

Fisher, Justin. (2006). On Higher-Order and Free-Floating Chances. *The British Journal for the Philosophy of Science* 57: 691-707.

On Higher-Order and Free-Floating Chances

Justin C. Fisher – Southern Methodist University

Abstract

Marc Lange ([forthcoming]) considers what I call *free-floating chances*—objective chances that obtain at a given time despite the fact that their values are not determined by the laws of nature together with the full history of non-chancy facts up to that time. I offer an intuitive example of this phenomenon, and use it to argue that free-floating chances are indeed possible. Their possibility violates three quite widely held principles about chances: the lawful magnitude principle, the principle that chances evolve by conditionalization, and a version of David Lewis’ principal principle. I argue that we should reject common formulations of each of these principles, though I offer revised understandings of each which retain much of the intuitive attractiveness of the originals and are consistent with the possibility of free-floating chances. I conclude by arguing that, while considerations of free-floating chances are important, they won’t sustain the extravagant conclusions Lange attempts to draw from them.

- 1 *Introduction*
 - 2 *First- and Higher-Order Chances*
 - 3 *Free-Floating Chances*
 - 4 *Support for the Intuitive Assessment*
 - 5 *Three Principles Violated*
 - 6 *What to do?*
 - 7 *COND as a Default Hypothesis*
 - 8 *A More Principled Principal Principle*
 - 9 *Conclusion*
-

1 Introduction

Marc Lange ([forthcoming]) considers what I call *free-floating chances*—objective chances that obtain at a given time despite the fact that their values are not determined by the laws of nature together with the full history of non-chancy occurrences in the world up to that time. I agree with Lange that free-floating chances are indeed possible. My first task in this paper is to explain why. I think it is easiest to see the possibility of free-floating chances by considering them as the limiting case of a more general phenomenon involving higher-order chances. I offer an intuitive example exhibiting the more general phenomenon in section 2, and then describe the limiting case in section 3. In section 4, I use this intuitive example to argue that free-floating chances are indeed possible.

In section 5, I explain how the possibility of free-floating chances is inconsistent with three fairly widely accepted principles about chances: the lawful magnitude principle, the

principle that chances evolve by conditionalization, and a version of David Lewis' principal principle. In sections 6-8, I argue that we should reject the necessary truth of common formulations of each of these principles, though I offer revised understandings of these principles which retain much of the intuitive attractiveness of the originals and are consistent with the possibility of free-floating chances.

These findings embody important insights about the nature of objective chances. However, I will close (in section 9) with a cautionary note against Lange's attempts to move from these considerations to more extravagant conclusions about the relations between objective chances, objective probabilities, laws, and categorical (or non-chancy) properties. Considering free-floating chances does not give us new reasons to stake particular positions on these issues, but it will help us to see these issues more clearly as we continue to debate them.

2 First- and Higher-Order Chances

Some free-floating chances may be seen as a limiting case of a more general phenomenon. In this section, I offer an intuitive example of such a phenomenon.

My example is somewhat fanciful, involving spell-casting and genies coming out of bottles. For those opposed to fanciful cases, it would be straightforward to translate my example into a structurally identical case involving fundamental physical particles and fundamental physical properties in a possible world with indeterministic natural laws perhaps not all that different from our own. However, these issues are complicated enough as it is, and I think we will do well to consider a case involving objects and events that are easy to visualize and keep straight in our minds.

At time t_0 Merlin casts a risky enchantment on a bottle. Even though he does everything exactly right, there is a 50% chance that his enchantment will turn out to have been a dark one and a 50% chance that it will turn out to have been a light one. If the enchantment turns out to have been a light one, then by time t_1 the bottle will have turned a slightly milky color (just how milky depends on how much flourish Merlin puts on his




Figure 1. Merlin's Flourish-full Enchantment.

wand in casting the spell) and there is a 40% chance that at time t_2 a white genie will emerge from the bottle and a 60% chance that a mere puff of white smoke will emerge instead. If, instead, the enchantment turns out to have been a dark one, then by t_1 the bottle will have turned a slightly murky color (how murky depends upon Merlin's flourish), and there will be a 60% chance that at t_2 a black genie will emerge from the bottle and a 40% chance that a mere puff of black smoke will emerge instead. (See figure 1.)

Following Lange ([forthcoming]), let $\text{Ch}_{t_1}(A_{t_2})$ denote the chance at time t_1 of event A happening at time t_2 . A *first-order chance* is a chance of an event that does not itself involve any chances of further events. Depending on whether Merlin's enchantment turns out to be light or dark, the first-order chance at t_1 of the event G_{t_2} of a genie emerging from the bottle at t_2 will be either 40% or 60%.

If Merlin's enchantment is light then $\text{Ch}_{t_1}(G_{t_2}) = 40\%$.

If Merlin's enchantment is dark then $\text{Ch}_{t_1}(G_{t_2}) = 60\%$.

A *second-order chance* is a chance of some event's occurring which itself involves a first-order chance of some other event occurring. Since the chanciness of Merlin's enchantment serves to determine the subsequent chances of a genie emerging, this case involves second-order chances.

$\text{Ch}_{t_0} [\text{Ch}_{t_1}(G_{t_2}) = 40\%] = 50\%$.

$\text{Ch}_{t_0} [\text{Ch}_{t_1}(G_{t_2}) = 60\%] = 50\%$.

There are a number of interesting puzzles involving the nature of higher-order chances and their relations to lower-order chances, laws, and non-chancy properties. But, perhaps the most interesting puzzles arise when we consider a limiting case of this example.

3 Free-Floating Chances

Recall that when Merlin casts what turns out to be a light enchantment, the bottle turns a milky color at time t_1 , where the amount of milkiess depends upon the amount of flourish Merlin gives his wand in casting the enchantment (and similarly for the murkiess of the bottle when he casts what turns out to be a dark enchantment). Let us suppose that the amount of flourish is directly proportional to the amount of milkiess (or murkiess) that the enchantment produces. And, just to be clear, let us suppose that differences in flourish have no effect whatsoever upon the first-order chances at t_1 of getting a genie or smoke from the bottle at t_2 .

Now, the puzzling case, which shares the structure of a case proposed in abstract terms by Lange, ([forthcoming]): Suppose Merlin learns to cast the enchantment without any flourish at all, so that, at time t_1 , the bottle will not have changed color to reveal

whether the enchantment will turn out to have been light or dark. Indeed, let us suppose that up to and including time t_1 *all categorical* (or non-chancy) *properties* would be exactly the same if the enchantment turns out to be light as they would be if it turns out to be dark. But, let us suppose, repeated experimentation with flourish-less enchantments yields the same frequencies of outcomes at t_2 as with flourish-full enchantments—20% white genies, 30% white smoke, 30% black genies, and 20% black smoke.

At first blush, it seems reasonable to say that the flourish-less case would be (or at least *could be*) just like more flourish-full cases in all respects except for the coloration of the bottle. In particular, it seems reasonable to say that, in any of these cases, at time t_1 the first-order chance of a genie emerging at t_2 is either 40% or 60%, depending upon whether Merlin's enchantment turns out to have been light or dark—i.e., depending upon the 50-50 second-order chancy event at t_0 . In the flourish-full cases, the color of the bottle at t_1 serves to indicate what this first-order chance at t_1 is, while in the flourish-less case there is not yet *any* categorical indication at t_1 of what this first-order chance is.¹ In the flourish-full cases, the chance at t_1 of a genie appearing at t_2 is tethered to the color of the bottle, while in flourish-less case, this chance *floats free* from the color of the bottle, and from all other categorical facts obtaining at or before t_1 .

4 Support for the Intuitive Assessment

One might find some support for this intuitive assessment and its attribution of free-floating chances by attending to the spirit of David Lewis' ([1994]) best-systems account of chances. Lewis holds that there is an important relation between chances and systematic descriptions of the distribution of categorical (non-chancy) properties throughout the history and future of a world. In particular, he thinks that the chances that obtain in a world are precisely those chances that would be posited in the systematic description that strikes the best balance between *strength* (saying accurate things about lots of phenomena), *fit* (attributing chances that differ not too much from actual frequencies) and *simplicity*.

Lewis doesn't consider the possibility of second-order chances, but they fit nicely with the spirit of his view. Take, for example, Merlin's world. One might offer a systematic description just in terms of first-order chances.

When a flourish-full enchantment is cast, there is a 50% chance of the bottle becoming milkier (proportional to the amount of flourish), and a 50% chance of its becoming murkier (proportional to the amount of flourish) instead. There is a 40% chance of a milky bottle yielding a white genie, and a 60% chance of it yielding white smoke instead. And there is a 60% chance of a murky bottle yielding a black genie, and a 40% chance of it yielding black smoke instead. So much for flourish-full enchantments. When a *flourish-less* enchantment is cast,

¹ It is quite natural to phrase this in terms of 'indication' or 'tell-tale signs', but it is important to resist the epistemic reading that these terms might suggest. What's important is not just that *agents like us* wouldn't be able to find evidence at or prior to t_1 indicating the color of Merlin's flourish-less enchantment. What's important is that *there is* no such evidence – *not even God* could find categorical facts at or prior to t_1 that would indicate whether Merlin's enchantment was light or dark. (This leaves open whether there might be *irreducibly chancy* properties at or prior to t_1 that would indicate which sort of enchantment it was, or whether even God would have to wait for subsequent categorical events to reveal this.)

the bottle will surely remain its original color, and then there is a 20% chance of a white genie, a 30% chance of white smoke, a 30% chance of a black genie, and a 20% chance of black smoke.

Now contrast this against a systematic description employing second-order and free-floating chances.

When an enchantment is cast, there is a 50% chance that *both* the bottle will be made milkier (proportional to the amount of flourish) *and* there will then be a 40% chance of a white genie and a 60% chance of white smoke; and there is a 50% chance that instead *both* the bottle will be made murkier (proportional to the amount of flourish) *and* there will then be a 60% chance of a black genie and a 40% chance of black smoke.

These descriptions have equal *strength* and *fit*, but the latter is much *simpler* (shorter) than the former. This is because the latter subsumes the similarities between flourish-less and flourish-full cases under a single generalization about second-order chances, while the former has to treat flourish-less cases quite differently from flourish-full cases. For anyone who thinks (with Lewis) that there is a close link between chances and simple descriptions, this gain in simplicity should weigh in favor of our intuitive attribution of free-floating chances in the case of Merlin's flourish-less enchantment.

One may accept the preceding reasoning while remaining neutral about Lewis' metaphysical claim that for certain chances to obtain *just is* for a certain sort of systematic description of categorical properties to be best. E.g., one might hold instead that the link between simplicity and chances is merely epistemic—that the most plausible explanation for observed similarities between flourish-less and flourish-full cases is the simplest one, and, as we have seen, this simplest explanation is the one that invokes free-floating chances. This much-more-epistemic argument would also lend support to our intuitive assessment above.

Perhaps the strongest intuitive reason for doubting the possibility of free-floating chances is the worry that there are too few constraints upon our positing them. We are now in a position to answer this concern.

First, we should note that, in cases like Merlin's, chances float free only *temporarily*—their exact values will be pinned down again by *subsequent* categorical events, e.g., the emergence at t_2 of something white or something black from the bottle. Indeed, *some* such chances' eventually being pinned down by *some* categorical events *sometimes* is a pre-requisite for their fitting the spirit of Lewis' best-system account. Granted, by definition, a free-floating chance at a given time is not determined by the laws and categorical history of the world *up to that time*. But, when we take a longer view—one that considers the future as well—free-floating chances may be fully evidenced by empirical categorical evidence.²

Second, Occam's razor cuts both ways—while it sometimes calls upon us to posit free-floating chances, it also forbids positing free-floating chances except in cases where doing so affords a gain in simplicity. So long as we're wielding Occam's razor, we may

² Since the commonly-espoused dictum of *Humean supervenience* takes this longer view, free-floating chances are compatible with Humean supervenience.

admit that we should sometimes posit free-floating chances, without thereby committing ourselves to positing them willy-nilly. If one accepts a Lewis-style view that the chances just are the ones that would be posited by the best system, this settles the matter—there *will be* no free-floating chances except in the relatively few cases where simplicity demands positing them. If one rejects Lewis' view, and holds that chances are irreducible and *sui generis*, then one may need to admit the possibility of free-floating chances (unlike those in Merlin's world) that will *never* be empirically revealed. But still, one should no more believe in the presence of such ephemeral chances than one should believe in the presence of undetectable goblins.

Either way, we can and should endorse the metaphysical possibility of free-floating chances, while still demanding a great deal of empirical evidence before we will become convinced that any such chances *actually obtain* here in our world. Like Lange, I take it to be an open empirical question whether such evidence will be forthcoming.

5 Three Principles Violated

Despite its intuitive attractiveness and relative simplicity, our countenancing the possibility of free-floating chances raises several difficult puzzles. In particular, it violates three fairly widely accepted principles about chances. Lange ([forthcoming]) notes the first two of these violations, but not the third. Free-floating chances violate each of these principles for essentially the same reason, so I will first state the three principles, then point out what they have in common, and then explain why that common feature is inconsistent with free-floating chances.

First, our intuitive assessment violates what Schaffer ([2003], pp. 36-7) calls *the lawful magnitude principle* (hereafter, LMP).

LMP: All chances at a given time are determined by the laws of nature together with the complete history of the world's categorical (i.e., non-chancy) properties up to and including that time.

LMP is attractive in part because it offers a straight-forward picture of the dynamic relationship between laws, chances, and categorical properties—although the world may be chancy, LMP insists that these chances be closely tethered to changing categorical conditions by the natural laws that determine the chances of further categorical developments.

Second, our intuitive assessment violates the principle that *chances evolve by conditionalization* (hereafter COND). COND holds that 'a later chance distribution comes from an earlier one by conditionalizing on the complete history of the interval in between' (Lewis 1986, pg. 201). We may put this more formally:

COND: For any times t_2 following t_1 following t_0 and for any event A, if H describes all categorical changes that occur from t_0 up to and including t_1 , then

$$\text{Ch}_{t_1}(A_{t_2}) = \text{Ch}_{t_0}(A_{t_2} \mid H).$$

Third, our intuitive assessment violates the principle that David Lewis ([1986]) takes to be so central to our pre-theoretic notion of chance that he calls it the *Principal*

Principle. This principle holds that if there is a chance at t_1 of future event G_{t_2} , this chance must be equal to the rational credence (or degree of belief) that a *suitably well-informed agent* would give to G_{t_2} . In Lewis' original formulation, a *suitably well-informed agent* would be one who knows all the chances at t_1 and knows no 'inadmissible' information about events after that.

PP-ORIGINAL: If there is a chance at t_1 of future event G_{t_2} , this chance must be equal to the rational credence that an agent would give to G_{t_2} if she knew all the chances at t_1 and has no 'inadmissible' information about events after that.

Our intuitive assessment does not violate PP-ORIGINAL. It does, however, violate what Lewis ([1986], [1994]) calls 'the Principal Principle reformulated,' which instead defines a *suitably well-informed agent* as one whose empirical knowledge includes all and only the laws of nature and the full categorical history leading up to and including t_1 .³

PP: If there is a chance at t_1 of future event G_{t_2} , this chance must be equal to the rational credence that an agent would give to G_{t_2} if she knew only the laws of nature and the categorical history up to and including t_1 .

Since PP explicitly (albeit indirectly) links chances to laws and categorical properties, it places more explicit constraints on our theory of chance than did PP-ORIGINAL. This makes PP somewhat more interesting than the PP-ORIGINAL. PP is also interesting because Lewis ([1994]) went to pains to defend its (approximate) truth.

When it is determinate what credences a suitably well-informed agent rationally should assign, PP demands that, if there are corresponding chances, then they must equal these credences. In the case of Merlin's flourish-less enchantment, it is pretty clear that, for an agent who knows only the natural laws and the categorical history leading up to and including t_1 , the rational credence in the claim that a genie will emerge at t_2 should be 50%.⁴ Hence, PP demands that if there is a $Ch_{t_1}(G_{t_2})$, then this chance must also be 50%.

LMP, COND, and PP are alike in that they demand a tight tethering between chances and the sequence of categorical properties leading up to them. Given a full set of prior chances, a full set of natural laws, and a categorical history leading up to and including t_1 , LMP, COND and PP all hold that the chances at t_1 must be uniquely determined.⁵

³ In thinking that this 'reformulation' is equivalent to PP-ORIGINAL, Lewis effectively presumes that LMP is a truth knowable to any rational agent, and hence that an agent who is 'suitably well-informed' in the reformulated sense could use this information about laws and categorical history to derive facts about chances, thereby making herself 'suitably well-informed' in the original sense as well. Since the 'reformulation' effectively presupposes LMP, it shouldn't be surprising that our intuitive assessment of the Merlin case violates the reformulation for much the same reasons that it violates LMP.

⁴ Or at least this credence matches the expected frequency of G in such circumstances, and there is no plausible argument for any particular alternative credence.

⁵ This may be less than perfectly clear in the case of PP, in large part because it may be less than perfectly clear what credences a rational agent should have given various sorts of evidence. But, for reasons noted above, it is quite clear what PP demands in the case of Merlin's flourish-less enchantment. So even if PP places less robust demands for determinacy of chances than LMP or COND, it still places enough demands to run counter to our intuitive assessment of Merlin's flourish-less enchantment.

Contrary to this, our intuitive assessment of Merlin's flourish-less enchantment held that the chance at t_1 of a genie appearing at t_2 is *not* fully determined by preceding factors—this chance might be either 40% or 60%, even though no categorical properties at or prior to t_1 serve to determine which. Our intuitive assessment countenances chances at t_1 that *float free* of any categorical changes that have happened by that time—a possibility that is ruled out by LMP, COND, and PP.

6 What to do?

We have seen intuitive and theoretical reasons for thinking that, in cases like Merlin's flourish-less enchantment, first-order chances might temporarily float free from categorical properties. But we've seen that this intuitive assessment violates three fairly widely accepted principles. What to do?

I think it's easiest to see what to do in the case of LMP. I grant that it might be *nice* if the world displayed the neat dynamical relationship between laws, first-order chances, and categorical properties that LMP describes, and that, so far as we currently know, we may be lucky enough to live in a nice world like this. Still, it is *at least coherent* to suppose that in some possible worlds—worlds like Merlin's—the relationship between laws, chances, and categorical properties isn't quite so straightforward as LMP suggests. It would take something more than an intuitive attraction to the simplicity of LMP to convince us that these apparently coherent cases *aren't even metaphysically possible*.

Barring further reasons for thinking these violations of LMP *aren't* possible, we should conclude (with Lange, forthcoming) that LMP is not a necessary truth about all metaphysically possible worlds. It may turn out that LMP is a contingent truth about our world, and (perhaps for reasons of parsimony) we may wish to retain LMP as a working hypothesis about our world. But empirical findings (akin to the ones that Merlin might make) might conceivably force us to reject this hypothesis. And, in thinking about the metaphysics of chances, we should not mistake this contingent empirical hypothesis for a necessary truth.

COND and PP are more challenging, perhaps because they have stronger intuitive ties than LMP to the essential nature of chances. I will consider them carefully in the coming sections.

7 COND as a Default Hypothesis

If we accept the intuitive assessment proposed above, by t_1 the chance of G (i.e., of a genie coming out of the bottle at t_2) is either 40% or 60%. This chance can't have been brought about by starting with whatever the chance of G was at t_0 and conditionalizing upon the intervening categorical events, for these events would be the same in the 40% case as they would be in the 60% case. Hence we need to allow that this chance came about in some other way than conditionalization. So we must accept some limitations to the generality of COND. But which limitations?

Lange ([forthcoming]) suggests that we deny that there was any chance of G prior to the time when the 40% or 60% chance *pops into existence*. This proposal allows us to hold that, once chances pop into existence, they always evolve by COND. However, this

proposal comes with many costs, not least of which is the strangeness of chances popping into existence.

Lange is suggesting that we should be *somewhat sparse* in our positing of chances. My own sympathies go in two directions regarding the sparseness of chances. On one hand, I am sympathetic to the thought that chances are *extremely sparse*—that the only *real* chances in our world are chances that obtain at the microphysical level involving basic physical properties and very short periods of time. This view might be motivated by the intuition that it is these chances, and no others, that a god would have to build into a world to get a world that would behave like ours.

But, on the other hand, I am often willing to go along with common usage, which allows that chances are *not at all sparse*. For *any* nomologically possible future event—no matter how far in the future it would be, and no matter whether it is given a macro- or micro-level description—there is a present chance that that event will occur. This view has a satisfying completeness about it, and it fits well with our willingness to talk about the chances of one sort of event following some other sort of event, without worrying about the micro-details of the intervening process.

Lange's position lies uncomfortably between these two intuitively tenable extremes. Lange allows that there are very many macro-level chances, and very many chances of temporally distant events. He denies only that such chances obtain in cases where there is a potential for higher-order jostling of first-order chances that would not be reflected in immediate changes in categorical properties. He posits gaps in the chances where, and only where, those gaps are needed to preserve COND as the sole way chances evolve (once they've popped into existence).

This strikes me as an *ad hoc* and unprincipled maneuver. Common usage allows that we can have good evidence that the chance at t_0 of a Merlin-style enchantment producing a genie is 50%, without having to worry about the question of whether there are immediate categorical indicators (like the changing color of the bottle) that reveal whatever fine-grained changes this chance might undergo after the enchantment is cast and before the genie would appear. Lange's proposed way of restricting COND forces us to say that it *does* matter whether there are such immediate categorical indicators along the way. It would be preferable if we could find a way of restricting COND that didn't conflict with common usage in this way, and that allowed us to retain the satisfying completeness of always positing chances for all nomologically possible eventualities.

There is a closely related objection. Lange notes that, on his proposal, objective chances cannot be objective probabilities because these chances are more sparse than objective probabilities. There is an objective probability, namely 50%, at time t_0 that a genie will emerge at t_2 , but on Lange's proposal, there is no *chance* at t_0 that a genie will emerge at t_2 —this chance won't pop into existence until t_1 . Lange touts this as a surprising consequence of his view, but one might just as easily take it to be an objection. We usually think of chances as being probabilities, and it would be preferable to find a way of restricting COND that would not force us to reject this common presumption.

I think there is a much more palatable way of restricting COND, one that does not require us to hold that chances are somewhat sparse. I think my preferred proposal amounts to the same thing as a proposal Lange considers in the final section of his paper, but I think my way of describing it is much more natural. I will first present Lange's later proposal, and then show how it may be simplified to the formulation I prefer.

Lange's later proposal holds that chances evolve by conditionalization not upon *categorical properties*, but instead upon *the outcomes of chancy processes*. In cases with no free-floating chances, all outcomes of chancy processes will be (or will be accompanied by) categorical changes, so this proposal will yield the same results as the original COND. In cases like Merlin's that do involve free-floating chances, when the chances of some future event initially float free, this must itself be the outcome of some higher-order chancy process. Hence Lange's later proposal says that the new values of these chances should be included in the body of information, conditionalization upon which would determine their new values. But this is just a fancier way of saying that when chances initially float free, *their new values will be whatever their new values are*.⁶ Since this is utterly trivial, the real work here must not be being done by the stuff about conditionalization, but instead by the preceding observation that the new values of newly-free-floating chances will have been determined stochastically by the higher-order-chancy processes that led them to float free. I propose that we just say this straight up, without introducing a potentially confusing and ultimately trivial layer of conditionalization along the way.

Let's see how this proposal applies to Merlin's case. At time t_0 , when Merlin is casting his spell, there is a 50% first-order chance that a genie will emerge from the bottle at t_2 . By time t_1 this first-order chance has changed to either 40% or 60%, as is evidenced by the milkiness or murkiness of the bottle in cases where the enchantment is flourish-full. But if the enchantment was flourish-less, then there will have been no tell-tale categorical signs by t_1 , and hence no way for this change in first-order chances to be underwritten by conditionalization. Instead, the laws of second-order chances in Merlin's world underwrite the possibility of second-order-chancy events like Merlin's flourish-less enchantment jostling first-order chances so that they float free (in stochastically determined ways) from the tethers of simple conditionalization.

More generally, I think we should not think of COND as a necessary truth, but instead as a default hypothesis about the evolution of chances in our world. In the simplest chancy worlds (perhaps including ours), first-order chances evolve in accordance with COND. But in more complicated worlds (like Merlin's) the laws governing the evolution of first-order chances will be more complicated—sometimes first-order chances will evolve in lock-step with categorical properties (as per COND), but sometimes they will (temporarily) be jostled free of categorical properties under the stochastic guidance of second-order chances. For all I've said about Merlin's world, it would be reasonable to guess that at least the *second-order* chances there evolve in accordance with COND,⁷ but we could imagine even more complicated worlds where second-order chances might also sometimes be jostled by even-higher-order chancy events.

Insofar as it's reasonable to favor simpler hypotheses, we should adopt COND as a default hypothesis about chances in our world. But possible evidence—like the evidence

⁶ I am drawing upon the common presumption that conditionalizing upon a piece of information results in that information's being incorporated as a certain truth in the resulting probability distribution. E.g., conditionalizing upon the information that $\text{Ch}^{(G_{t_2})}=40\%$ should yield a probability distribution that assigns to G_{t_2} the value 40%. This presumption holds true under Jeffrey-conditionalization, which is probably the most plausible formal framework within which one might represent Lange's later proposal.

⁷ Or at least in accordance with the modified version of COND that says that chances of a given order evolve by conditionalization upon *both* categorical changes *and* changes in lower-order chances.

Merlin might gather—could lead us to conclude that COND fails for first-order chances in our world. Such findings may force us to retreat to the fallback hypothesis that COND holds only for second-and-higher-order chances, while first-order chances sometimes evolve under the stochastic guidance of higher-order chances. And further findings might lead us to retreat to thinking that COND applies only at even-higher orders...

8 A More Principled Principal Principle

Lewis ([1994], pp. 245-46) suggests that something like PP is essentially linked to our notion of chances:

A feature of Reality deserves the name of chance to the extent that it occupies the definitive role of chance; and occupying that role means obeying [PP].

It would follow that, if we depart from PP, we'll risk losing any motivation for thinking that what we're talking about deserves the label 'chance'. Let us consider how great a departure we must make.

Above we saw that PP demanded a close correspondence between chances (where they obtain) and rational credences for suitably well-informed agents. Our intuitive assessment of Merlin's flourish-less enchantment held that chances may temporarily float free of categorical events, taking on values which will not be revealed by categorical changes until later. Since a 'suitably well-informed agent' would not yet be privy to these later categorical changes, her credences cannot be expected to track these chances. Instead, such an agent would need to set her credence at some weighted average of the possible jostlings, and await further evidence that might inform her of how exactly the chances had been jostled. So, if we accept the intuitive assessment of Merlin's flourish-less enchantment, we should expect PP to have exceptions.

One option here would be to retreat to PP-ORIGINAL, which linked chances only to the credences of an agent who is well-informed *about chances*. PP-ORIGINAL is consistent with our puzzling cases, but, unfortunately, it told us quite little about how chances are linked to credences *about laws* or *about categorical properties*. It would be nice to be able to say something substantive about this as well.

One might note that violations to ('reformulated') PP will occur only when there are higher-order-chancy events that jostle lower-order chances and make them temporarily float free of any categorical indicators of their values. So, a second option would be to adopt a weakened form of PP that remains silent about these problematic cases. This weakened version of PP would not suffer from exceptions of the sorts I've raised. And insofar as such cases are rare, this weakened PP will be correspondingly useful.

But it would be preferable to have a version of PP that could also say something positive about the relation between chances and credences in these cases, rather than remaining silent. We'd like a version of PP that spells out formally what I said in intuitive terms a moment ago—when first-order chances are jostled free of available categorical evidence, a rational agent should set her credences to an appropriately weighted average of the potential values that the chances might have been jostled to. Here is my proposal:

PP*: Given t_1 and A_{t_2} , let H be the full categorical history up to and including t_1 , and let t_α designate a time prior to t_1 and prior to any evidence-less jostlings⁸ of the chances of A_{t_2} if there is such a time.⁹ At time t_1 , a suitably well-informed agent rationally must assign credence to A_{t_2} equal to the value of $\text{Ch}_{t_\alpha}(A_{t_2} | H)$ if it has a value.^{10, 11}

PP* takes advantage of three reasonable assumptions. First, I assume that, at t_α , prior to their being jostled, the chances of A_{t_2} would take appropriately into account all potential future jostlings of them. E.g., prior to the jostling caused by Merlin's enchantment, the chances of a genie coming out of the bottle (given that he casts the enchantment) are 50%—which takes into account the potential 40% chance of a white genie and the potential 60% chance of a black genie. In intuitive terms, this first assumption tells us that appropriate relative weightings for the weighted average that we want were present in the chances at t_α .

Second, I assume that a rational agent who knew only the laws and the categorical history up to and including t_α would have credences that match the chances at t_α . This assumption should be unproblematic for anyone who accepted the original version of PP. In the Merlin case, this tells us that such an agent would have credence 50% in the claim that a genie will come out of the bottle at t_2 , given that Merlin casts the enchantment at t_0 .

The first two assumptions provide a link between chances at t_α and rational credences at t_α . Now we need a third assumption linking these to rational credences at t_1 . To this purpose, I assume that rational agents would update by conditionalization upon available evidence, even when evidence-less jostling makes chances float free of categorical changes. It is commonly held that conditionalization is the rationally appropriate way to update one's credences. In cases where there is *no* evidence-less jostling of chances, we may assume chances evolve by conditionalization (see above) so by conditionalizing her credences, a suitably well-informed agent will keep them in line with the chances. When there *is* evidence-less jostling of chances, the best a 'suitably well-informed' agent can hope to do is get the correct weighted average, and conditionalization does this by preserving the relative weightings between possible jostlings that (by our first assumption) were encoded correctly in the (conditional) chances at t_α .

⁸ Technically, this should say *less-than-fully-evidenced jostlings*, for in more complicated cases there could be jostlings that are initially evidence-less, but then some categorical evidence comes in that serves to *partially delimit* the space of possible outcomes that the jostling might have had.

⁹ If there is a beginning of time, and there were no free-floating chances at the beginning of time, then t_α could designate the beginning of time.

¹⁰ A clause should perhaps be added allowing for the assignment of credences at the beginning of time, something that PP* as stated remains silent about, for there is no t_α prior to the beginning of time. I will ignore this case in what follows.

¹¹ Lewis ([1994]) argues that, if we accept Humean Supervenience as he does, we should accept another revision to PP, one that holds that credences should reflect not *unconditional chances* but instead *chances that are conditional upon the truth of the theory T of chances that would be included in the best system describing this world*. People who are worried about these problems may adopt Lewis' solution into PP* by replacing $\text{Ch}_{t_\alpha}(A_{t_2} | H)$ with $\text{Ch}_{t_\alpha}(A_{t_2} | H \& T)$.




Figure 2. Curve B illustrates, for various times t , the rational credence that an agent informed of the laws and the categorical history up to t would have in the claim that a genie will appear at t_2 . From t_α to t_0 , as it becomes more and more probable that Merlin will cast the enchantment, B evolves in lockstep with the *chances* that a genie will appear at t_2 , depicted by both C_D and C_L . When Merlin casts his flourish-less enchantment at t_0 , these chances are jostled free of any categorical indicators of their values, going either to 60% (depicted in C_D) or 40% (depicted in C_L). Since a ‘suitably well-informed’ agent would not know at t_1 which of these jostlings occurred, she would instead conditionalize upon what categorical changes had occurred by then, yielding a weighted average of C_D and C_L .

In Merlin’s case, an agent who knew that Merlin would cast the enchantment at t_0 would already know that the bottle would stay colorless either way, so conditionalizing on this information leaves her *credence* in ${}^G_{t_2}$ unchanged at 50%, even while the *chance* of ${}^G_{t_2}$ has been jostled to either 40% or 60%. (See Figure 2.)

It is worth noting that, given two related assumptions, PP* is equivalent to PP in worlds where there is no evidence-less jostling of chances. First, we may assume that for any relevant t_1 there is a t_α .¹² Second, we may assume that in such a world chances evolve by conditionalization on categorical changes. Given these assumptions, we get

$$\text{Ch}_{t_1}({}^A_{t_2}) = \text{Ch}_{t_\alpha}({}^A_{t_2} | H).$$

PP says a suitably well-informed rational agent must conform her credence in ${}^A_{t_2}$ to the left-hand side of this equality; PP* says she must conform it to the right-hand side. Given the equality, the two are equivalent.

Since PP is equivalent to PP* when we restrict our attention to the cases that most theorists have thought about the most, and since the move from PP to PP* is intuitively well-motivated, I think chances’ conformance to PP* should satisfy any intuitive demand

¹² This will be true so long as t_1 is not at the beginning of time – see note 10 above.

that chances conform to something like PP. As Lewis ([1994], p. 246) puts it, after noting that even his own account of chances does not perfectly satisfy the most common-sense formulation of PP:

Nothing perfectly occupies the role [of chance], so nothing perfectly deserves the name. But near enough is good enough. [...] [A]n imperfect candidate may deserve the name quite well enough.

9 Conclusion

Throughout the preceding discussion I have remained neutral regarding many difficult questions.

It is consistent with everything I've said that Lewis is right in thinking that talk of chances is akin to talk of lifetime batting averages—that this talk refers to interesting patterns in the distribution of categorical properties, rather than to some independent furniture in the ontology of the world. But it is also consistent with everything I've said that chancy properties might exist in a more robust sense—that when Merlin casts his second-order-chancy enchantment, a new property might be instantiated—the 40% chance of a genie appearing—where this property is not constitutively determined by the distribution of categorical properties throughout the history and future of the world.

Similarly, everything I've said is neutral regarding many of the conclusions that Lange ([forthcoming]) hopes to draw regarding the constitutive relations between chances, laws, and categorical properties. For all I've said, Lewis might be right that chances and natural laws are both constitutively determined by a Humean mosaic of instantiations of perfectly natural categorical properties. Or Maudlin ([forthcoming]) might be right that what laws there are is independent of the Humean mosaic, but these together determine what chances there are. Or Lange might be right that chancy properties robustly exist alongside the Humean mosaic (and perhaps alongside primitive subjunctive properties as well), and that all these properties together constitutively determine the natural laws.

I think that consideration of Merlin's flourish-less enchantment should prompt us to revise or retreat from LMP, COND, and PP, and I think this is an interesting and striking conclusion. I don't think, however, that cases like this give us any new reason to stake a particular stance regarding any of the difficult questions just mentioned. Such cases help us to see more clearly what is at issue when we ask these difficult questions, and they help us avoid some deceptively alluring pitfalls as we seek answers, but they will not force us to accept any particular answers.

Acknowledgements

I am especially indebted to Marc Lange whose work in progress inspired my interest in this topic, and whose conversations greatly influenced my thinking about it. I am also thankful for helpful comments from Marshall Abrams, John Carroll, Charles Morgan, John Roberts, Alex Rosenburg, Sarah Wright, and an audience at the Society for Exact Philosophy. Finally, I owe thanks to the wonderful philosophical community in the research triangle, especially UNC and the Triangle Ellipse, for hosting me for the 2004-2005 academic year and offering great venues for discussing this and other topics.

References

- Lange, M. [forthcoming]: 'Do Chances Receive Equal Treatment Under the Laws? Or: Must Chances Be Probabilities?', *The British Journal for Philosophy of Science*.
- Lewis, D. [1986]: 'A Subjectivist's Guide to Objective Chance', in D. Lewis, *Philosophical Papers—Volume 2*. New York: Oxford UP, pp. 83-113.
- _____. [1994]: 'Humean Supevenience Debugged', in D. Lewis, 1999, *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge UP, pp. 224-47.
- Maudlin, T. [forthcoming]: 'Why Be Humean?', in *The Physics Within Metaphysics*, Oxford UP.
- Schaffer, J. [2003]: 'Principled Chances', *The British Journal for the Philosophy of Science*, **54(1)**, pp. 27-41.