Chance and Necessity

Daniel Nolan, Draft

To appear in *Philosophical Perspectives* 30.1, December 2016

Many people, by now, are used to the idea that it makes sense to be uncertain about various necessary truths, even logical truths. And not just that it makes sense as a psychological matter, but that it can be *rational* to be less than completely certain of some necessary truths, or *rational* to be less than completely certain of the falsehood of impossibilities. Cases of the necessary *a posteriori* are like this: it was rational for chemists in the early nineteenth century to be uncertain about the chemical composition of water, for example, even if it is necessary that water is H_2O. If I am uncertain of "The Blues will win the next chariot race" I should be equally uncertain of "Actually, the Blues will win the next chariot race" even though, given the technical use of "actually" popular among philosophers, the second sentence is necessary if true at all.

Even apart from these sorts of cases, it seems very sensible to be less than completely confident of all necessary truths. The standard view is that the identities in basic arithmetic are necessary (e.g. 78+113=191), but many rational people are not completely certain they have added up their restaurant bills correctly. (If someone offers you a bet on a restaurant bill sum, giving you a dollar if you are right and taking everything you own if you are wrong, don't take that bet.) Even those who think that mathematical truths are contingent will typically believe there are logically true sentences that are about as uncertain as mathematical ones: the claim that the first eight axioms of Peano arithmetic materially imply that 78+113=191, for example can be expressed in predicate logic. This is even before we get to questions such as whether it is rational to be uncertain about whether classical logic or the axioms of classical probability: if it is rational to be less than completely certain of our logic or probability theory, then this will be another sort of reason to have a credence of less than 1 in the theorems of logic and probability. These sorts of considerations also show that contingent claims that are necessarily equivalent to each other should not always receive the same credence. One simple example: "Obama is in Washington and water is H_2O" is necessarily equivalent to "Obama is in Washington", but any reasonable uncertainty about "water is H_2O" will make the first claim less credible than the second.

Coping with this requires some modification of a standard interpretation of credences, on which they are measures over possibilities, or possible states, outcomes, or events. A conservative way to do this is to understand credence measures as measures over both possible and impossible worlds/states/outcomes/events. Once we realise that people can believe impossibilities, and rationally do so, extending rational degrees of belief to impossibilities seems difficult to resist. A natural consequence of extending probability theory in this way is to deny that necessarily equivalent claims always share the same credence for an agent: an agent can be sure that water is water but unsure that water is H_2O, sure that it is raining or not raining but unsure of 78+113=191, or unsure whether it is actually raining, even though one of "actually it is raining" or "actually it is not raining" is a necessary truth.

However, even those who are willing to accept that credences may vary between logically equivalent propositions may well think the case of *objective chance* is different. What could it mean for necessarily equivalent propositions to have different *chances*? Surely there is no chance that one could happen without the other? In addition, objective chance is often thought of as (objective) possibility coming in degrees. Necessarily equivalent propositions cannot differ from each other with respect to whether they are possible, or necessary, so it can seem natural that they could not differ in what chances they receive either.

A final motivation to think that necessary equivalents cannot differ in their chances is the thought that our motive for distinguishing necessary equivalents is only to handle issues about our representation of the world, not the world itself (e.g. that we may have distinct necessarily equivalent representations of the world because of our ignorance, not because we are tracking different states in the world itself). I have argued elsewhere that this is a mistaken view (see Nolan 2014), but those tempted by a view like this may well think that objective chance is about the world, not about how we represent it, and so should not be sensitive to any distinction between necessary equivalents.

The view I will be arguing against in this paper is a principle about the substitution of necessary equivalents in chance statements. In particular, the view I will be arguing against is the view that necessarily equivalent propositions always receive the same objective chance. For those who prefer to think of chances as features of events rather than of propositions, the target will be the claim that events which necessarily co-occur always receive the same chance. To make the following discussion easier, let me label these views (where $Ch_x()$ is a chance function relative to parameters x):

SUB: For any propositions P and Q, if P and Q strictly imply each other, then $Ch_x(P)=Ch_x(Q)$

Or if we prefer to state chances as features of events rather than values of propositions,

SUB(events): For any events A and B such that necessarily, A occurs if and only if B occurs, then $Ch_x(A)=Ch_x(B)$.

Of course, even the same proposition bearing different chances would be unsurprising if there were more than one *kind* of genuine objective chance. And if chances are relative to some parameter (times, chance setups, worlds, or whatever), again it would be no surprise if even one and the same proposition bore different chances relative to different times, different chance setups, different worlds, or whatever. *A fortiori,* necessarily equivalent propositions could easily have different chances of different kinds, or different chances relative to different parameters However, these types of variation are not my concern here. Keep the kind of objective chance fixed, and the parameters, if any, fixed. With the same choices made for kind and parameter, can necessarily equivalent propositions differ in their chances? SUB says not, but I will argue in this paper that we should say that SUB is mistaken, even when this specification about kinds and parameters is taken as read.

Despite this appearance, and despite the almost universal treatment of chance as a measure only over possibilities (whether possible worlds, possible events or possible states), there is a paradox about objective chance that is best resolved by allowing necessarily equivalent propositions to vary with respect to chance. Or so I will argue: avoiding this result would also come with serious revisions to received wisdom about objective chances.

Principal Principles Go Both Ways

To set up this paradox, we need to invoke the Principal Principle, or one of its variants. David Lewis's original Principal Principle is a claim about the connection between rational prior probabilities and objective chances. He assumes that chances are relative to both times and worlds (according to Lewis chances are contingent, and evolve over time, with propositions about earlier times no longer being chancy at later ones). The very general idea is that certainty of the chances at a time is as good a guide as we can have to outcomes of events, at least when our only other information is about events up to that time and the general principles about chance (e.g. the probabilistic laws of nature). Lewis's preferred way of stating his Principal Principle is in terms of *initial* credence functions: those ideal credence functions that (perhaps *per impossibile*) ideally rational creatures would have before any information received through "updating": for example, before any information from the senses. These initial credence functions build in conditional credences, which specify, in effect, what credence the individual would come to have were she to update by conditionalisation (by becoming certain of various things as "evidence"). Lewis's first formulation of an unmodified Principal Principle is (Lewis 1983 p 87):

$C(A/XE)=x$

Where C is an initial credence function, A is any proposition, X is the proposition that the objective chance of A is a number x, and E is any "admissible" proposition (e.g. one about the past or the laws of nature). This says that in an initial credence function, the conditional credence of a proposition A on the proposition that A has chance of x, together with as much or as little allowable information you like, will be x. A way of expressing a similar idea is to put it in terms of unconditional (posterior) credences to have: being told (for certain) that the objective chance of A is X is enough to ensure a rational person has a credence of X in A, almost regardless of any other information she might have. Information about chance has an overriding role in determining rational belief, according to Lewis.

A "reformulation" Lewis offers of the Principal Principle suggests that the best inquirers can do with perfect information about the past and the "theory of chance" is to set their credences in propositions as equal to the chances of those propositions (p 97):

$Ch_{tw}(A) = C(A/H_{wt}T_{w})$[1]

That is, the conditional credence a perfectly rational agent has on a proposition A, given full information about the history of a world up to time *t* together with the "theory of chance" that is true in that world, is equal to the objective chance of A at *t*, in *w*. (Lewis derives this formulation from his original formulation on pp 96-97, though some of the steps in the derivation require substantial assumptions.) This reformulation also provides a recipe to derive objective chances from ideal credences: in principle, if you know enough about what conditional credences ideally rational agents have in various contingent propositions, given facts about the history and principles like the laws of nature, you know enough to tell what the objective chances are.

This reformulation is in terms of conditional credences of certain kinds of ideal agents, and Lewis has reasons to present things this way rather than in terms of what credences agents *would* have were they to become certain of the story about the history and laws of an event (and they were to lack other kinds of information, such as direct information about whether A is true, for example). See Meacham 2010 for a clear presentation of why Lewis preferred this way of formulating a chance-credence link. However, Lewis's principle about initial credence functions has the following

[1] I have replaced Lewis's "P" as the symbol used to represent the objective chance function with "Ch".

consequence: if an ideal reasoner with such an initial credence function updates by conditionalisation on the conjunction of H_{tw} and T_w, and learns nothing else, her credence in A will be equal to the objective chance of A at *t*. It will be convenient to talk about what credence an ideal reasoner would have when she learns H_{tw} and T_w with certainty, and the link between that an objective chance, so I will often put things that way below, though a scrupulous reader who instead wishes to consider initial conditional credences such a reasoner might have should be able to recast the cases below in that light, *mutatis mutandis*.

There is a significant literature on proposals to tweak principles to do roughly the same job for which Lewis proposed his Principal Principle. Famously, Lewis himself thought it needed to be modified to go better with his Humean theories of law and chances in Lewis 1994. I do not want to rely on any peculiarities of Lewis's account, however, and alternatives to his Principal Principle will serve my purposes just as well. The idea which is the common core of the family of principles can be put in one of two ways: either chance is very close to the credence that an ideal, suitably informed, inquirer would have if she was certain of specific information and had no extra, "inadmissible" information; or chance is close to the *conditional credence* an ideal inquirer would have, as a prior probability, *conditional* on the relevant information. At least when the chance and credence concern events in worlds like ours, now and hereabouts, and given the kinds of chances that are likely to be found in a world like ours. Those familiar with the range of variant proposals about the link between credence and chance should be able to reconstruct the cases I will present in those frameworks: and I would risk confusing those unfamiliar with the variety of alternatives if I were to complicate my discussion to be compatible with them all.[2]

[2] One dimension of variation that is becoming more popular is to allow that some objective chance functions are not temporally asymmetric in the way suggested in the text, and that there is no particular privilege that we should give to the *history* of a chance event rather than the chance of it given other sorts of information. (See Handfield 2012 and the papers in Wilson 2014 for recent discussions, and for references to the literature on time-symmetry.) So the reader may wish to see cases of ideal reasoners with information about the history and theory of chance as a special case of how the credences, and conditional credences, of ideal reasoners should interact with objective chance. Seeing the cases discussed as special cases should not matter for their cogency, though, and even more examples of the kind of problem I have in mind can be constructed if we have objective chance functions that are not time-asymmetric in the way that Lewis assumes they are.

For our purposes, we can think of an ideal scientist armed with all the information about the laws of nature, and all the information about the past, and ask what credence she would have in a future outcome: that a given atom decays, or that it rains tomorrow, for example. I will help myself to the "reformulation" of the original Principal Principle below ($Ch_{tw}(A) = C(A/H_{wt}T_{w})$), labelling it "PP" for convenience. To remind the reader that alternatives to Lewis's Principal Principle will deliver slightly different conditional and unconditional probability verdicts, I will often signal that the application of the PP delivers a verdict on credence and chance that is at least *close* to correct in the circumstances to be discussed: all the leading rival principles to Lewis's original account give similar probabilistic verdicts in the cases below.

Notice the reformulated PP gives you two guides: not only a guide to what initial credences to ideally have, in terms of the chances, but also a guide to what the chances are, in terms of what ideal investigators would have their credences (or conditional credences) set to. The value of a PP is both that it explains why information about chances could be useful as a guide to what to believe about whether chancy outcomes will occur, *and* as a guide to what quantity in the world could turn out to be the objective chances of outcomes.

The Paradox for SUB

With the PP in place, some of the cases that cause puzzles for rational credences over possibilities will also cause puzzles for objective chance. I will label the first puzzle *Actuality*. Consider an everyday chancy proposition: e.g. that it will rain in Berlin sometime on June 3rd 2025. Let us suppose this has an objective chance, at noon on June 2nd 2025, of 0.8. Our hypothetical ideal scientist, given all the information about the laws of nature and the past up until noon on June 2nd 2025, but no other information, would have a credence of (close to) 0.8 that it rains sometime on June 3rd. Alternatively, *conditional* on that body of information, our hypothetical ideal scientist at the (hypothetical, ideal) start of inquiry would have a *conditional* credence of (close to) 0.8 of rain on June 3rd.

Now consider the proposition that *actually,* it will rain sometime on June 3rd 2025 in Berlin. (I mean to be using "actually" in the sense familiar from modal logic: that *actually p* is true in all possible worlds when *p* is true at the actual world. This may or may not correspond to any non-jargon use of the expression "actually".) Since our idealised scientist can be assumed to know how "actually" works, it looks like she has just as good information about whether it will *actually* rain as she has about whether it will rain. On the face of it, the best credence to have in whether it will actually rain, given all the information about the past and the laws on June 2nd, will also be (close to) 0.8. And likewise, the best conditional credence for our ideal starting scientist, on the relevant proposition about the past and the laws as of June 2nd, will be (close to) 0.8.

But given the rule for the application of "actually", the proposition that *actually* it rains on June 3rd is either necessarily true or necessarily false. So it is either necessarily equivalent to a logical tautology or a logical contradiction. Let us suppose that it does rain on June 3rd, so that "actually, it rains on June 3rd 2025 in Berlin" is true, and so necessarily equivalent to the proposition that either it rains on June 3rd 2025 in Berlin or it does not. The probability calculus ensures that the triviality that it will rain or it will not gets chance 1. So now we face a challenge for SUB: by applying the principal principle in the way I suggested, one proposition has chance 0.8, while another, necessarily equivalent to it, has chance 1.

The technique used for "actually, it rains in Berlin sometime on June 3rd" can be easily generalised: indeed, if there were two true propositions *at all* with different chances, there would be necessary propositions with different chances on the above assumptions, since the actualised variants of those two truths would both be necessary and would have different chances by a variation of the above argument. This is the first version of the kind of paradox I want to resolve here.

Some will suspect there is something specifically peculiar about the 'actually' operator, as usually understood in discussions of modal logic, which is responsible for this problem. After all, one might doubt that there can be an operator "actually" that functions in the way needed; or one might bite the bullet and insist that, counter-intuitively, the credence we should have in *p* and the actualised variant of *p* often varies dramatically; or one might suspect a formal trick even if one cannot quite put one's finger on where it is. So I will present a different, and hopefully even more intuitively compelling argument to the conclusion that something should give when we combine the Principle Principal (or one of its cousins) with SUB: and I will argue that the most tempting way to resolve the tension is the abandonment of SUB. Let us call the next case to be discussed *Charlie*.

Charlie

Suppose I have 1000 atoms of ^{17}N: an isotope with a half-life of a little less than five seconds. Give the next atom of this batch to decay the name "Charlie". Suppose, for simplicity, somethings which may in fact be true: that the internal structure of the atoms is relevantly the same and there are not unknown quantities which give some of the atoms a higher chance of decay than others, and that half-lives express genuine objective chances and not just our ignorance. Now consider two questions:

What is the chance of Charlie decaying in the next five seconds?

What credence should an ideally rational agent have that Charlie will decay in the next five seconds, conditional on complete information about the laws and the history?

I think the answers to these two questions diverge greatly. The chance that Charlie will decay in the next five seconds is a little over .5: after all, the atoms each have an objective chance of a bit over .5 of decaying in the next five seconds. However, I think the credence we should assign to the claim that Charlie will decay in the next five seconds is very high: something in the vicinity of $1 - (0.5)^{1000}$. The conditional credence given the laws and the history should be roughly our credence that at least one of the atoms will decay in the next five seconds, given the way "Charlie" was introduced. $Ch_t(A)$ diverges greatly from $C(A/ H_{tw}T_w)$.

We might add something else to the story to make it clearer that our idealised theoriser can entertain the relevant hypothesis about Charlie's decay. Let us suppose that an experimenter is about to start an experiment monitoring 1000 atoms of ^{17}N, each with a half-life of just under five seconds. The experimenter has already labelled them a_1, a_2, ... a_{1000}. Let the dubbing of Charlie be one of the events that happens shortly before the relevant 5 seconds, for example in the laboratory. Then the occurrence of the dubbing will be part of the information about the history that our idealised inquirer will either be certain of, or be able to have a conditional credence on the basis of.

If you wish, we could introduce other stipulations associated with the introduction of the name "Charlie" to deal with the possibility of exact ties of decaying and the (remote) possibility that none of the atoms ever decay. We could introduce the name "Charlie" to name "the next atom of this batch to decay, if there is a unique one, and Jupiter otherwise" – or if you are hung up on trying to work out Jupiter's chance of decay, you may of course select an entity other than Jupiter.

We can then contrast two propositions which apparently should get quite different credences. Suppose a_{17} is in fact the unique next atom to decay in the set-up, and the experimenter dubs "Charlie" just before the start of the observation. Then we have two claims:

(1) a_{17} will decay in the next five seconds.
(2) Charlie will decay in the next five seconds.

(1) and (2) are apparently necessarily equivalent, in as strong a sense of necessity as you like. (I am assuming that names are rigid designators and something like the necessity of identity is true, or at least that any counterexamples to it are irrelevant here.) Given the PP, (1) receives a chance of (about) 0.5 and (2) receives a chance very close to 1 (greater at any rate than 0.99). It seems that again we should not insist that necessarily equivalent propositions receive the same objective chance, and so reject SUB.

Alternative Diagnoses

We could block this argument, as before, by rejecting the PP, though if we chose to take this route in order to preserve SUB, we would also need to reject many of the rivals in the literature to Lewis's Principal Principle as well as the letter of the Principal Principle.

What other alternatives do we have available, should we wish to keep both SUB and a PP-like principle? One option would be to reject the commitments about the name "Charlie" needed to get the problem going. We could, for example, abandon any theory of names that lets us conclude that (1) and (2) are necessarily equivalent. Or we could deny that "Charlie" is a name at all in the above example: though I cannot see why conferrals of names cannot occur in this way. Some philosophers I have talked to wonder whether "Charlie" even refers before the next atom decays, on the grounds that indeterminacy in the future infects dubbings like this: if so, you could recast this as a discussion of what the chances *were* of (1) and (2), seen from a vantage point well after a_{17} decayed.

Another kind of option would be to say that in employing an actuality operator, or in introducing the name Charlie in the way that we have, we have introduced "inadmissible information" in the sense introduced by Lewis, and so we cannot use the conditional credences of our ideal scientist who is aware of that information to determine the relevant objective chance. I have not seen anyone develop an account of how introducing an actuality operator into our language gives us inadmissible information, but Schwarz 2014 p 90 suggests about a move like the introduction of "Charlie" above that it introduces an "inadmissible identifier": conditional credences about when Charlie decays are, in effect, not just credences that some particular atom will decay in the next five seconds, but on a proposition that an atom, *being the next one to decay*, will decay in the next five seconds. John Hawthorne and Maria Lasonen-Aarnio 2009 also talk about a case like *Charlie* and diagnose it as involving knowledge of contingent a priori truths, and suggest that the truth of connected "belief episodes" is "inadmissibly connected to the future": though they have little to say about knowing what we ordinarily know when we witness this sort of dubbing would yield inadmissible information.

The thought is that the introduction of “Charlie” is something like finding out, in advance, *which* atom will be the next to decay. If, for example, a time-traveller returned from the future with that information to add to an ideal scientist’s opinions that would enable the ideal scientist to do better in forming her credences about decay than just conforming her credences to the objective chances: but this is no threat to the Principal Principle, because Principal Principles only tie the objective chances to idealised credence, or conditional credence, of informed inquirers who have all admissible information but lack "inadmissible" information. After all, whatever you know about the chances of a chancy coin flip, if you also knew with complete certainty that the coin land heads on this trial your credence in heads should be 1, not whatever the chance of heads is.

As far as I can see, though, the introduction of the name “Charlie” does not give our ideal scientist any information, in any direct way at least, other than information about the past and the laws. Consider a situation where Wayne the experimenter utters the words “I hereby dub the next atom in the experimental setup to decay ‘Charlie’”. He makes some noises, and some mental activity goes on: referential intention forming, for example. Our ideal scientist should be able to learn that these things went on with Wayne near the apparatus just before the start of the observations. What people say near apparatuses looks about as ordinary a thing to learn as I can think of. It has all the usual marks of admissible information—it is just information about relatively ordinary, contingent goings-on in the past, and is the sort that we might expect ordinary inquirers to possess. Likewise, when we have a conditional credence on Charlie being the next one to decay plus other facts about the past, including Wayne's dubbing, the fact about the dubbing is just one more ordinary proposition about the past leading up to the experiment to conditionalise on. While the option to radically revise what sorts of facts are admissible is available, there does not yet seem to be a very satisfying philosophical explanation of why things we learn from observing such dubbings should count as inadmissible in the way that facts about the future of the chance episode might.

Yet another option would be to decide that chances are not features of propositions at all, but rather something else. There are several versions of this strategy: for example, we could insist that objective chances are features of sets of possible worlds directly, whether or not those are suitable contents of belief: then necessarily equivalent objects of chance would be *identical*, leaving, apparently, no room for a puzzle about assigning different chances to necessarily equivalent objects of chance. Or we could do an equivalent move in the vocabulary of assigning objective chances to events, treating necessarily co-occurring events as identical. If we went either of these routes, a Principal Principle will no longer give us a direct road from ideal credence to chance: given the ideal credences, there will be a further question about what chances are had by sets of worlds or events.

If we deny that chances are features of propositions, then it will be much trickier to say what sort of Principal Principle we should endorse: or indeed what information about chances should do to our degrees of belief, given that our beliefs and credences, at least, concern propositions.[3] However, many ways of pursuing this strategy will have the same effect as my proposed solution: if we associate *chances* with something other than propositions, the kind of problem I have raised here will strongly suggest that when it comes to specifying the link between credences and chances, necessarily equivalent propositions will not always be treated the same as each other by whatever principle links chances and rational degrees of belief in propositions.

Another option would be to assign chances to epistemically individuated propositions, such as David Chalmers's "sets of epistemically possible scenarios" associated with

[3] Salmon, manuscript, in response to cases he developed independently with similar structures to *Actuality* and *Charlie*, and proposes that chances of propositions are relative to propositional guises. He denies SUB, allowing disparate objective chances to obtain of the same proposition. If I have understood his view, he would also say credences in propositions are also relative to propositional guises, thus treating chances and credences as involving the same objects of probability, in both cases relative to propositional guises.

sentences (see Chalmers 2011). Chalmers argues that credences should attach to these sets of epistemic scenarios, and plausibly, "it rains tomorrow" and "actually it rains tomorrow" are epistemically equivalent. At least for those aware of how "Charlie" is introduced, "Charlie decays in the next five minutes" and "the next atom to decay in the experimental setup" are also plausibly epistemically equivalent (or near enough). Whatever the merits of treating sets of epistemic scenarios as the objects of credence, however, it would be odd to treat them as the objects of objective chance unless we were idealists about 'objective' chance: at the very least we would be owed a story about why objective chances such as those embodied in atomic half-lives have anything very much to do with what idealised versions of us can distinguish *a priori*. It is also not clear whether to interpret this suggestion as accepting or rejecting SUB, since sentences with different modal profiles are supposed to be associated with the same sets of epistemic scenarios: on one very natural way of interpreting the proposal that chances are assigned to these sets, it is a way of modelling necessarily equivalent sentences diverging in their chances, and so is a way of *implementing* the failure of SUB for the entities that are most naturally ascribed features such as necessary equivalence (i.e. Chalmers's "secondary intensions").

How Best to Proceed Without SUB?

Even if we do as I suggest and reject SUB, the current situation can still seem unsatisfactory. After all, we have an argument that (1) should receive a chance of about 0.5 and (2) should receive a chance near 1 (<0.99). But it may seem strange to think that Charlie/a_{17}/whatever we call it should have such different chances of decay—can we even talk of objects having chances of decay at all in this scenario?

The most satisfying way to resolve this uneasiness, I think, is to restrict the PP.[4] If only some propositions have chances, it would make sense to restrict the PP to a

4 Even in Lewis's original presentation, he does not insist that the Principal Principle applies to all propositions, though the exceptions he envisages are quite different from the ones discussed here (Lewis 1986 p 132). Since Lewis identifies necessarily equivalent propositions with each other, he will not be able to have cases where necessarily equivalent propositions vary with respect to whether they have chances. Despite Schwarz's different diagnosis of what is wrong with cases like *Charlie*, his positive proposal about the Principal Principle ends up being related in some ways to the one I propose

chance/credence link *for the propositions that have chances at all*. Once we do this, we have the option of denying that both (1) and (2) have chances: and indeed once we do this, it is natural to think that (2) is not the sort of proposition we need to assign a chance to in the first place. (To lack a chance altogether is different from having a chance of 0: 0 is a perfectly respectable value for a chance to have, while propositions that lack chances (if there are any) would not even have chance 0.) Propositions like (1) are the sort that we naturally infer chance propositions from, given the theory of radioactive decay, not (2). To be sure, we naturally apply chances to some propositions that are superficially like (2), such as the proposition that "some atom in the sample will decay in the next five seconds". But that proposition is not equivalent to the proposition that Charlie will decay in the next five seconds: the proposition that some atom in the sample decays in the next five seconds is true in many states where Charlie takes longer to decay than it actually does, and is beaten by some other atom.

Notice that if we allow (1) has a chance of 0.5 but (2) lacks a chance altogether, we would still have to reject the principle that necessarily equivalent propositions have the same chance. The pair of (1) and (2) would still be a perfectly good counterexample to SUB. Allowing that some troublesome propositions lack chances altogether, however, would let us, if we wished, adopt a more restricted form of SUB: we could continue to maintain that any pair of necessarily equivalent propositions *which have chances at all* share the same chance. The problems for SUB raised in this paper do not cause trouble for this more restricted form of SUB.

Restricting the PP to only some propositions also allows for some other options. We could allow, for example, that some propositions that did not obey the PP nevertheless had chances, for example. A theory that did so could assign (1) and (2) the same chance (presumably 0.5). I think it is more appealing to say that the PP (or whichever relative we adopt) applies to all propositions *with chances*: we lose some of our grip on chance itself if information about chances sometimes constrains credence in a PP-like way and sometimes does not. So I do not recommend assigning (2) a chance

---

here, restricting it only to certain of the properties we might hypothesise that objects possess. (Schwarz 2014 pp 90-92 despite holding that PP-like principles do not apply to it. Still, I think it is interesting to notice that restricting the PP leaves some questions about chance unsettled without some further principle, such as requiring that all propositions with chances at all conform to the PP (or whichever relative we endorse).[5]

Taking the option that (2) lacks a chance altogether does leave us with one puzzle. Consider the atom Charlie (or a_{17}, or by any other name). We might want to ask what chance *it* has of decaying, nevermind what it is called. Chances seem *de re*, especially for those who think that single-case chance is in entirely good order. But if (1) and (2) receive different chances (or one receives a chance and the other does not), we may need to be more careful. And the puzzle does not stop here. It seems very natural to talk about an object or an event in a number of different ways and feel entitled to retain chance ascriptions: if a_{17} is Ludwig's favourite atom, and a_{17} has a 0.5 chance of decaying in the next 5 seconds, then Ludwig's favourite atom has a 0.5 chance of decaying in the next five seconds (and so we had better be ready for Ludwig's disappointment). Or if a_{17} is the only atom in partition 20 of our apparatus, and a_{17} has a 0.5 chance of decaying in the next 5 seconds, then the only atom in partition 20 of our apparatus has a 0.5 chance of decaying in the next five seconds. Both the inference about Ludwig's favourite atom and the atom in partition 20 seem entirely fine—deductively valid even!—but it is hard to see how they could be good inferences if atoms do not have chances of decay *simpliciter.*

We would have a solution to this problem that would be somewhat satisfactory if, for each object and property it might have, there is at most one chance function that assigns a chance to a proposition saying that object has that property. When there is a proposition like this for an object and a property that receives a chance C, then, we could allow ourselves to make attributions of C to that object having that property however we picked out the object (or the property). We could then help ourselves to locutions like "Ludwig's favourite atom has a 0.5 chance of decay in the next five seconds", even if the proposition assigned a chance designates that atom in some

5 My current suspicion is that there are reasons to reject even the restricted principle that any two necessarily equivalent propositions that have objective chances have the same objective chance: but these reasons are quite different from any touched on here.

other way, and does not mention Ludwig or his preferences at all. If this solution was pursued, we would need to be careful not to confuse these de re attributions with chances of propositions: if it turned out that a_{17} had a 0.5 chance of decaying, and so we could correctly say that Charlie did, it would still not do to hold that the *proposition* that Charlie decays in the next 5 seconds received a 0.5 chance, on pain of abandoning the proposed solution to the main puzzle posed in this paper. Still, this suggestion would recover some of the *de re* behaviour of chance ascriptions, at least provided objects rarely or never were ascribed a given property by different propositions which each received a chance relative to a particular chance function (and set of parameters).

Conclusions

The solution I have recommended to the problems raised in this paper is that only some propositions get objective chances, for any given objective chance function. Furthermore, these propositions are not just strange propositions that we cannot easily express: they include ones we can easily evaluate and act upon. (We could bet whether it will actually rain on a given date, or when exactly Charlie will decay.) This is potentially trouble for decision theories that assume outcomes always have well-defined chances.

There are other motives for being even more radical than I have done in this paper and insisting that, for example, distinct metaphysical impossibilities get distinct non-zero chances. (This is how Nathan Salmon reacts to the kinds of cases I have discussed in Salmon, manuscript.) Some notions of probability used in mathematics seem to have this feature: the probability that a given number is prime can be non-zero even when, in fact, it is composite, if we take with full seriousness the talk of "probability" used in some statements of the prime number theorem, for example. And there are a number of other applications of the thought that impossibilities may have non-zero objective probabilities, though a discussion of these should be left for another occasion.[6] If a more radical approach to objective probabilities and impossibility is required, the conclusions of this paper may need to be revisited: though at present I think the preferred diagnosis of this paper for the particular cases discussed will survive consideration of more exotic cases that may require different treatments.

Finally, the case of objective chance is one more argument for a more general metaphysical contention I have been defending: that hyperintensional theoretical resources, that is, theoretical resources which distinguish necessary equivalents, are needed for adequate theories of the non-representational world and not merely of our representations. That is, as well as needing to recognise hyperintensionality to account for belief, meaning, fiction, and so on, we need resources to theorise about metaphysical topics that are not matters of representation: in some good sense, there is hyperintensionality in the non-representational world (Nolan 2014). If objective chance ascriptions are hyperintensional: if, that is, substituting necessarily equivalent propositions into operators attributing objective chance can change truth-value of claims, then we have another case where hyperintensional notions are needed beyond our theory of representation. If objective chance is in the world and not merely in our representations of it, then so is at least one worldly hyperintensional phenomenon, and not merely representational one.

One topic for future work, if this proposal is on the right track, is to determine *which* propositions get genuine objective chances and which do not. It would be good to

[6] Let me mention two other potential applications for non-zero probabilities of metaphysical impossibilities in metaphysical theorising. One is to allow that there is a non-zero physical chance of metaphysically impossible autoinfanticide in time-travel cases, as in the view defended in Horacek 2005. A recent development of that line of thought is in Effingham, in preparation. Another would be as a way of reconciling macroscopic indeterminism with microscopic indeterminism: when an event E receives a macroscopic chance at a time of, for example, 0.5, even though E is incompatible with the past and the microscopic laws, it may be convenient for the relevant macroscopic chance function to give, at that time, a non-zero chance to the *conjunction* of (i) a proposition completely specifying a time in the past of E, (ii) a specification of the microscopic laws, and (iii) the claim that E occurs. (The idea is that the macroscopic principles encoded in the chances leave E open, and do so even when past microscopic principles are specified.) This option would fit the spirit of Glynn 2010's treatment of macroscopic chance, though I do not think Glynn himself is committed to it in that paper. Other treatments of macroscopic chance such as List and Pivato 2015 require that macroscopic chance functions are defined over different algebras of propositions than microscopic chance functions when macroscopic chance co-exists with microscopic determinism, and so would avoid ascribing non-zero macroscopic chances to impossibilities.

have something principled to say about what the problematic propositions in *Actuality* and *Charlie* have in common in virtue of which they lack chances. I can see two broad strategies for pursuing this project. One would be a "top-down" one, where we begin with the assumption that all propositions are fit to get chances and then come up with "monster barring" principles to rule out apparently troublesome ones. Another, which seems more promising, would be a "bottom up" process by which we identify some propositions deserving of objective chance, and then articulate principles that take us from the chances of those propositions to chances of others in ways that do not lead to trouble. For example, we might begin with propositions about events that are explicitly attributed chances in our best physical theory, then extend chances to boolean combinations and other logical combinations of these, and look for other principled ways to extend chance attributions further in such a way that (among other things) do not lead to ascribing conflicting objective chances to necessarily equivalent propositions, at least not in the sort of cases pointed out above.

Of these two strategies, I would prefer the latter: it would not only provide more security that there are not more problems awaiting discovery, but it would give us insight into why the various propositions with objective chances have the chances that they do. But the details of such a construction are for another time.[7]

Daniel Nolan
Department of Philosophy
University of Notre Dame
Notre Dame, Indiana 46656
USA

References

7 Thanks to Sara Bernstein, Catharine Diehl, Alan Hájek, Elizabeth Harman, John Hawthorne, Nathan Salmon, Alexander Sandgren, Jason Turner; and audiences at the 2014 Australasian Association of Philosophy meeting, the University of North Carolina at Chapel Hill, the Humbolt University and the University of Virginia for discussion. This research was supported by the National Humanities Center and by the Australian Research Council's *Discovery Projects* funding scheme (project number DP130104665).

Chalmers, D. J. 2011. "Frege's Puzzle and the Objects of Credence". *Mind* 120.479: 587-635

Effingham, N. in preparation. *The Philosophy of Time Travel*.

Glynn, L. 2010. "Deterministic Chance". *British Journal for the Philosophy of Science*. 61: 51-80

Handfield, T. 2012. *A Philosophical Guide to Chance*. Cambridge: Cambridge University Press

Hawthorne, J. and Lasonen-Aarnio, M. 2009. "Knowledge and Objective Chance" in Greenough, P. and Pritchard, D. (eds.) *Williamson on Knowledge*. Oxford: Oxford University Press pp 92-108

Horacek, David. (2005) ‘Time Travel in Indeterministic Worlds’. Monist, 88, 423–36.

Lewis, D. 1986. "A Subjectivist's Guide to Objective Chance", with postscripts, in Lewis, D. *Philosophical Papers Volume II*. Oxford: Oxford University Press, pp 83-132

Lewis, D. 1994. "Humean Supervenience Debugged". *Mind* 103.412: 473-90

List, C. and Pivato, M. 2015. "Emergent Chance". *Philosophical Review* 124.1: 119-152

Meacham, C. 2010. "Two Mistakes Regarding the Principal Principle". *British Journal for the Philosophy of Science* 61.2: 407-431

Nolan, D. 2014. “Hyperintensional Metaphysics”. *Philosophical Studies*. 171.1: 149-160

Salmon, N. manuscript. "Impossible Odds".

Schwarz, W. 2014. "Proving the Principal Principle" in Wilson, A. (ed). *Chance and Temporal Asymmetry*. Oxford: Oxford University Press pp 81-99

Wilson, A. (ed). 2014. *Chance and Temporal Asymmetry*. Oxford: Oxford University Press