



Humean Supervenience Debugged

Author(s): David Lewis

Source: *Mind*, New Series, Vol. 103, No. 412, (Oct., 1994), pp. 473-490

Published by: Oxford University Press on behalf of the Mind Association

Stable URL: <http://www.jstor.org/stable/2254396>

Accessed: 06/06/2008 11:57

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=oup>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We enable the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.

SYMPOSIUM: CHANCE AND CREDENCE

Humean Supervenience Debugged

DAVID LEWIS

Years ago, I wrote that much of my work could be seen in hindsight as a campaign on behalf of “Humean Supervenience”: the thesis that the whole truth about a world like ours supervenes on the spatiotemporal distribution of local qualities (1986, pp. ix-xvi). I thought this campaign had been mostly successful. Despite some unfinished business with causation, especially the problem presented in Menzies (1989), I think so still. But I wrote that “There is one big bad bug: chance. It is here, and here alone, that I fear defeat” (1986, p. xiv). I think I can say at last how to beat the bug. But first I’ll have to take a lot of time reviewing old ground. I’ll reintroduce Humean supervenience, with some afterthoughts. I’ll say what a Humean analysis of chance might look like. I’ll say why Humean analyses of chance are in bad trouble, and why unHumean analyses are not an acceptable refuge. I’ll give the beginning of a solution to the problem that plagues Humean analyses, and I’ll say why that beginning is not good enough. And then I’ll come at last to the good news: thanks to a suggestion by Michael Thau, I think I know how to complete the solution. The resulting rescue of Humean chance won’t give us all we might wish, but I think it gives us enough.

So the key idea in this paper is Thau’s. I thank him for kindly permitting me to use it here in my own way. But it can be used in other ways as well. Thau himself has not joined my campaign on behalf of Humean supervenience. But there are other theses about chance, weaker and less contentious than Humean supervenience itself, that are bitten by their own versions of the big bad bug. As Thau (1994) explains in his companion paper, his idea can be used also in defence of those weaker theses.

1. Humean Supervenience

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 (1988, pp. 132-3 and 158-9). If two possible worlds are discernible in any way at all, it must be because they differ in what things there 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.

From this starting point, we can go on to add various further theses about the basis on which all else supervenes. As an anti-haecceitist, I myself would drop the “what things there are” clause; I claim that all contingent truth supervenes just on the pattern of coinstantiation, never mind which particular hooks the properties and relations are hanging on. (On my view, the hooks are never identical from one world to another, but that by itself doesn’t make the worlds discernible.) Some will wish to add that the fundamental properties and relations are Armstrong’s immanent universals; or maybe that they are Williams’s families of exactly resembling tropes. And we may reasonably hope that physics—present-day physics, or anyway some not-too-distant improvement thereof—will give us the inventory of all the perfectly natural properties and relations that ever appear in this world.

Humean Supervenience is yet another speculative addition to the thesis that truth supervenes on being. It says that in a world like ours, the fundamental relations are exactly the spatiotemporal relations: distance relations, both spacelike and timelike, and perhaps also occupancy relations between point-sized things and spacetime points. And it says that in a world like ours, the fundamental properties are local qualities: perfectly natural intrinsic properties of points, or of point-sized occupants of points. Therefore it says that all else supervenes on the spatiotemporal arrangement of local qualities throughout all of history, past and present and future.

The picture is inspired by classical physics. Humean Supervenience doesn’t actually say that physics is right about what local qualities there are, but that’s the case to keep in mind. But if we keep physics in mind, we’d better remember that physics isn’t really classical. For instance, a rival picture inspired by waves in state-space might say that many fundamental properties are instantiated not at points but at point-tuples (Forrest 1988, p. 155). The point of defending Humean Supervenience is not to support reactionary physics, but rather to resist philosophical arguments that there are more things in heaven and earth than physics has dreamt of. Therefore if I defend the *philosophical* tenability of Humean Supervenience, that defence can doubtless be adapted to whatever better supervenience thesis may emerge from better physics.

Even classical electromagnetism raises a question for Humean Supervenience as I stated it. Denis Robinson (1989) has asked: is a vector field an arrangement of local qualities? I said qualities were intrinsic; that means they can never differ between duplicates; and I would have said offhand that two things can be duplicates even if they point in different directions. Maybe this last opinion should be reconsidered, so that vector-valued magnitudes may count as intrinsic properties. What else could they be? Any attempt to reconstrue them as relational properties seems seriously artificial.

Humean Supervenience is meant to be contingent: it says that among *worlds like ours*, no two differ without difference in the arrangement of qualities. But

when is a world like ours? I used to say: when it's a world of the "inner sphere", free of fundamental properties or relations that are alien to our world. Sally Haslanger (forthcoming) has shown that this answer probably won't do. One lesson of the Armstrong (1980)¹ spinning sphere (also known as the Kripke² spinning disk) is that one way to get a difference between worlds with the exact same arrangement of local qualities is to have things that are bilocated in spacetime. Take two worlds containing spheres of homogeneous matter, unlike the particulate matter of our world; in one world the sphere spins and in the other it doesn't; but the arrangement of local qualities is just the same. These are worlds in which things persist through time not by consisting of distinct temporal parts, but rather by bilocation in spacetime: persisting things are wholly present in their entirety at different times. The difference between the spinning and the stationary spheres is a difference in the pattern of bilocation. No worries for Humean Supervenience, so I thought: I believe that ours is a temporal-parts-world, therefore neither of the worlds in the story is a world like ours. But why assume that things that indulge in bilocation must differ in their fundamental nature from things that don't? Why think that if ours is a temporal-parts-world, then otherworldly bilocated things must have properties alien to our world? No good reason, I fear. Haslanger's point seems well taken. I still want to insist that if ours is a temporal-parts-world, then bilocation-worlds don't count as "worlds like ours", but I think I must abandon my former reason why not.

2. Symmetry and frequency

Chance is objective single-case probability: for instance, the 50% probability that a certain particular tritium atom will decay sometime in the next 12.26 years. Chance is not the same thing as degree of belief, or *credence* as I'll call it; chance is neither anyone's actual credence nor the credence warranted by our total available evidence. If there were no believers, or if our total evidence came from misleadingly unrepresentative samples, that wouldn't affect chance in any way.

Nevertheless, the chance of decay is connected as follows to credence: if a rational believer knew that the chance of decay was 50%, then almost no matter what else he might or might not know as well, he would believe to degree 50% that decay was going to occur. *Almost* no matter; because if he had reliable news from the future about whether decay would occur, then of course that news would legitimately affect his credence.

¹ This paper was presented at the 1976 conference of the Australasian Association of Philosophy.

² Saul Kripke, "Identity Through Time", presented at the 1979 conference of the American Philosophical Association, Eastern Division, and elsewhere.

This connection between chance and credence is an instance of what I call the *Principal Principle* (Lewis 1980).³ We shall see much more of it as we go on, both as the key to our concept of chance and as an obstacle to Humean analyses.

If Humean Supervenience is true, then contingent truths about chance are in the same boat as all other contingent truths: they must be made true, somehow, by the spatiotemporal arrangement of local qualities. How might this be? Any satisfactory answer must meet a severe test. The Principal Principle requires that the chancemaking pattern in the arrangement of qualities must be something that would, if known, correspondingly constrain rational credence. Whatever makes it true that the chance of decay is 50% must also, if known, make it rational to believe to degree 50% that decay will occur.

In simple cases, two candidates for chancemakers come to mind: symmetries and frequencies. Take symmetries first. Suppose a drunkard is wandering through a maze of T-junctions, and at each junction we can find nothing that looks like a relevant difference between the case that he turns left and the case that he turns right. We could well understand if rational credence had to treat the cases alike, for lack of a relevant difference. If the symmetry is something that would, if known, constrain credence, then it is suitable to serve as a chancemaker. In short, Humean chances might be based on a principle of indifference.

We know, of course, that an unrestricted principle of indifference is inconsistent. If we define partitions of alternative cases by means of ingeniously hoked-up properties, we can get the principle to say almost anything we like. But let us assume, as we should and as we already have done, that we can somehow distinguish natural properties from hoked-up gerrymanders. And let us apply the principle only to natural partitions. We still have no guarantee that we will get univocal answers, but we no longer know that we will not.

So far, so good; but I still I have two reservations about symmetries as chancemakers. For one thing, there is no reason to think we have symmetries to underlie all the chance phenomena we think there are. It would be nice to think that each tritium atom contains a tiny drunkard in a maze of symmetrical T-junctions, and the atom decays when its drunkard finds his way out. But this is sheer fantasy. So far as we know, nothing remotely like it is true.

More important, symmetries are only defeasible constrainters of rational credence. Therefore they can be only defeasible chancemakers. The symmetry of the T-junctions would no longer require 50-50 division of credence if we also knew that, despite this symmetry, the drunkards turn right nine times out of ten. Now we do have a relevant difference between left and right turns. Frequencies can defeat symmetries. And when symmetries are undefeated, that is because the frequencies are such as not to defeat them. So now it looks as if frequencies are the real chancemakers.

A frequency is the right sort of thing to be a Humean chancemaker: it is a pattern in the spatiotemporal arrangement of qualities. And we can well understand how frequencies, if known, could constrain rational credence. So far, so good.

³ Reprinted with added postscripts in Lewis (1986). See also Mellor (1971).

The simplest frequency analysis of single-case chance will just say that the chance of a given outcome in a given case equals the frequency of similar outcomes in all cases of exactly the same kind.

Again, this would be worse than useless if we couldn't distinguish natural from gerrymandered kinds; again, we could get the analysis to yield almost any answer we liked. But we can distinguish. (If we could not, puzzles about chance would be the least of our worries.) Further, nature has been kind to us. Large chance systems seem to be put together out of many copies of very small chance systems; and very small chance systems often do come in enormous classes of exact copies. You see one tritium atom, you've seen them all.

In some cases, I think this simple frequency analysis is near enough right. But it has its limits, and even when it works well, I'd like to see it subsumed under something more general. It is only plausible when we do have the enormous classes of exact copies. Tritium atoms are abundant; not so for atoms of unobtainium³⁴⁶. It's hard to make the stuff; in fact in all of space and time, past present and future, there only ever exist two Un³⁴⁶ atoms: one with a lifetime of 4.8 microseconds, as chance would have it; the other with a lifetime of 6.1 microseconds. So exactly half of all Un³⁴⁶ atoms decay in 4.8 microseconds. What does this frequency make true concerning the half-life of Un³⁴⁶, in other words concerning the chance that an atom of it will decay in a given time? Next to nothing, I should think.

Further, consider unobtainium³⁴⁹. This isotope is even harder to make, and in fact there is not one atom of it in all of space and time. Its frequency of decay in a given time is undefined: 0/0. If there's any truth about its chance of decay, this undefined frequency cannot be the truthmaker.

The problem of unobtainium, in both versions, may set us thinking of frequencies not in our actual world, but rather in counterfactual situations where unobtainium is abundant. (Maybe even infinitely abundant, but let's ignore the issues that will raise.) I think that's a blind alley. Different abundant-unobtainium worlds will have different decay frequencies; to ask what the frequency would be if unobtainium were abundant, we have to select those of the abundant-unobtainium worlds that are closest to actuality. Suppose indeed that we have the same decay frequency in all these worlds. I ask: what makes them closest? It must be something *X* that just these abundant-unobtainium worlds, some of the ones with the right decay frequency, have in common with our actual world. But then it is *X*, here in our actual world where unobtainium is scarce or absent, that is the real chancemaker. The abundant-unobtainium worlds that also have *X* are just a side-show.

Another well-known problem for simple frequentism: if spacetime is finite and chance systems get only just so small, then all frequencies are rational numbers. Yet the great majority of real numbers are irrational, and we have no pre-philosophical reason to doubt that chances in a finite world can take irrational values.

The answer to our problems about unobtainium lies in remembering that single-case chances follow from general probabilistic laws of nature. (At least, the

ones we know about do; and I think it's a spoils-to-the-victor question whether the same goes in general. The analysis I shall put forward can't handle lawless chances, and I take that to be no problem.) There are general laws of radioactive decay that apply to all atoms. These laws yield the chance of decay in a given time, and hence the half-life, as a function of the nuclear structure of the atom in question. (Or rather, they would yield the chance but for the intractability of the required calculation.) Unobtainium atoms have their chances of decay not in virtue of decay frequencies for unobtainium, but rather in virtue of these general laws. The appeal to laws also solves our problem about irrational values. Whether or not the world is finite, there is no reason why the function of nuclear structure that is built into the law can only yield rational values.

In general, probabilistic laws yield history-to-chance conditionals. For any given moment, these conditionals tell us the chance distribution over alternative future histories from that moment on, as a function of the previous history of particular facts up to and including that moment. The historical antecedents are of course given by the arrangement of qualities. The laws do the rest.

So the appeal to laws just postpones our problem. What pattern in the arrangement of qualities makes the chances? In part, features of history up to the moment in question. For the rest, it is the pattern that makes the probabilistic laws, whatever that is. So now we must turn to a Humean analysis of laws, and see whether we can extend that to cover probabilistic laws. I think we can.

3. The best-system analysis of law

Ramsey once thought that laws were "consequences of those propositions which we should take as axioms if we knew everything and organized it as simply as possible in a deductive system" (1990, p. 150).⁴ I trust that by "it" he meant not everything, but only as much of everything as admits of simple organization; else everything would count as a law. I would expand Ramsey's idea thus (see Lewis 1973, p. 73). 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.

Some familiar complaints seem to me question-begging. (See Armstrong (1983, pp. 40-59); van Fraassen (1989, pp. 45-51).) If you're prepared to agree that theorems of the best system are rightly called laws, presumably you'll also want to say that they underlie causal explanations; that they support counterfactuals; that they are not mere coincidences; that they and their consequences are in

⁴ Here I quote Ramsey's later summary of his former view.

some good sense necessary; and that they may be confirmed by their instances. If not, not. It's a standoff—spoils to the victor. Other complaints are more worrisome. Like any regularity theory, the best-system analysis says that laws hold in virtue of patterns spread over all of space and time. If laws underlie causation, that means that we are wrong if we think, for instance, that the causal roles of my brain states here and now are an entirely local matter. That's an unpleasant surprise, but I'm prepared to bite the bullet.

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. I can admit that *if* nature were unkind, and *if* disagreeing rival systems were running neck-and-neck, than lawhood might be a psychological matter, and that would be very peculiar. I can even concede that in that case the theorems of the barely-best system would not very well deserve the name of laws. But I'd blame the trouble on unkind nature, not on the analysis; and I suggest we not cross these bridges unless we come to them.

(Likewise for the threat that two very different systems are tied for best. (See Armstrong (1983, pp. 70-71); van Fraassen (1989, pp. 48-49).) I used to say that the laws are then the theorems common to both systems, which could leave us with next to no laws. Now I'll admit that in this unfortunate case there would be no very good deservers of the name of laws. But what of it? We haven't the slightest reason to think the case really arises.)

The best-system analysis is Humean. The arrangement of qualities provides the candidate true systems, and considerations of simplicity and strength and balance do the rest.

But so far, we don't have probabilistic laws. If chances were somehow given, we could just include them in the subject matter of the competing true systems, and go on as before.⁵ But chances are not yet given. We decided that the chance-making patterns in the arrangement of qualities had to include the lawmaking patterns for the probabilistic laws that determine the chances in all the different cases.

4. The best-system analysis of law and chance together

So we modify the best-system analysis to make it deliver the chances and the laws that govern them in one package deal. 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. We cannot yet require them to be true in what they say about chance, because we have yet to say what chance means; our systems are as yet not fully interpreted. 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.

As before, some systems will be simpler than others. Almost as before, some will be stronger than others: some will say either what will happen or what the chances will be when situations of a certain kind arise, whereas others will fall silent both about the outcomes and about the chances. And further, some will fit the actual course of history better than others. That is, the chance of that course of history will be higher according to some systems than according to others. (Though it may well turn out that no otherwise satisfactory system makes the chance of the actual course of history very high; for this chance will come out as a product of chances for astronomically many chance events.) Insofar as a system falls silent, of course it fits whatever happens.

The virtues of simplicity, strength, and fit trade off. The best system is the system that gets the best balance of all three. As before, the laws are those regularities that are theorems of the best system. But now some of the laws are probabilistic. So now we can analyse chance: the chances are what the probabilistic laws of the best system say they are. (See Lewis (1986, pp. 128-9).)

As before, we may reasonably hope that the best system is very far ahead of the rest; and very robustly ahead, so that the winner of the race does not depend

⁵ Well, not quite as before: we'd have to impose constraints that would, for instance, disqualify a system according to which some law had once had some chance of not being true. See Lewis (1986, pp. 124-8).

on how we happen to weigh the various desiderata. How well the laws and chances deserve their names should depend on how kind nature has been in providing a decisive front runner. The prospect is best if the chance events are not too few and not too miscellaneous.

In the simplest case, the best-system analysis reduces to frequentism. Suppose that all chance events—more precisely, all events that the leading candidates for best system would deem to be chancy—fall into one large and homogeneous class. To fall silent about the chances of these events would cost too much in strength. To subdivide the class and assign different chances to different cases would cost too much in simplicity. To assign equal single-case chances that differed from the actual frequency of the outcomes would cost too much in fit. For we get the best fit by equating the chances to the frequency; and the larger the class is, the more decisively is this so.

But suppose the class is not so very large; and suppose the frequency is very close to some especially simple value—say, 50-50. Then the system that assigns uniform chances of 50% exactly gains in simplicity at not too much cost in fit. The decisive front-runner might therefore be a system that rounds off the actual frequency.

Or suppose the class is not so very homogeneous. It divides, in virtue of not-too-gruesome classifications, into a few large subclasses that exhibit different frequencies of outcomes. (To take an extreme case, a class of *J*s that are 50% *K*s might exhibit a regular alternation in time: *K*, not-*K*, *K*, not-*K*,....) Then a system that assigns unequal chances in different subclasses will gain greatly in fit at not too much cost in simplicity.

Or suppose the class is inhomogeneous in a different way. Each member is associated with some value of a continuously variable magnitude *M*. (In an extreme case, no two members have exactly the same value of *M*.) If we divide into as many subclasses as there are values of *M*, the subclasses will be too numerous for simplicity and too small for the frequencies in them to mean much. Instead, the best system will contain a functional law whereby chance depends on the value of *M* in that particular case. Different candidate systems will use different functions. Some functions, for instance a constant function, will go too far in gaining simplicity at the expense of fit. Others will do the opposite. We may hope that some function—and hence the system that employs it—will be just right, and hence the decisive front runner.

In this last case, and others that combine several of the complications we've considered, frequentism has been left behind altogether. That's how we can get decay chances for Un^{346} , and even for Un^{349} , in virtue of chancemaking patterns that don't involve decay frequencies for unobtainium itself.

Despite appearances and the odd metaphor, this is not epistemology! You're welcome to spot an analogy, but I insist that I am not talking about how evidence determines what's reasonable to believe about laws and chances. Rather, I'm talking about how nature—the Humean arrangement of qualities—determines

what's true about the laws and chances. Whether there are any believers living in the lawful and chancy world has nothing to do with it.

It is this best-system analysis of law and chance together that I've wanted to believe for many years. Until very recently, I thought I knew a decisive reason why it couldn't be true. Hence my lamentation about the big bad bug. But now that Michael Thau has shown me the way out, I can endorse the best-system analysis with a clear conscience.

5. *Undermining*

The big bad bug bites a range of different Humean analyses of chance. Simple frequentism falls in that range; so does the best-system analysis. Let's suppose that we have a Humean analysis which says that present chances supervene upon the whole of history, future as well as present and past; but not upon the past and present alone. That's so if present chances are given by frequencies throughout all of time. That's so also if present chances are given by probabilistic laws, plus present conditions to which those laws are applicable, and if those laws obtain in virtue of the fit of candidate systems to the whole of history.

Then different alternative future histories would determine different present chances. (Else the future would be irrelevant, and present chances would be determined by the past and present alone.) And let's suppose, further, that the differences between these alternative futures are differences in the outcomes of present or future chance events. Then each of these futures will have some non-zero present chance of coming about.

Let F be some particular one of these alternative futures: one that determines different present chances than the actual future does. F will not come about, since it differs from the actual future. But there is some present chance of F . That is, there is some present chance that events would go in such a way as to complete a chancemaking pattern that would make the present chances different from what they actually are. The present chances *undermine* themselves.

For instance, there is some minute present chance that far more tritium atoms will exist in the future than have existed hitherto, and each one of them will decay in only a few minutes. If this unlikely future came to pass, presumably it would complete a chancemaking pattern on which the half-life of tritium would be very much less than the actual 12.26 years. Certainly that's so under simple frequentism, and most likely it's so under the best-system analysis as well. Could it come to pass, given the present chances? Well, yes and no. It could, in the sense that there's non-zero present chance of it. It couldn't, in the sense that its coming to pass contradicts the truth about present chances. If it came to pass, the truth about present chances would be different. Although there is a certain chance that this future will come about, there is no chance that it will come about while still having the same present chance it actually has. It's not that if this future came about,

the truth about the present would change retrospectively. Rather, it would never have been what it actually is, and would always have been something different.

This undermining is certainly very peculiar. But I think that, so far, it is no worse than peculiar. I would not join Bigelow, Collins, and Pargetter (1993, pp. 443-62) when they intuit a “basic chance principle” to exclude it outright. For I think the only basic principle we have about chance, the principle that tells us all we know, is the Principal Principle. And at first sight the Principal Principle says nothing against undermining. It concerns, rather, the connection between chance and credence.

But look again, and it seems that the Principal Principle does rule out undermining. It was this discovery that led me to despair of a Humean analysis of chance. (See Lewis 1986, pp. xiv-xvii, 111-3, 130).

Now is the time to take a closer look at what the Principle says.⁶ Above, I applied it to the case of someone who knows the chances, but that is a special case. The general case involves not knowledge but conditioning. Take some particular time—I’ll call it “the present”, but in fact it could be any time. Let C be a rational credence function for someone whose evidence is limited to the past and present—that is, for anyone who doesn’t have access to some very remarkable channels of information. Let P be the function that gives the present chances of all propositions. Let A be any proposition. Let E be any proposition that satisfies two conditions. First, it specifies the present chance of A , in accordance with the function P . Second, it contains no “inadmissible” information about future history; that is, it does not give any information about how chance events in the present and future will turn out. (We don’t assume that E is known; that extra assumption would yield the special case considered earlier.) Then the Principal Principle is the equation

$$C(A/E) = P(A).$$

Now take A to be F , our alternative future history that would yield present chances different from the actual ones; and let E be the whole truth about the present chances as they actually are. We recall that F had some present chance of coming about, so by the Principal Principle, $C(F/E) \neq 0$. But F is inconsistent with E , so $C(F/E) = 0$. Contradiction. I could tolerate undermine as merely peculiar. But not contradiction!

6. No refuge

It is because the chancemaking pattern lies partly in the future that we have some chance of getting a future that would undermine present chances. This problem would go away if we could assume that the chancemaking pattern lay entirely in the past. That would be so if all of our history-to-chance conditionals, which

⁶ I have simplified my formulation in (1980) by combining the propositions there called X and E .

specify exactly what chances would follow any given initial segment of history, were necessarily true. (See Lewis 1980, final section.)

We dare not assume this. First, because of the problem of the early moment. There might be a beginning of time; or at least a beginning of the part of time in which certain kinds of chancy phenomena go on. What could make the chances at a moment not long after the beginning? (Or at the beginning itself?) There's not much room for any sort of chancemaking pattern in the time before this early moment. To the extent that chancemaking patterns are just frequencies in large and uniform classes, the problem is that the relevant classes, if confined to the time before the early moment, may be ridiculously small. But the problem won't go away if instead we take the different sort of chancemaking pattern envisaged by the best-system analysis.

Second, because of the problem of fluctuation. We usually think that there are laws, and hence regularities, of uniform chances. All tritium atoms throughout space and time have precisely the same chance of decaying in a given period. But the different atoms are preceded by different initial segments of history, and it is not to be expected that the different chancemaking patterns in these different segments will all make precisely the same chance of decay.

(Taking these two problems together, we get what I'll call the problem of drift. For simple frequentism, it goes as follows. Suppose that early on, *J*s divide about 50-50 between *K*s and not-*K*s, but so far there haven't been many *J*s altogether. Then we should expect that there might chance to be a run of *K*s, or not-*K*s, that would significantly raise, or lower, the chance that the next *J* would be a *K*. Such a run might perpetuate itself, with the result that the chance of a *J* being a *K* would drift almost to one, or zero, and remain there for a long time after. But we do not at all expect the chances in radioactive decay, say, to undergo any such drift.)

At this point, my opponents will doubtless say that I have done their work for them. I have refuted the position I wanted to hold. It only remains for me to concede defeat, and agree that the chancemakers are not, after all, patterns in the arrangement of qualities. They are something else altogether: special chancemaking relations of universals, primitive facts about chance, or what have you. (See Armstrong (1980, pp. 128-36); Tooley (1987, pp. 142-60).)

But I think there is no refuge here. 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.⁷

Unless, of course, you can convince me first that this special relation is a chancemaking relation: that the fact that $N^*(J,K)$ makes it so, for instance, that

⁷ Thanks here to Alex Byrne.

each J has 50% chance of being K . But you can't just tell me so. You have to show me. Only if I already see—dimly will do!—how knowing the fact that $N^*(J,K)$ should make me believe to degree 50% that the next J will be a K , will I agree that the N^* relation deserves the name of chancemaker that you have given it.

The same complaint applies, by the way, to theories that qualify technically as Humean because the special primitive chancemaking whatnots they posit are said to be qualities instantiated at points. Again, I can only agree that the whatnots deserve the name of chancemakers if I can already see, disregarding the names they allegedly deserve, how knowledge of them constrains rational credence in accordance with the Principal Principle.

7. *The beginning of a solution*

Our problem, where F is an unactualized future that would undermine the actual present chances given by E , is that $C(F/E) = 0$ because F and E are inconsistent, but $C(F/E) \neq 0$ by the Principal Principle because E specifies that F has non-zero chance of coming about. If that use of the Principal Principle is fallacious, the contradiction goes away. We're left with nothing worse than peculiar undermining.

But that use of the Principal Principle *is* fallacious, if indeed the present chances are made by a pattern that extends into the future. Then E bears inadmissible information about future history: it excludes the future F , and it likewise excludes all other futures that would undermine the present chances given by E . Since E is inadmissible, the Principal Principle does not apply. The fatal move that led from Humeanism to contradiction is no better than the obvious blunder:

$C(\text{the coin will fall heads/it is fair and will fall heads in 99 of the next 100 tosses}) = 1/2$

or even

$C(\text{the coin will fall heads/it is fair and it will fall heads}) = 1/2.$

Victory!—Another such victory and I am undone. What we have just seen is that if chancemaking patterns extend into the future, then *any* use of the Principal Principle is fallacious. For *any* proposition that bears information about present chances thereby bears information about future history. The Principal Principle never applies; and yet without it I deny that we have any handle at all on the concept of chance.

I saw this, but I didn't dare to believe my eyes. I wrote

If anyone wants to defend the best-system theory of laws and chances ... I suppose the right move would be to cripple the Principal Principle by declaring that information about the chances at a time is not, in general, admissible at that time; and hence that hypothetical information about chances, which can join with admissible historical information to imply chances at a time, is likewise inadmissible. The reason would be that, under the proposed analysis of chances, information about present chances is a disguised form of inadmissible information about future

history—to some extent, it reveals the outcomes of matters that are presently chancy. That crippling stops all versions of our *reductio*... I think the cost is excessive; in ordinary calculations with chances, it seems intuitively right to rely on this ... information. So, much as I would like to use the best-system approach in defence of Humean supervenience, I cannot support this way out of our difficulty. (1986, pp. 130-1)

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

8. *The solution completed*

Now we're ready at last to hear from Thau. As follows.

First, admissibility admits of degree. A proposition E may be imperfectly admissible because it reveals something or other about future history; and yet it may be very nearly admissible, because it reveals so little as to make a negligible impact on rational credence.

Second, degrees of admissibility are a relative matter. The imperfectly admissible E may carry lots of inadmissible information that is relevant to whether B , but very little that is relevant to whether A .

Third, near-admissibility may be good enough. If E specifies that the present chance of A is $P(A)$, and if E is nearly admissible relative to A , then the conclusion that $C(A/E) = P(A)$ will hold, if not exactly, at least to a very good approximation. If information about present chances is never perfectly admissible, then the Principal Principle never can apply strictly. But the Principle applied loosely will very often come very close, so our ordinary reasoning about chance and credence will be unimpaired. Only a few peculiar applications will be so badly crippled that they cannot be regained even as approximations.

And one of these applications that cannot be regained will be the one that figured in our *reductio*. If F is a future that would undermine the chances specified in E , then, relative to F , E is as inadmissible as it could possibly be. For E flatly contradicts F . Our use of the Principal Principle to conclude that $C(F/E)$ is non-zero was neither strictly nor loosely correct. Hence it no longer stands in the way of the correct conclusion that $C(F/E) = 0$.

9. *Correcting the Principal Principle*

We face a question. If the old Principal Principle applies only as an approximation, what is it an approximation to? How, exactly, is chance related to credence?

Can we find a new, corrected Principal Principle that works exactly when the old one works only approximately?

Here I shall present only a special case of the correction I favour, applicable only to a special case of the old Principal Principle. (However, the old Principle in full generality follows from the special case; and in parallel fashion, a new general Principal Principle follows from the corrected treatment of the special case.) Ned Hall and I proposed the same correction independently. In his companion paper (1994), he gives a much fuller discussion of the correction, its motivation, and its consequences. To reach the special case I want to present,⁸ we consider a certain very informative proposition, as follows. Let H_w be the proposition giving the complete history of world w up to, and including, time t —as it might be, the actual world up to the present. Let T_w be the complete theory of chance for world w —a proposition giving all the probabilistic laws, and therefore all the true history-to-chance conditionals, that hold at w . And let P_w be the chance distribution at time t and world w . If T_w were admissible, then the conjunction $H_w T_w$ also would be admissible. Further, it would specify the chance (at t , at w) of any proposition A . So we could put it for E in the old Principal Principle. Dropping the subscripts henceforth by way of abbreviation, we would have:

$$(OP) \quad C(A/HT) = P(A).$$

The correction I favour is to replace (OP) by this new version:

$$(NP) \quad C(A/HT) = P(A/T).$$

If T were perfectly admissible, then (OP) would be correctly derived as an instance of the old Principal Principle. If so, then also the change from old to new would be no change at all. For in that case, we would have $P(T) = 1$. (Proof: for A in (OP), put T itself.) Hence we would have $P(A) = P(A/T)$, making (OP) and (NP) equivalent.

But if, as the Humean thinks, there are undermining futures with non-zero present chance that make T false, then T rules out these undermining futures. If so, then, T and HT are not perfectly admissible; (OP) is not correctly derived; the old Principal Principle cannot be applied to determine $C(A/HT)$; $P(T) \neq 1$; and, exceptional cases aside, it will turn out that $P(A) \neq P(A/T)$. If so, then I say we should accept (NP) as our new, corrected Principal Principle. We should abandon (OP) except insofar as it is a convenient approximation to (NP).

By conditionalising credence or chance on T , we ignore undermining futures. The trouble with (OP) is that on the left-hand side we do conditionalise on T , but on the right-hand side we don't. No harm would come of this discrepancy if there weren't any undermining futures to worry about. But if there are, then it is only to be expected that the discrepancy will cause trouble. We've already seen the trouble: taking A as an undermining future F , $C(F/HT) = 0$, but $P(F) \neq 0$. To remove the discrepancy, we conditionalise the right-hand side as well as the left:

⁸ It is this special case that appears as "the Principal Principle Reformulated" in Lewis (1980).

$P(F/T) = 0$, so all's well. The Principal Principle, thus corrected, no longer stands in the way of Humean Supervenience. By my lights as a Humean, that is reason enough to correct it. In his companion paper, Hall argues that we have other good reasons as well.

In formulating (OP) and (NP), we didn't mention admissibility or near-admissibility. Rather, we fixed on the particular propositions H and T , noting only in passing that H is uncontroversially admissible while T 's admissibility is questionable. But we can regain contact with Thau's lesson about imperfect and relative admissibility, as follows. Define the *admissibility quotient* of T (and of HT), relative to A , as the quotient $P(A/T)/P(A)$. Then (OP) is a good approximation just to the extent that the admissibility quotient is close to one. (Or if $P(A/T)$ and $P(A)$ are both zero, making the quotient undefined. Let us ignore this case.) We have perfect admissibility when the quotient is one exactly; that's so just when A and T are probabilistically independent with respect to P . T is perfectly admissible relative to A if (but not only if) A is any proposition that is entirely about past and present history. For instance, T is perfectly admissible relative to H .

If, on the other hand, A is an undermining future that contradicts T , then the admissibility quotient is zero, and (OP) is not at all a good approximation to (NP).

More commonplace applications of the Principal Principle, where A typically concerns a small sample from a big population, are quite another story: we can expect an admissibility quotient very close to one. Here's an easy example. Hitherto, exactly $2/3$ of J s have been K s; and we know somehow, that there are 10,002 more J s still to come. Our complete theory of chance, T , says that every J , anywhere and anywhen, has $2/3$ chance of being a K . Simple frequentism—a simplistic form of Humeanism, but good enough for a toy example—says that T holds iff exactly $2/3$ of all J s are K s. Equivalently, given what we know about the past: iff $2/3$ of the future J s are K s. Now let A say that three out of the next four J s will be K s. It is a routine matter to calculate the admissibility quotient of T with respect to A .⁹ It turns out to be about 1.00015. So this time, though (OP) is not exactly right, it is indeed a very good approximation. If our example had concerned the next four J s from a population of astronomical size, the approximation would have been much better still.

$${}^9 P(A) = \left(\frac{2}{3}\right)^4 ; P(A/T) = \frac{P(AT)}{P(T)} = \frac{P(A) \binom{2}{3}^{6665} \binom{1}{3}^{3333} \left(\frac{9998!}{6665! 3333!}\right)}{\binom{2}{3}^{6668} \binom{1}{3}^{3334} \left(\frac{10002!}{6668! 3334!}\right)}$$

When we divide out to obtain $P(A/T)/P(A)$, most terms cancel and we are left with a product of a few fractions. The number of these fractions is the sample size, 4, and each of them differs from one by something of the order of the reciprocal of the population of future J 's, 10,002.

10. Against perfectionism

Our new version of the Principal Principle is better by Humean lights; but for myself, I still find the old one more intuitive. (Once we return to the general case, the new version gets quite messy.¹⁰) So I still say that the old Principle is “the key to our concept of chance”. And yet it’s only approximately right, and that only sometimes. Chance can be defined as that feature of Reality that obeys the old Principle, yet chance doesn’t quite obey it! Isn’t this incoherent?

Not at all. Let’s put it this way. A feature of Reality deserves the name of chance to the extent that it occupies the definitive role of chance; and occupying the role means obeying the old Principle, applied as if information about present chances, and the complete theory of chance, were perfectly admissible. Because of undermining, nothing perfectly occupies the role, so nothing perfectly deserves the name. But near enough is good enough. If nature is kind to us, the chances ascribed by the probabilistic laws of the best system will obey the old Principle to a very good approximation in commonplace applications. They will thereby occupy the chance-role well enough to deserve the name. To deny that they are *really* chances would be just silly.

It’s an old story. Maybe nothing could perfectly deserve the name “sensation” unless it were infallibly introspectible; or the name “simultaneity” unless it were a frame-independent equivalence relation; or the name “value” unless it couldn’t possibly fail to attract anyone who was well acquainted with it. If so, then there are no perfect deservers of these names to be had. But it would be silly to lose our Moorings and deny that there existed any such things as sensations, simultaneity, and values. In each case, an imperfect candidate may deserve the name quite well enough.¹¹

Department of Philosophy
Princeton University
Princeton
New Jersey
08544
USA

DAVID LEWIS

¹⁰ Here is how we regain generality. By probability theory, including a suitable form of additivity—as it might be, a principle of infinite additivity of infinitesimal quantities in a suitable non-standard model—we find that $C(A/E)$ equals the $C(-/E)$ -expectation of the quantity whose value at each world w is $C(A/EH_wT_w)$. Assuming that proposition E is entirely about history up to t and chances at t , we find that whenever $C(-/E)$ assigns w positive credence, EH_wT_w simplifies to H_wT_w . So according to (OP), $C(A/E)$ is the $C(-/E)$ -expectation of the quantity whose value at each w is $P_w(A)$; whereas according to (NP), it is the $C(-/E)$ -expectation of the quantity whose value at each w is $P_w(A/T_w)$. Assuming further that E specifies the chance at t of A , so that $P_w(A)$ has a constant value at all E -worlds, we find that the expectation of chance derived from (OP) just equals this same value. Sad to say, the parallel expression derived from (NP) remains unsimplified.

¹¹ Thanks, above all, to Michael Thau and Ned Hall; and to many others who have helped me by discussions of this material.

REFERENCES

- Armstrong, D.M. 1980: "Identity Through Time", in Peter van Inwagen, ed., *Time and Cause: Essays Presented to Richard Taylor*. Dordrecht: Reidel.
- 1983: *What is a Law of Nature?*. Cambridge: Cambridge University Press.
- Bigelow, John 1988: *The Reality of Numbers*. Oxford: Oxford University Press.
- Bigelow, John, Collins, John and Pargetter, Robert 1993: "The Big Bad Bug: What are the Humean's Chances". *British Journal of Philosophy of Science*, 44, pp. 443-62.
- Forrest, Peter 1988: *Quantum Metaphysics*. Oxford: Blackwell.
- Hall, Ned 1994: "Correcting the Guide to Objective Chance". *Mind*, 103, pp. 505-518.
- Haslanger, Sally (forthcoming): "Humean Supervenience and Enduring Things". *Australasian Journal of Philosophy*.
- Lewis, David 1973: *Counterfactuals*. Oxford: Blackwell.
- 1980: "A Subjectivist's Guide to Objective Chance", in Richard C. Jeffrey, ed., *Studies in Inductive Logic and Probability*, vol. II. Berkeley: University of California. Reprinted with added postscripts in Lewis 1986, pp. 83-132.
- 1986: *Philosophical Papers*, vol. II. Oxford: Oxford University Press.
- Mellor, D.H. 1971: *The Matter of Chance*. Cambridge: Cambridge University Press.
- Menzies, Peter 1989: "Probabilistic Causation and Causal Processes: A Critique of Lewis". *Philosophy of Science*, 56, pp. 642-63.
- Ramsey, F.P. 1990: *Philosophical Papers*. Cambridge: Cambridge University Press.
- Robinson, Denis 1989: "Matter, Motion and Humean Supervenience". *Australasian Journal of Philosophy*, 67, pp. 394-409.
- Thau, Michael 1994: "Undermining and Admissibility". *Mind*, 103, pp. 491-503.
- Tooley, Michael 1987: *Causation: A Realist Approach*. Oxford: Oxford University Press.
- van Fraassen, Bas 1989: *Laws and Symmetry*. Oxford: Oxford University Press.