Even if the physical facts metaphysically determine the mental facts (i.e., physicalism is true), it may be that we cannot epistemically determine the relation between the two. For a recent survey of attempts to explain intentional mental properties in terms of non-intentional properties, see Barry Loewer, "Naturalizing Semantics," in Crispin Wright and Robert Hale, eds., *Companion to the Philosophy of Language* (Cambridge: Cambridge University Press, 1997).

12. Nonphysical Humean properties would be epiphenomenal with respect to physical properties.

13. See Tim Maudlin, *Quantum Nonlocality and Relativity* (Oxford: Basil Blackwell, 1994) for in-depth discussion of the Bell inequalities and quantum mechanics.

14. An example of an entangled state is the EPRB state $1/2$ ($||>||> + ||>||>$). In this state neither electron 1 nor electron 2 possesses a well-defined spin, but the state also entails that the probability of a measurement of any component of spin yielding an "up" result is $1/2$. It also entails that if the spin component in the x direction of one electron is measured and yields "up" then a measurement of the same component of spin on the other electron will certainly yield "down."

15. Lewis, *Philosophical Papers,* vol. 2, xi.

16. In addition to Bohm's theory, two other noncollapse versions of quantum theory, the modal interpretations, and the many minds version of the many worlds interpretation are purged of the defects Lewis mentions (although each has its own peculiarities). There are also collapse versions—the GRW theory being the most promising—that are purged of the features that Lewis rightly finds unacceptable. For a survey of most of these versions, see David Albert, *Quantum Mechanics and Experience* (Cambridge, Mass.: Harvard University Press, 1992).

17. See Lewis, "Humean Supervenience Debugged."

18. This neglects spin. Particles with spin add to the dimensionality of configuration space.

19. Rendering the wave function fully Humean involves a couple of maneuvers. One is that the fact that the sum of the amplitudes of disjoint regions is less than or equal to 1 cannot be understood as following from the nature of the quantum state, since that would mean that the values of the field at some points have implications for its values at other points. Instead, this constraint has to be construed as an initial condition or law. The Schrödinger dynamics entails that if it is satisfied at one time, it is satisfied at all

times. A second problem is that the exact form of the Schrödinger equation depends on the Hamiltonian of the universe. If the Hamiltonian is understood as a property, it is not a Humean property, since it is instantiated by nothing smaller than the whole universe and doesn't supervene on Humean properties. To overcome this, it must be built into the Schrödinger law. In other words, Schrödinger's law formulated with the Hamiltonian of our universe is the fundamental dynamical law governing the evolution of the quantum field in configuration space.

20. The velocity of the world particle depends on the value of the quantum field at the point the particle occupies in accordance with a deterministic law. (See David Bohm and B. J. Hiley, *The Undivided Universe* [New York: Routledge, 1993]; Albert, op. cit.; Loewer, "An Argument for Strong Supervenience.")

21. See David Albert, "Elementary Quantum Metaphysics," in James T. Cushing, Arthur Fine, Sheldon Goldstein, eds., *Bohmian Mechanics and Quantum Theory: An Appraisal* (Dordrecht: Kluwer, 1996) for a discussion of the claim that configuration space is the fundamental space of quantum theory.

22. This version of Bohm's theory certainly isn't true. An adequate version would be one that is compatible with quantum field theory and gravitational theory. Such a theory has yet to be created.

23. Other threats come from mental properties, especially consciousness properties, and from persisting entities both physical and mental. If consciousness properties are not compatible with physicalism, then unless they were instantiated at points they would not be compatible with HS either. Whether or not instances of an *F* persist as the same *F* through time is not, at least for some *F*s (e.g., persons), obviously supervenient on Humean properties. Lewis's approach to persistence through time is to reduce it to causal relations among temporal parts.

24. Nomic nonfactualism takes two familiar forms. Noncognitivism about laws says that saying that a generalization is a law is to express an epistemic attitude toward it. Eliminativism about laws says that our concept of laws fails to refer to anything real. Blackburn defends the position that nomic concepts involve projections of our attitudes (see Simon Blackburn, *Essays in Quasi-Realism* [New York: Oxford University Press, 1993]). Bruno de Finetti advocates a noncognitivist account of chance in his "Foresight: Its Logical Laws, Its Subjective Sources," in Henry E. Kyburg, Jr., and Howard E. Smokler, eds., *Studies in Subjective Probability* (New York: Wiley, 1964). Van Fraassen argues that the place of laws within science has been greatly overestimated by philosophy (see Bas C. van Fraassen, *Laws and Symmetries* [Oxford: Oxford University Press, 1989]).

25. Carroll (op. cit.) emphasizes that nomic elimitivism is implausible for these reasons.

26. Lewis, "Humean Supervenience Debugged," 478.

27. Ibid., 480.

28. If the laws are Newton's laws or Bohm's laws (both of which are deterministic), then Lewis's account is in trouble. The problem is that because these laws are time reversible, there is a world that differs from the actual world in its history, a world in which Nixon pushes the button but a small violation of the laws of the actual world (no bigger than the violation incurred by a world that matches the actual world up until a short

time before Nixon pushes the button) leads to a world that matches the actual world from a short time after Nixon pushes the button. Because of this, "if Nixon had pushed the button, there would have been a nuclear holocaust" comes out as false on Lewis's account. One remedy is to count past match as more important than future match.

29. Lewis introduces further complications to handle probabilistic causation, preemption, and overdetermination (see Lewis, *Philosophical Papers,* vol. 2).

30. Stephen Stich, in *Deconstructing the Mind* (Cambridge, Mass.: MIT Press, 1996), discusses the difficulty of drawing a line between reduction and elimination.

31. Lewis discusses the possibility of defining natural properties as ones whose sharing makes for objective resemblance, in "New Work for the Theory of Universals," *Australasian Journal of Philosophy* 59 (1983): 5–30. But this isn't very enlightening without an account of objective resemblance.

32. Van Fraassen (op. cit.) develops this point as a criticism of Lewis's account.

33. These features are presented in van Fraassen, op. cit.

34. For sources of traditional regularity accounts, see Ernest Nagel, *The Structure of Science* (New York: Harcourt, Brace, and World, 1961); Carl Hempel, *Aspects of Scientific Explanation* (New York: The Free Press, 1965); and Nelson Goodman, *Fact, Fiction, and Forecast* (Cambridge, Mass.: Harvard University Press, 1983).

35. Cited in van Fraassen, op. cit., 27.

36. The connection between L-laws and L-counterfactuals is a bit complicated. On Lewis's account of counterfactuals, it could turn out that the most similar world or worlds in which an *F* occurs are ones in which "*F*s are followed by *G*s" is not a law and a *G* would not follow the *F*.

37. See Dretske, op. cit.; and Armstrong, op. cit.

38. Armstrong, op. cit., 40.

39. Dretske, op. cit., 258.

40. Non-Humean accounts of laws have no way of guaranteeing that scientists possess probability distributions that permit the confirmation of laws either. Of course, they could hold that confirmation is an objective notion and that scientists should have probability distributions on which laws are confirmed. But that is something Lewis could say as well. The only account of laws I know of that makes their confirmation by instances an essential feature of them is Goodman's Humean account (see Goodman op. cit.). But it fails to satisfy most of the other conditions on our list.

41. Armstrong, op. cit., 52.

42. For an examination of Armstrong's argument, see van Fraassen, op. cit., 128.

43. If, on her probability distribution, theories are underdetermined by evidence, then the best she could do (even by her own lights) is to end up allocating her probabilities among observationally equivalent theories; see John Barman, *Bayes or Bust* (Cambridge, Mass.: MIT Press, 1992).

44. Lewis jokes that "if we don't like the misfortunes of nature that the laws visit upon us, we can change the way we think!" ("Humean Supervenience Debugged," 479). Of course, at most we could change whether the misfortunes were the result of laws or were mere accidents. Changing the standards won't alleviate the misfortunes themselves.

45. Suppose that there are distinct fundamental systems that have the same observa-

tional consequences but differ with respect to some generalizations concerning unobservables. At most one is true, but even if we had all the observational evidence and followed the scientific method, it doesn't follow that we could know which is true. Hence in this situation we wouldn't know all the laws.

46. But Fred Dretske comes close.

> Consider the complex set of legal relationships defining the authority, responsibilities, and powers of the three branches of government. . . . The legal code lays down a set of relationships between the various *offices* of government, and this set of relationships imposes legal constraints. . . . Natural laws may be thought of as a set of relationships that exist between the various "offices" that objects sometimes occupy (Dretske, op. cit., 264).

47. Earman, op. cit.

48. The view that our intuitions involving a concept must be satisfied by the concept's reference may rely on a certain view of concepts and intuitions. Specifically, it may rely on the view that concepts are analytically tied to certain descriptions, that anyone who grasps a concept knows what these descriptions are, and, further, that robust intuitions involving the concept provide access to these descriptions. This view of concepts has not fared especially well in recent discussions (see Stich, op. cit.) and is, in any case, a shaky foundation on which to base a rejection of HS.

49. See Armstrong, op. cit.

50. These authors differ somewhat about the nature of the necessitation relation. Armstrong thinks of it as a causal relation.

51. See Carroll, op. cit., and Tim Maudlin, "Laws: A Modest Proposal" (manuscript, 1994).

52. For such objections, see van Fraassen, op. cit.

53. Armstrong, op. cit, 92.

54. Ibid., 41.

55. Van Fraassen (op. cit.) makes this point against "inference to the best explanation."

56. Carroll, op. cit., 91.

57. Carroll concludes from this argument that there are laws and then argues on the basis of thought experiments that laws are A-laws.

58. Lewis once thought that this was not so. (See his *Philosophical Papers,* vol. 2.) He called chance the "one big bad bug" in his HS program. The problem with L-chance is its apparent incompatibility with what he took to be the definitive principle connecting chance and belief, i.e., "the Principle Principle":

$$\text{(PP)}\ P_t(A/\text{ch}_t(A) = x \,\&\, E_t) = x,$$

where P_t is one's degree-of-belief function at time t, $\text{ch}_t(A)$ is the chance of A at t, and E_t is any proposition that is admissible at t. (PP) supplies the likelihoods in Bayes's theorem. Lewis pointed out that (PP) is incompatible with his account of chance. The trouble is this: Suppose that h is the history of the actual world up through time t, f is the actual future after t, and f^* is a non-actual future after t. The best system for $h \,\&\, f$ may endorse the history-to-chance conditional $h \rightarrow \text{ch}(f^*) = x$, even though on Lewis's account of chance

the world $h \,\&\, f^*$ is incompatible with $\text{ch}(f^*) = x$. That is, the best system for $h \,\&\, f^*$ and h entails that $\text{ch}(f^*) = x$ is false. This being so, the rational degree of belief to assign to f^* given that the L-chance of f is x is 0; not x, as (PP) counsels.

Lewis recently responded (in "Humean Supervenience Debugged") to his own objection to HS by adopting a suggestion due to Michael Thau to replace (PP) by

(NP) $P(A/\text{ch}_t(A) = x \,\&\, E) = \text{ch}_t(A/B)$,

where B is the proposition that $\text{ch}_t(A) = x$ [i.e., the set of possible worlds at which $\text{ch}_t(A) = x$]. (See Michael Thau, "Undermining and Admissibility," *Mind* 103 [1994]: 495–503.) For most propositions A, the chances of A and B at t are independent, so (NP) reduces to (PP). That solves the problem.

59. See, e.g., Wilfrid Sellars, "Concepts As Involving Laws and Inconceivable without Them," *Philosophy of Science* 15 (1948): 287–315; Sydney Shoemaker, "Causality and Properties," in Peter van Inwagen, ed., *Time and Cause* (Dordrecht: D. Reidel, 1980); Chris Swoyer, "The Nature of Natural Laws," *Australasian Journal of Philosophy* 60 (1982): 203–23.

60. See Blackburn, op. cit.

61. I am grateful to David Albert, John Carroll, Brian Loar, Tim Maudlin, and Scott Sturgeon for comments on earlier versions of this paper.