Brit. J. Phil. Sci. **54** (2003), 27–41

Principled Chances

Jonathan Schaffer

ABSTRACT

There are at least three core principles that define the chance role: (1) the Principal Principle, (2) the Basic Chance Principle, and (3) the Humean Principle. These principles seem mutually incompatible. At least, no extant account of chance meets more than one of them. I offer an account of chance which meets all three: L*-chance. So the good news is that L*-chance meets (1)–(3). The bad news is that L*-chance turns out unlawful and unstable. But perhaps this is not such bad news: L*-chance turns out to at least approximate plausible additional core principles concerning lawfulness and stability. And perhaps there is better news: one may treat 'chance' as vague, in a way that allows every core principle of chance to be met.

1 Introduction

There are at least three core principles that define the chance role: (1) the Principal Principle (David Lewis [1980]); (2) the Basic Chance Principle (John Bigelow, John Collins, and Robert Pargetter [1993]), and (3) the Humean Principle (Lewis [1986]). Lewis ([1994]) offers a conception of chance, L-chance, that meets (3), approximates (1), and violates (2). The Necessitarian offers a conception of chance, N-chance, that meets (2), but violates (1) and (3). I offer a conception of chance, L*-chance, that meets (1)–(3).

So is L*-chance chance? Perhaps so, since only L*-chance meets (1)–(3). Or perhaps not, since L*-chance turns out to be unlawful and unstable. Or perhaps so again, since L*-chance turns out to at least approximate plausible additional core principles concerning lawfulness and stability. Or perhaps *maybe*: why must one even decide between L-chance and L*-chance?

© British Society for the Philosophy of Science 2003

2 Chance principles

There are at least three core principles that define the chance role: (1) the Principal Principle, (2) the Basic Chance Principle, and (3) the Humean Principle. These principles connect the concept of chance to the concepts of credence, possibility, and being (respectively). To a first approximation, information about chances should constrain rational credences, the existence of nonzero chances should entail realizing possibilities, and the magnitude of chance should be determined by certain basic features of reality. Principles (1)–(3) are plausible ways of being precise about these connections.

The Principal Principle: The Principal Principle (PP) connects information about chances to rational credences, via the formula C(A/XE)=x. Here C is any reasonable initial credence function, A is the target proposition, X is the proposition that the chance of A (at time t) is x, and E is any admissible evidence. The PP says that, given the information that the chance of A (at t) is x, as well as any other admissible information, one's credence in A should be x.

Lewis (the principal behind the principle) maintains that the PP captures 'all we know about chance' ([1980], pp. 266–7; see also [1986], pp. xv–xvi, [1994], pp. 484–5, p. 489). While I think that we know more about chance, such as its further connections to possibility and to being, I would agree that the PP is one of the core principles that define the chance role.

The Basic Chance Principle: The Basic Chance Principle (BCP) connects the existence of nonzero chances to realizing possibilities, via the conditional: If $P_{tw}(A)=x$ (for $x>0$), then there exists a world w' such that (i) A is true at w'; (ii) w' matches w up to t, and (iii) $P_{tw'}(A)=x$. Here P_{tw} is a probability function,[1] with a time and world index. The BCP says that, if the chance of A is nonzero (at t in w), then A is possible, in the sense that there is a world w' with the same history up until t and the same chance of A, at which A occurs.

Bigelow, Collins, and Pargetter (the BCP behind the BCP) maintain 'anything that failed to satisfy the BCP would not deserve to be called *chance*' ([1993], p. 459). The BCP, they explain, is a necessary condition for 'grounding' the positive chance of A—if A couldn't obtain, or could only obtain given a different history or difference chances, then it seems that its alleged positive chance is spurious.

Which of the PP and BCP is more central to the chance role? Bigelow *et al* clearly consider the BCP at least as central as the PP. Lewis maintains that the PP is all that matters, calling violations of the BCP 'no worse than peculiar' ([1994], p. 237). This dispute seems rather unprincipled—what theoretical basis is there for determining the centrality of a given platitude? (For whatever it is worth, my judgment is that both the PP and the BCP are

[1] That is, P_{tw} is a function from a time t, world w, and argument A onto the unit interval [0, 1] in accord with the axioms of probability, so that, for instance, if $P_{tw}(A)=x$, then $P_{tw}(\sim A)=1\text{-}x$.

central.) Fortunately, there is no need to resolve this here. I take it that all hands would agree that a conception of chance that met both the PP and the BCP would be *ipso facto* preferable to one that didn't.

The Humean Principle: Humean Supervenience (HS) connects the magnitude of chances at a world to its arrangement of occurrent facts, via the supervenience claim: $P_{tw}(\text{-})$ supervenes on the distribution of occurrent facts at *w*.[2] Here an occurrent fact is a categorical, intrinsic quality of a region (or a field, or a portion of matter, or whatnot). HS says that the chances are determined by the occurrences.

Lewis (the head spokesman for HS) regards HS as (i) a plausible empirical claim, and (ii) a prerequisite for meeting the PP. HS is of course highly suspect. The main ground for suspicion, though, is the thought that HS is too restrictive to allow for an adequate account of modal concepts like chance.[3] Thus Necessitarians such as David Armstrong ([1983], [1997]) have postulated additional basic features of reality over and above the occurrent facts, to serve as the chance-makers. For Armstrong, there are more basic features of reality than just the occurrent facts. There are also basic law-and-chance facts. For Armstrong, the occurrent facts are the states-of-affairs: the inherence pattern of universals {F, G, H, . . .} in particular substances {a, b, c, . . .}. The law and chance facts are second-order degree-of-necessitation universals: N:*x* (F, G) (to be read as: F-occurrences necessitate G-occurrences to degree *x*).

I take it that all hands would agree that HS is the more economical hypothesis. If it works, there is no (or at least less) need to posit basic necessitation facts. So all hands should agree that a conception of chance that met HS would be *ceteris paribus* preferable to one that didn't. The only issue is whether this preference is satisfiable. HS should be regarded as a theoretical virtue, albeit not a decisive one.

Assuming that there are no further core principles of chance (see §5, however), a given relation deserves the name 'chance' to the extent that it meets (1) the PP, (2) the BCP, and (3) HS. A relation that meets all three would be a perfect realizer of the chance role. A relation that meets one or two may still be an imperfect realizer of the chance role, if no other relation does better. If there are multiple imperfect realizers of the chance role then one must determine which of (1)–(3) have priority. For whatever it is worth, my judgment is that (1) and (2) are roughly tied for first honors, and that (3) is a distant third. But no need to rank-order: *there is a perfect realizer of (1)–(3).*

[2] Lewis maintains that the supervenience should be on the *local* occurrent facts through *w*. Since locality is (i) empirically problematic, and (ii) not relevant here, I suppress it in the main text. Lewis also restricts the scope of HS to *inner worlds*, without alien properties.

[3] Another reason for being suspicious about HS is the concern that the Humean base of occurrent facts cannot be made sense of prior to the notions of law, cause, and chance (Sydney Shoemaker [1980]). I lack the space to discuss this concern here.

3 Chance principals

There are two principal conceptions of chance in the literature. They are (1) Lewisian chance, and (2) Necessitarian chance. Lewisian chance summarizes the frequencies, while Necessitarian chance is an irreducible theoretical posit.

Theories of chance are typically divided into (1) frequency theories, (2) propensity theories, and (3) subjectivist theories. Lewisian chance represents the most sophisticated expression of frequency theories. Necessitarian chance represents propensities (I use 'Necessitarian' broadly, to include all manner of propensity theories). I do not consider any representative of subjective chances here: subjective chances draw no connections at all between chance, possibility, and being, and so ought to be disqualified from the start.[4]

Lewisian chance: Lewisian chance (L-chance) is generated by the Best System Theory (BEST) of laws and chances. The BEST inaugurates a contest between axiom systems. The requirements for entry are that the system is (i) *true* in what it says about occurrent history, and (ii) *cautious* in that it never says both that A and that L-ch(A) < 1. The conditions for victory are that the system possesses an optimal balance of simplicity, strength, and fit, where fit measures the L-chance the system assigns to total history H: the greater L-ch(H), the better the fit.

Is L-chance chance? In other words, does $P_{tw}(A)$=L-$ch_{tw}(A)$? To begin with, L-chance respects HS by construction. It connects the magnitude of the chances to the pattern of the occurrences, by taking the former to summarize the latter.

L-chance does not meet the BCP or the PP, but it at least approximates the latter in mundane cases. The reason why L-chance does not meet either the BCP or the PP emerges from Lewis's *undermining* argument. Suppose that there is now some nonzero present chance that F: the future will hold far more 'coin flips' than the past, all or most of which will land heads.[5] On the

[4] The most sophisticated subjectivist theories reconstruct 'chance'-talk in terms of exchangeable credences over infinite repetitions of the trial (see Brian Skyrms [1984], and Frank Arntzenius [ms]) This goes some way to forging a connection between 'chance'-talk and being, since the subjectivist who updates her credences by conditionalizing on observed *results* will enjoy credences that mirror the 'chances' assigned by her objectivist counterpart, at least under most conditions. (Though what justifies *exchangeable* credences but the expectation that *the same chances* obtain on each trial?) In any case, I see no connection at all between 'chance' and possibility here. How could mere 'chance'-talk really constrain what is possible? Or, to anticipate some of the discussion in §5, I see no connection between 'chance' and lawhood either. (See note 19 for a little more on subjectivism). The thoroughgoing subjectivist might retreat to a subjectivist, antirealist construal of possibility and lawhood (and perhaps even of being). So I should perhaps offer a far more qualified dismissal of subjectivism: subjective chances do not appear to draw the right connections between 'chance', possibility-realistically-construed, and lawhood-realistically-construed. In the main text I assume a generally realist position.

[5] Here I am treating coin flips as basic indeterministic mechanisms with L-ch_{tw}(heads)=.5. The example is chosen purely for ease of exposition: the reader may feel free to substitute wave-function collapses, or radioactive decays, or some other more serious candidate for a basic indeterministic mechanism.

BEST, L-ch_{tw}(F) is nonzero. But if F came to pass, then the resulting pattern of occurrences would presumably be different enough to require different laws, under which L-ch_{tw} (heads) > .5.

Here the BCP is violated. L-ch_{tw}(F)=j (where j > 0), but there is no world w' at which (i) F is true, (ii) w' matches w up to t, and (iii) L-$ch_{tw'}$(F)=j. This is because satisfaction of conditions (i) and (ii) entails violation of (iii). If (i) and (ii) hold at w', the resulting difference in laws entails L-$ch_{tw'}$(F) > j.

The PP is violated as well. Let the target proposition be F, and consider the special case in which T_w is the complete theory of chance for w, and in which H_{tw} is the complete history of w through to t. Here, given that information about the laws is admissible, T_w can serve as X, since T_w specifies the chance of F; and here, given that information about the past is admissible, H_{tw} can serve as E. Then $C(F/T_wH_{tw}) > 0$. But since the truth of F would entail the falsity of the actual laws, F and T_w are incompatible, so $C(F/T_wH_{tw})=0$.—Contradiction.[6]

This argument *seems* to suggest that HS and the PP are incompatible. Thus Jenann Ismael maintains: '[O]n any view according to which a statement about the t-chances places some constraint—no matter how weak—on post-t history, the statement will be undermined by a history which violates the constraint.' She continues: 'To disallow undermining altogether, an account of chance must completely sever the link with actual history; worlds with identical histories must sometimes differ with respect to the chances' ([1996], pp. 82–3).

But the PP may at least be approximated. As Lewis, drawing on Michael Thau ([1994]) and Ned Hall ([1994]), came to see, one can define a New Principle (NP): $C(A/\ T_wH_{tw})=P_{tw}(A/T_w)$. L-chance is said to fit the NP. The contradiction is averted: since F and T are incompatible, $P_{tw}(F/T_w)= 0=C(F/T_wH_{tw})$. The PP is approximated: to the extent A concerns a relatively small sample of history, or a relatively predictable future, the NP and the PP converge (Lewis [1994], pp. 486–9). So one can still apply the PP for most

[6] Some recent commentators have questioned the incompatibility of F and T_w. Peter Vranas ([1998]) argues that, since Lewis regards HS as contingent, he need not regard F and T_w as incompatible, only as false in the 'inner sphere' through which HS holds. But this is a misinterpretation of Lewis. For Lewis, while HS is indeed contingent, BEST is a necessary analytic truth. And it is BEST that renders F and T_w incompatible. Barry Loewer (forthcoming) argues that F and T_w are not incompatible unless one also conditionalizes on the totality claim Q: H_{tw} and F comprise the total history of w. For without Q there might be a still further future F* which regrounds T_w by compensating for F. Yet, says Loewer, one cannot conditionalize on Q because Q is inadmissible. I reply, first, that Q is at least sometimes admissible. It might be settled in advance that a limited number of trials will occur. For example, it might be settled in advance that certain high-energy structures can only occur in the first second of a post-Big-Bang universe. I reply, second, that one need not conditionalize on Q. One may equally conjoin Q to F, and consider $C(FQ/\ T_wH_{tw})$ (in this position, admissibility does not matter). Now suppose that the world has some nonzero lawful probability of ending at any given time. Then L-ch_{tw}(FQ) is nonzero, and so $C(FQ/\ T_wH_{tw}) > 0$. But FQ and T_w are incompatible, so $C(FQ/T_wH_{tw})=0$. —Contradiction regained.

everyday reasoning tasks without going too far wrong. Lewis concludes that given that (i) N-chance violates the PP (see below), (ii) L-chance at least approximates the PP, and (iii) the PP is definitive of the chance role, L-chance should be regarded as chance.

Why think that L-chance fits the NP, though? It is one thing to block the undermining argument, but another thing to show that L-chance can explain why the NP(/PP) is a requirement of rationality. Why should information about the BEST constrain rational credence? Here Lewis claims to see 'dimly but well enough, how knowledge of frequencies and symmetries and best systems could constrain rational credence' ([1994], p. 484). I think Lewis is right, though I am afraid I do not have anything substantial to add here. Suppose the only relevant information one has is that there are ten positives in one hundred trials. What credence ought one to assign to there being a positive result on the seventeenth trial? I think the answer is .1, though I am afraid I do not have a deeper explanation to offer. But it is enough that a frequency-credence constraint be seen, however it is ultimately to be explained.

Why think that the NP quantity $P_{tw}(A/T_w)$ must be defined? There is, after all, no guarantee that L-ch(T) will be defined. Carl Hoefer expresses this worry: 'It's not clear to me, for example, that it is reasonable to require T_w to give us probabilistic laws so strong and comprehensive that they determine their own probability of truth' ([1997], p. 327).[7] But this is no real worry at all: at worlds at which L-ch(T) is defined (like, one should hope, ours), the L-chances will still be in the running for the chance role. At worlds at which L-ch(T) is undefined, Lewis may well say that there is no fact of the matter about chance; or that if there is, it is in virtue of some even less perfect realizer of the chance role.

Necessitarian chance: Necessitarian chance (N-chance) is an irreducible theoretical posit, over and above the occurrent facts. I will treat David Armstrong as an exemplary Necessitarian. For Armstrong, the law and chance facts are second-order degree-of-necessitation universals: N:*x* (F, G) (§2). For other Necessitarians, chance may be a brute necessity *in re*, or some other irreducible relation between natural properties.

Is N-chance chance? To begin with, it violates HS by construction. Continuum-many distinct worlds are countenanced, which agree on all the occurrences but diverge as to the magnitude of the chances.

N-chance does not meet the PP, but it does meet the BCP. N-chance does not meet the PP because there seems to be no relation at all between the value of N:*x* and reasonable credence, as Lewis explains:

[7] Hoefer's main argument is that L-ch(T), even if defined, is 'quite un-Humean in spirit', so that 'the quantity $P_{tw}(T_w)$ should be regarded by the Humean as an amusing bit of nonsense' ([1997], p. 327–8). I fail to see the force of this. If L-ch(T) is indeed definable by a Humean recipe, then it is available to the Humean: the intuitions of Hoefer's 'good Humean' are not sacred.

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

As Bas van Fraassen ([1989], pp. 81–6) argues, since N-chance is logically independent of actual frequency, the best the Necessitarian can do is to maintain that divergence is unlikely. But if 'likely' is meant subjectively, then the question is begged: a connection between chance and credence is assumed. Whereas if 'likely' is meant objectively (as N-chance), then the question is not even addressed: no connection between N-chance and credence is drawn.[8]

Perhaps the best the Necessitarian can do is (i) to introduce some logical connection to actual frequency via an implicit definition of N-chance, and then (ii) to parrot Lewis's claim that information about frequencies should (somehow!) constrain rational credence. For instance, N:x might be implicitly defined as 'that property which disposes the set-up to yield a unique limiting relative frequency x on infinite repetitions'. But either (i) this disposition is itself chancy, in which case it cannot be invoked in defining chance, or (ii) this disposition is non-chancy, in which case it entails that all infinite repetition worlds with the same N-chances must exhibit the same frequencies. Such a conclusion would be anathema for the Necessitarian.[9]

N-chance is capable of meeting the BCP because, since it is distinct from the occurrent facts, it is amenable to BCP-respecting recombination. That is, N:x (F, G) can coexist with any occurrent history, as long as $0<x<1$ ($x=0$ and $x=1$ may be dismissed as not genuinely chancy, except by mathematical

[8] Robert Black, defending Necessitarianism, concedes that the PP shall have to be taken as 'a primitive principle of reason' ([1998], p. 384). But here the theory has come unhinged. One might as well associate chances with any collection of numbers, say those corresponding to every third letter of the bible. These numbers have just as much claim (namely: none) to constrain rational credence. Eric Hiddleston ([2001]) suggests that N-chance can constrain rational credence, on the cheap. He maintains that (i) it is rational to believe our best theories, and (ii) to believe a theory requires fitting one's credences to the chances the theory projects. But (ii) only holds if the theory actually projects real chances, which is just the issue. Believing a theory does not require fitting one's credences to the magnitude of N:x, unless one already has reason to think that N-chance is chance.

[9] Such a conclusion, moreover, would violate the BCP. Let w be a world with infinitely many coin flips, exhibiting a .5 limiting relative frequency of heads; let G be the proposition: there are infinitely many coin flips, all of which land heads; let t be any time before the first tails landing. On one view of the probabilistic calculus, N-ch_{tw}(G) is infinitesimal. This yields a violation of the letter of the BCP: N-$ch_{tw}(G)>0$, yet there is no world w' such that (i) G is true, (ii) w' matches w up to t, and (iii) N-$ch_{tw'}$(G)=N-ch_{tw}(G). On another view of the probabilistic calculus, N-ch_{tw}(G)=0, but only '0' in the measure-theoretic sense in which it is still *possible*. This preserves the letter of the BCP, since the antecedent of the BCP is restricted to nonzero chances; but it still violates the spirit of the BCP, since, intuitively, the ban on zero chances ought not to cover *realizable* zero chances.

courtesy). Thus if N-ch_{tw}(A)=x (for x > 0), then there exists a world w' at which (i) A is true, (ii) w' matches w up to t, and (iii) N-$ch_{tw'}$(A)=x.[10]

These considerations *seem* to suggest that the PP and the BCP are incompatible: to meet the PP an account must generate chances that conform closely to the actual frequencies, whereas to meet the BCP an account must allow chances to veer widely from the actual frequencies. Were this so, then our very conception of chance would be deeply incoherent.

I conclude that L-chance meets HS, approximates the PP, and violates the BCP; and that N-chance meets the BCP, but violates the PP and HS. If these were the only options, I would call it a draw: I would give the PP and the BCP equal weight, which gives N-chance the edge; but I would give HS some weight too, enough to bring the contest back in balance. (I would, however, have no idea how to argue against someone who rated this outcome a victory for one side or the other.) I would also conclude that all our core principles about chance *seem* mutually incompatible: *our concept of chance seems deeply incoherent*.

4 Principled chance?

Is there any conception of chance that meets all our core principles? L-chance and N-chance are not the only options. I now introduce a third conception of chance, L*-chance, so named because it is derived from L-chance, as follows: L*-ch_{tw}(A)=L-ch_{tw}(A/T). In words, the L*-chance of A is equal to the L-chance of: A conditional on the theory of chance.[11]

Perhaps the best way to intuit the difference between L-chance and L*-chance is to think of the difference between *drawing with replacement* and *without* (Arntzenius and Hall, [ms]). Suppose that at t one is to make a million random draws from an urn with a million marbles, half black, half white. Suppose L-ch_{tw}(black)=.5. And suppose that nothing else occurs at w, or at least nothing else relevant to determining L-ch_{tw}(black). Then the L-chance of drawing all-black is $.5^{1000000}$, which is, intuitively speaking, the chance of drawing all-black in a drawing in which the drawn marble is replaced in the urn for the next drawing. But the L*-chance of drawing all-black is generated

[10] Actually, Armstrong himself is committed to violations of the BCP in virtue of his denial of uninstantiated universals (Armstrong [1997], pp. 38–43). Suppose that N:.5 (F, G), that F occurs just once, and that G occurs just that once. By the BCP there must be a comparable world w' in which F occurs and G does not. But Armstrong cannot allow this because G will be uninstantiated at w', and so N:.5 (F, G) cannot exist at w'. Since this problem involves (i) a fairly exotic case, and (ii) optional features of Armstrong's metaphysics, I will not pursue this complaint further.

[11] The original idea is due to Frank Arntzenius, in collaboration with Hall (Arntzenius and Hall [ms]). Arntzenius and Hall employ L*-ch to argue that the PP should not be taken as the sole chance principle. But (for reasons to be discussed in §5) they dismiss L*-chance as a technical trick which makes trouble for PP-driven conceptions of chance, rather than as a serious competitor for the chance role.

by conditionalizing on the fact that L-ch_{tw}(black)=.5, which (on the supposition that nothing else is relevant to determining L-ch_{tw}(black) at *w*) is equivalent to conditionalizing on the fact that half the draws will be black.[12] Thus the L*-chance of drawing all-black is zero, while the L*-chance of drawing half-black is one, which is, intuitively speaking, the chance that one gets if one thinks of each drawn marble as discarded without replacement.

L*-chance meets HS, the PP, and the BCP. Since L*-chance is defined from L-chance, and since L-chance meets HS, L*-chance meets HS as well. L-chance and L*-chance are different Humean recipes for determining P_{tw}(-).

L*-chance meets the PP, given that L-chance meets the NP (Lewis [1994]). Recall that according to the NP, $C(A/T_wH_{tw})=P_{tw}(A/T_w)$. So given that the NP holds for L-chance, $C(A/T_wH_{tw})$=L-$ch_{tw}(A/T_w)$=L*-ch_{tw}(A). But that is just to say that L*-chance meets the PP perfectly![13] In particular, L*-chance cannot be undermined. Suppose (for *reductio*) that F is an undermining future with respect to T_w. That is just to say that F and T_w are incompatible, in which case L-ch(F/T_w)=0, which is just to say that L*-ch(F)=0. (To put the point in terms of drawing without replacement, if one starts towards an undermining future by a run of black drawings, one is guaranteed some compensating run of whites as one reaches the bottom of the urn—there aren't enough black marbles for undermining.)

L*-chance, finally, meets the BCP, in virtue of the fact that the L*-chance of an undermining future F is 0. So suppose the L*-ch(A)=j (for $j>0$). Then L-ch(A/T)=j. Since H is admissible, L-ch(A/TH)=j. So A is compatible with T and H. So there exists a world w' at which (i) A is true, (ii) w' matches *w* up to *t*: H is true, and (iii) L*-$ch_{tw'}$(A)=j: T is true. In fact, L*-chance meets an even stronger principle than the BCP, namely that 'the complete theory of chance is not chancy' (Bigelow *et al.* [1993], p. 456): L*-ch(T)=L-ch(T/T)=1. (Intuitively, if one makes a million draws without replacement from an urn with a million marbles, one can be certain that the ratio of black-to-white drawings will equal the ratio of black-to-white marbles.)

Good news: HS, the PP, and the BCP are demonstrably compatible, which is surprising given the appearances reported in §3. *Our concept of chance* seems *coherent*. There is a perfect realizer of the chance role as defined in §2. So, assuming that there are no further core principles of chance, the case is closed: *chance is L*-chance*.

[12] More precisely, conditionalizing on the fact that L-ch_{tw}(black)=.5 is equivalent to conditionalizing on the fact that *roughly* half the draws will be black, where 'roughly' means: 'close enough that the constraint of simplicity will lead to rounding'. If 500,001 draws are blacks, then that should suffice for L-ch_{tw}(black)=.5, while if 600,000 draws are blacks, then not (see Lewis [1994], p. 481). There need be no fact of the matter as to where the 'roughly' line is drawn.

[13] By Lewis's own lights, this suffices to show that L*-chance is chance.

5 Unprincipled chance?

The good news so far is that there is a conception of chance, L*-chance, that meets HS, the PP, and the BCP. Here comes the bad news.

Bad news: L*-chance is both unlawful and unstable. L*-chance is unlawful because L*-ch(A) is defined in terms of L-ch(A/T), where T is the theory of chance generated by the BEST. So the laws will still entail history-to-L-chance conditionals of the form: if the history through t is H, then L-ch$_{tw}$(A)=x. But L*-ch$_{tw}$(A) will generally not be equal to x, since L*-ch(A) will generally not be equal to L-ch(A). It is L-ch(-) rather than L*-ch(-) which appears in the laws of nature. No contradiction here: the history to L-chance conditionals need not be interpreted as being about real chances. L-ch(-) might be left uninterpreted, or interpreted as something like 'the systematic frequency of -'. Rather the problem is that L*-ch(-) and lawhood have disconnected.

The second piece of bad news is that L*-chance is unstable. Leaving out the special case of an infinite history, L-ch(T) can be expected to increase with time: as more of the mosaic unfolds, the picture grows clearer. This renders the L*-chances unstable over time: the chance that an earlier flip lands heads will generally differ from the chance that a later flip lands heads.[14] Probabilistic independence fails: the L*-chance of the next landing will be influenced by the outcomes of the previous landings. If one thinks in terms of drawing without replacement, then a run of black drawings will entail a greater chance of a white drawing.[15]

How bad is the bad news? In particular, is it enough to disqualify L*-chance from serving as chance? More precisely, should one (i) add further principles to the chance role which connect the concept of chance to lawhood and to stability, and (ii) rank-order these further principles above (1) the PP, (2) the BCP, and (3) HS, in such a way that L*-chance will no longer outstrip L-chance and/or N-chance, which have no analogous problems with lawhood or stability?

Lawhood: Can one formulate a principle connecting chance and lawhood that (i) L*-chance violates, and (ii) deserves to be accorded high rank in defining the chance role? Consider the *Lawful Magnitude Principle* (LMP): If

[14] This is why Arntzenius and Hall reject L*-chance: '[*R*]*eal* chances don't vary in this way—or if they do, it's not up to the philosopher to say so. Any philosophical position that says that [...] the chances *must* vary in this way is unacceptable' ([ms]).

[15] Here the metaphor of drawing without replacement is a bit misleading, in that it makes it seem as if L*-chance licenses the Gambler's Fallacy, by making a white drawing cause the next drawing to be biased to black. But really the Gambler's Fallacy presupposes that the trials in play are stable and independent. What L*-chance actually entails is that there are subtle constitutional dependencies between most trials. So it is not as if the one outcome 'causes' a bias in the next drawing; rather it is that the chances are *constituted* by the global distributions of occurrences in a way that adjusts for local deviations from global frequencies.

$P_{tw}(A)=x$, then there is a lawfully entailed history-to-chance conditional of the form: if the history through t is H, then $ch_{tw}(A)=x$. The LMP connects the magnitude of chance to the magnitude of a lawfully projected quantity. L*-chance does not meet the LMP, because (i) L*-chance maintains the BEST, and (ii) the magnitude of L^*_{tw}(-) will generally differ from the magnitude L_{tw}(-) projected by the BEST.

But the LMP is at least approximated. Given that the lawfully projected magnitude is L-ch(A), and given that L-ch(A/T) approximates L-ch(A), it follows that L*-ch(A) approximates the lawful magnitude (just as the NP approximates the PP). To the extent A concerns a relatively small sample from the middle of history, or a relatively predictable future, the L*-chances and the lawful magnitudes converge. In particular, all lawful predictions concerning localized experiments will go through.

Moreover, other connections between chance and lawhood are sustained. First, while the L*-chances do not themselves appear in the laws, it is not as if they are completely independent of the lawful magnitudes: the L*-chances are derived directly from the lawful magnitudes. Second, L*-chance meets the weaker *Lawful Existence Principle* (LEP): If the laws of w are purely deterministic, then there are no chances at w (except for the degenerate cases of 0 and 1).[16] While the LMP connects the magnitude of chance to the magnitude of the lawfully projected quantity, the LEP connects the *existence* of chances to the *existence* of lawfully projected quantities. L*-chance meets the LEP, because if there are no L-chances then there will be no L*-chances by construction.

So the unlawfulness of L*-chance is not such bad news. Or, it would be only if (i) one insisted on the LMP over and above the LEP, and (ii) one insisted that the LMP had to be met exactly rather than approximated. For whatever it is worth, I see little justification for being so insistent here.

Stability: Can one formulate a principle connecting chance with stability that (i) is violated by L*-chance, and (ii) deserves to be accorded first rank in defining the chance role? Consider the *Stable Trial Principle* (STP): If (i) A concerns the outcome of an experimental setup E at t, and (ii) B concerns the same outcome of a perfect repetition of E at a later time t', then $P_{tw}(A)=x=P_{t'w}(B)$.[17] The STP predicts, for instance, that if one repeats a coin flip, the chance of heads should be the same on both trials.

But once again, the STP is approximated. Given that A and B are repeat trials, L-ch(A) and L-ch(B) will meet the STP, from which it follows

[16] An arguable exception: there might also be chance distributions over initial conditions even given a deterministic dynamics (see Barry Loewer [2001]). L*-chance is properly neutral here.

[17] Complication: the STP needs to be supplemented with something like a 'causal independence of trials' proviso. But the notion of causality is not available here, given that causality presupposes the notion of probability-raising. I shall ignore this complication in the main text.

that L*-ch(A) and L*-ch(B) will approximate the STP (just as the NP approximates the PP). To the extent that A and B concern a relatively small sample from the middle of history, or a relatively predictable result, L*-chance converges on stability. In particular, all localized trials will appear stable and independent.

In fact, part of the plausibility of the STP can be explained away. It turns out that L*-chance actually *entails* that one should set one's credences in accord with the STP, at least in the infinite case. This is because (i) L*-chance is *exchangeable*, and (ii) exchangeable distributions over infinite sequences are provably stable. A definition of chance satisfies exchangeability iff it assigns equal chances to every permutation of a sequence—exchangeability is order-invariance. So, for instance, if the definition assigns a 1/16 chance to the marble-drawing sequence BWWW, it should assign a 1/16 chance to all its permutations: WBWW, WWBW, WWWB. Since L-chance (and in particular L-ch(A/T)) is keyed into frequencies, and since frequencies are order-invariant (all permutations of a sequence are by definition frequency-preserving), it follows that L*-chance is order-invariant. And, as Bruno de Finetti proved, an exchangeable distribution over an infinite sequence is stable (see Brian Skyrms [1980], pp. 158–60).[18]

So the instability of L*-chances is not such bad news either. Or, it would be only if one insisted on exact stability for the finite case. For whatever it is worth, I still find such insistence very plausible, but not any more plausible than the PP, or the BCP.

So is L*-chance chance? The failure of stability, small and limited though it may be, still seems hard to swallow. But the alternatives (L-chance and N-chance) seem even worse. *If* these were the only options, I would conclude, with deep reservations, that L*-chance is indeed chance: it is the least imperfect of the lot.

6 Principled chances

So far I have considered three definitions of chance: L-chance, N-chance, and L*-chance. I have assessed these in respect to five main platitudes: (1) the PP, (2) the BCP, (3) HS, (4) the LMP, and (5) the STP. L-chance meets (3), (4), and (5), and approximates (1). N-chance meets (2), (4), and (5). L*-chance meets (1), (2), and (3), and approximates (4) and (5). So far it *seems* as if L*-chance is the best of a flawed lot.

[18] As Skyrms explains, De Finetti exploited this proof to defend subjectivism, on the grounds that the subjectivist whose allotment of credences respects exchangeability will behave (at least in the infinite case) *as if* she believes in genuine, stable objective chances. L*-chance does as well as subjective-'chance' in the infinite case, and better in the finite case, where L*-chance at least approximates stability.

So is L*-chance chance? *Maybe:* I think that the most principled approach of all is to regard 'chance' as *vague*, such that L-chance and L*-chance are admissible precisifications. It turns out that (i) L-chance and L*-chance fit the model of semantic indecision concerning ontological multiplicity, (ii) by treating 'chance' as vague between L-chance and L*-chance, one can explain how all of platitudes (1)–(5) are met, and (iii) vagueness about 'chance' should not trouble the analysis of 'causation' or lead to any other obvious problems.

L-chance and L*-chance fit the model of semantic indecision concerning ontological multiplicity, in that they constitute 'many but almost one' chance functions.[19] In saying that they are 'many', I mean that (i) both L-ch_{tw}(-) and L*-ch_{tw}(-) exist, since frequencies exist and since both functions are defined from the frequencies;[20] and (ii) L-$ch_{tw} \neq$ (-)L*-ch_{tw}(-). In saying that they are 'almost one', I mean that (i) they are approximately equal: L-ch_{tw}(-) $\approx$ L*-ch_{tw}(-), especially in the most salient case of local application to controlled experiment; and (ii) they are ontologically co-based, in that both supervene on the frequencies. L-chance and L*-chance are just barely different ways of applying frequency data to the local case. (This 'almost one' condition would fail, for instance, if N-chance were added to the mix.)

By treating 'chance' as vague over L-chance and L*-chance, one can explain how all of platitudes (1)–(5) are met. Chances meet HS (this is true on either L-chance or L*-chance); chances meet the PP and the BCP (this is true on the accommodating interpretation of 'chance' as L*-chance); and chances meet the LMP and the STP (this is true on the accommodating interpretation of 'chance' as L-chance). So while it turns out that no precise definition of 'chance' meets all of (1)–(5), it also turns out that there are 'many but almost one' chance functions which jointly meet (1)–(5).

Are there any disadvantages to treating 'chance' as vague? The obvious objection is that to do so would render all causal facts vague, which would in turn render all causally loaded facts (including, plausibly, facts concerning perception, agency, and reference) vague. But it is unclear whether this is in fact objectionable, especially given that L-chance and L*-chance are 'almost one': it is not obvious that there will be any straightforward cases in which these chance functions deliver different causal verdicts. And indeed virtually all Humeans are already committed to vagueness in causality: for regularity theorists, it may be a vague affair whether two events are of a type; and for

[19] This model of vagueness is defended in Lewis ([1993]). It is of course contentious (see Timothy Williamson ([1994])), though for present purposes I need only assume that the model is coherent, whether or not it can resolve the problem of the many, or the Sorites.

[20] Even the Necessitarian accepts that both L-ch_{tw}(-) and L*-ch_{tw}(-) exist; she merely denies that either is worthy of the name 'chance'.

counterfactual theorists, it may be a vague affair as to which are the near ~O(*c*)-worlds.[21]

In any case, whether L*-chance is chance, or whether L-chance and L*-chance both constitute admissible chance functions, Humean Supervenience is upheld.

Acknowledgements

Thanks to Phil Bricker, David Lewis, Barry Loewer, Michael Strevens, Brian Weatherson, and especially to Frank Arntzenius and Ned Hall.

Department of Philosophy
University of Massachusetts at Amherst
352 Bartlett Hall
Amherst, MA 01003
USA
schaffer@philos.umass.edu

References

Armstrong, D. [1983]: *What is a Law of Nature?*, Cambridge: Cambridge University Press.

Armstrong, D. [1997]: *A World of States of Affairs*, Cambridge: Cambridge University Press.

Arntzenius, F. [manuscript]: 'Subjectivism About Chances'.

Arntzenius, F. and Hall, N. [manuscript]: 'On What We Know about Chance'.

Bigelow, J., Collins, J. and Pargetter, R. [1993]: 'The Big Bad Bug: What are the Humean's Chances?', *The British Journal for the Philosophy of Science*, **44**, pp. 443–62.

Black, R. [1998]: 'Chance, Credence, and the Principal Principle', *The British Journal for the Philosophy of Science*, **49**, pp. 371–85.

Hall, N. [1994]: 'Correcting the Guide to Objective Chance', *Mind*, **103**, pp. 504–18.

Hiddleston, E. [2001]: 'Humean Supervenience, Chance, and Magic', presented at *Metaphysical Mayhem VI*, Syracuse, NY.

Hoefer, C. [1997]: 'On Lewis's Objective Chance: "Humean Supervenience Debugged"', *Mind*, **106**, pp. 321–34.

Ismael, J. [1996]: 'What Chances Could Not Be', *The British Journal for the Philosophy of Science*, **47**, pp. 79–91.

[21] For Lewis, causation requires probability-raising 'by a large factor' ([1986], pp. 176–7; Lewis's new influence theory ([2000], p. 182) is silent here), which is a further aspect of vagueness. What is interesting about this aspect of vagueness is that it would more or less wash out the minor differences between L-chance and L*-chance in assessing causality. So, for Lewis, treating 'chance' as vague might not engender *any* further vagueness in causality.

Lewis, D. [1980]: 'A Subjectivist's Guide to Objective Chance', in R. Jeffrey (*ed.*), *Studies in Inductive Logic and Probability*, Vol. 2, Berkeley: University of California Press, pp. 263–93.

Lewis, D. [1986]: *Philosophical Papers*, Vol. 2, Oxford: Oxford University Press.

Lewis, D. [1993]: 'Many, but Almost One', in K. Campbell, J. Bacon and L. Reinhardt (*eds*), *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong*, Cambridge: Cambridge University Press, pp. 23–38.

Lewis, D. [1994]: 'Humean Supervenience Debugged', *Mind*, **103**, pp. 473–90.

Lewis, D. [2000]: 'Causation as Influence', *The Journal of Philosophy*, **98**, pp. 182–97.

Loewer, B. [2001]: 'Determinism and Chance', *Studies in the History and Philosophy of Modern Physics*, **32**, pp. 609–20.

Loewer, B. [forthcoming]: 'David Lewis's Humean Theory of Objective Chance', *Philosophy of Science*.

Shoemaker, S. [1980]: 'Causality and Properties', in P. Van Inwagen (*ed.*), *Time and Cause*, Dordrecht: D. Reidel.

Skyrms, B. [1980]: *Causal Necessity*, New Haven: Yale University Press.

Skyrms, B. [1984]: *Pragmatics and Empiricism*, New Haven: Yale University Press.

Thau, M. [1994]: 'Undermining and Admissibility', *Mind*, **103**, pp. 491–503.

van Fraassen, B. [1989]: *Laws and Symmetry*, Oxford: Oxford University Press.

Vranas, P. [1998]: 'Who's Afraid of Undermining? Why the Principal Principle need not Contradict Humean Supervenience', presented at *Philosophy of Science Association 16*, Kansas City, MO.

Williamson, T. [1994]: *Vagueness*, London: Routledge.