

Finite Frequentism in a Big World

Nick Tosh

(Penultimate draft. Please cite published version, forthcoming in *The British Journal for the Philosophy of Science*.)

ABSTRACT

The view that chances are relative frequencies of occurrence within actual, finite reference classes has long been written off. I argue that it ought to be reconsidered. Focussing on non-deterministic chance, I defend a version of finite frequentism in which reference classmates are required to have *qualitatively identical pasts*. While my analysis can evade or resist several standard objections, it has a counterintuitive consequence: non-trivial chances entail the existence of light cones that are perfect intrinsic duplicates. In mitigation, I argue that our scientific knowledge is consistent with the hypothesis that there are many such duplicates in the actual world. Moreover, my analysis has some striking advantages: it is simple, it is metaphysically undemanding, and it makes possible a satisfying explanation of the chance–credence connection.

- 1 Introduction
- 2 Target
- 3 Sketch and Skirmish
- 4 Articulation
- 5 Standard Worries
 - 5.1 Frequency tolerance
 - 5.2 Leibniz’s dictum
 - 5.3 Single case chance
 - 5.4 Missing values
 - 5.5 Explanation
 - 5.6 Counterfactual chances
- 6 Chance and Credence
 - 6.1 An infeasible indifference intuition
 - 6.2 Articulating the credence principle
 - 6.3 Appraising the credence principle
- 7 Confirmation
- 8 Laws
- 9 Summing Up

1 Introduction

I will be defending a finite frequentist analysis of chance. The word ‘chance’ can, of course, be used in many ways. Roughly speaking, I am interested in the sorts of chances that (even) a Laplacean super-intelligence—a being with complete knowledge of the past, complete knowledge of the laws of physics, and unlimited computational resources—might care to know. Quantum chances are the chief motivating examples, but I should stress at the outset that this is not a paper on the interpretation of quantum *theory*. I will have more to say about the analysandum in the next section. In the

meantime, I need to apologise for my analytical strategy.

Finite frequentism is the view that ‘the probability of an attribute *A* in a finite reference class *B* is the relative frequency of actual occurrences of *A* within *B*’ (Hájek [1997], p. 212). (‘Actual’ should not be read rigidly. The idea is that probabilities in a given world are relative frequencies *in that world*.) Relative frequencies satisfy all the probability axioms. They are mathematically simple and metaphysically unmysterious. Moreover, measuring a probability is very often a matter of measuring a relative frequency, so it is not surprising that some ‘computer scientists, statisticians, physicists, economists... seem to speak as if probability simply *is* relative frequency’ (ibid., p. 210, original emphasis). However, finite frequentism has almost no support in the philosophical literature.¹ As a template for analyses of chance, it faces a battery of standard objections: (i) chances do not put *logical* constraints on frequencies; (ii) non-trivial chances may attach to one-off, *sui generis* trials; (iii) chance is an irreducibly modal notion; (iv) chances may take irrational values (e.g. $1/\sqrt{2}$); (v) frequencies are explanatorily impotent; (vi) frequencies do not display the kinds of counterfactual dependencies we associate with chances.² So far these are good grounds for caution but poor grounds for dismissal. The first three objections are as much of a threat to Lewisian analyses as they are to finite frequentism; Lewisian analyses are nonetheless taken seriously. The fourth counts specifically against finite frequentism, but it doesn’t look like a clincher. (Is there an intuition that chances can be irrational? Certainly. Is that intuition non-negotiable? Surely not.) The fifth and sixth will seem compelling only to those who assume that frequencies are not law-governed. But perhaps some frequencies *are* law-governed. I will have more to say about all these objections in §5. In the meantime, I want to formulate—and take a first stab at motivating—my specific proposal. It is best approached by way of the notorious *reference class problem*.

Consider the structure of our chance talk. We speak of the chance that a particular radium nucleus will decay in the next hour, period; we do not speak of the chance relative to this or that ensemble of nuclei. Anyone proposing to reduce chances to frequencies must be able to explain, in some tolerably general way, how chance-bearers pick out chance-defining reference classes. It is widely agreed that finite frequentists have no good explanation to give.³ Why the consensus? Let’s follow tradition and consider an epidemiological example. Joe is an obese white male smoker. Accordingly, the proportion of obese white male smokers who get cancer is of more interest to him than the proportion of all smokers who get cancer; it is a better (though

¹ Criticising it in 1997, Alan Hájek feared readers might complain that his target was ‘at best of historical interest as far as the philosophical literature is concerned ... unworthy of much scrutiny in these more enlightened times’ ([1997], p. 212). Gillies ([2000]) does not mention the position at all in a book-length survey of interpretations of probability. Handfield ([2012]) does, but describes it as ‘obviously incorrect’ (p. 108). I know of only one contemporary philosopher who has come close to endorsing finite frequentism: John Roberts ([unpublished]) defends ‘the possibility of interpreting—or perhaps, reconstructing—the probabilistic laws found in scientific theories as laws about frequencies’. Given the state of the literature, Roberts and I are natural allies. However, our projects are significantly different. The most obvious difference is one of subject matter: law-governed probabilities needn’t be *chances* in my sense; conversely, it is not clear that chances are *necessarily* law-governed. (I hope my analysis allows for chance laws, but it allows for lawless chances too.) There are also strategic differences. I offer an a priori solution to the reference class problem; Roberts delegates the task of defining appropriate reference classes to scientists. I justify Principal-Principle-shaped credence constraints by appealing to a self-location indifference intuition; Roberts appeals to a principle connecting confirmability and lawlikeness. I shall refer—gratefully—to Roberts’s paper in §8, when I discuss laws. Until then, it makes sense for me to develop my case independently.

² Hájek ([1997]) makes all these objections.

³ See, e.g., (Hájek [1997], p. 214, Handfield [2012], pp. 106–7).

still inadequate) proxy for *his* chance of getting cancer. Narrowing the reference class further might be expected to improve the proxy. However, one mustn't narrow it too far, or one ends up with just Joe. (The 'relative frequency of actual occurrences' of an attribute in the class {Joe} can take only trivial values, 0 or 1. That won't do.) The finite frequentist must claim that there is a sweet spot to be found here: a privileged reference class that is exactly narrow enough, and with respect to which relative frequencies are *analytically equivalent* to Joe's chances. The claim is implausible on its face. There is a strong intuition that, presented with some supposedly privileged reference class, we will always be able to stump the finite frequentist by asking 'why not a little narrower?', or 'why not a little broader?', or 'why not a similarly-sized class with different members?'

I shall argue that this intuition is mistaken. There is, in fact, a good solution to the reference class problem: two chance bearers are reference classmates iff they have *qualitatively identical pasts*. This condition is simple, general and pretty evidently not ad hoc: it is the natural way to enforce the principle that the past is not chancy. Admittedly, the condition is also rather demanding. Indeed, it is so demanding that many readers will want to dismiss it as a sure-fire route to singleton reference classes. They should think again. The universe is a big place. It might, for all we know, be much bigger than its presently surveyable portion. Joe might have many doppelgangers, and some of those doppelgangers might have past light cones that match Joe's in every detail. It is not *obvious* that Joe's reference class (so defined) is a singleton.⁴ Readers willing to concede this point—and I shall defend it properly in §3—may still be inclined to reject my analysis as a non-starter. 'According to your proposal', they will complain, 'the humdrum claim that Joe has a 43% chance of getting cancer entails the extraordinary claim that Joe has many actual, past-light-cone-matching doppelgangers. That is patently absurd.' My response is that the former claim is not humdrum. Indeed, were it not for the surprising success of a deeply puzzling physical theory—quantum mechanics—we would not even be discussing it.⁵

The structure of the paper is as follows. In §2, I clarify the analysandum, and set out a relativistic framework for handling chance talk. In §3, I sketch the analysis, and offer a preliminary defence. In §4, I articulate the analysis in detail. In §5, I argue that my version of finite frequentism evades many of the usual objections. One significant exception is the complaint that finite frequentism prohibits irrational chances; here I advocate bullet-biting. (I also defer discussion of chance laws to §8.) In §6, I show that my analysis makes possible a simple and satisfying explanation of the chance–credence link. That explanation is driven by a self-location indifference intuition with two signal strengths: (i) it is finitistic, so there are no doubts about the appropriate measure; (ii) it is indefeasible, thanks to the narrowness of the reference classes involved. In §7, I anticipate the objection that chances would be *unknowable* (at least in non-trivial cases) if my analysis were correct. In response, I sketch a reassuring confirmation-theoretic story. In §8, I argue that my analysis allows for the existence of chance laws. There is admittedly a countervailing intuition here: many readers will feel that chance regularities would be *mere coincidences* if my analysis were correct. I outline reasons for distrusting that intuition. Summing up in §9, I conclude that finite frequentism does not deserve its poor reputation: my analysis is,

⁴ Nor is it obvious that the class is finite. If it is infinite, then the analysis I defend implies that Joe's chance of getting cancer is either trivial or undefined (see §4).

⁵ 'Joe has a 43% chance of getting cancer' *would* be a humdrum claim if it was intended as shorthand for, e.g., '43% of the obese white male smokers surveyed in study X subsequently developed cancer', but 'chances' of that sort are not my topic.

in fact, a serious contender for the chance role.⁶

2 Target

I aim to give an analysis of non-deterministic chance in relativistic spacetimes. Most philosophers interested in chance have worked within a Newtonian framework. I break with that convention only because my particular strategy turns out to run more smoothly if non-trivial light cone structure is assumed from the outset. Before we get to the analysis, however, we need to marshal our categories, and here it does make sense to begin with the Newtonian case.

If we attach chances to (Newtonian) times, we can spell out the defining characteristic of *non-deterministic* chance as follows: if at time t an outcome is determined by history and law, then the chance of that outcome at t must be 1.⁷ An immediate implication is that the past is not chancy, for any outcome which has already occurred is determined by history, and hence, trivially, is determined by history and law. With minor tweaking, we can eliminate the implicit appeal to a physical arrow of time. For example, we can attach chances not to times simpliciter, but to *oriented* times. An oriented time is a logical construct comprising (i) a time and (ii) a formal tag designating one temporal direction as past. Let t range over oriented times, and let t 's *history* be the maximally specific intrinsic property of the region of spacetime t designates as its past. Then we can give a time-symmetric characterisation of non-deterministic chance: if at oriented time t an outcome is determined by history and law, then the chance of that outcome at t must be 1. We should prefer the time-symmetric version. It is more general, and it leaves open the possibility that the correct analysis of temporal asymmetry might refer, non-question-beggingly, to facts about chance.

How does this picture change when we introduce light cone structure? Suppose you wish to talk about the things that are nomically determined, or that are non-trivially chancy, 'now'. Presumably 'now' picks out a time slice of some sort. Let's tag the time slice with an orientation label that matches your psychological arrow of time. Then we can say that your past region is the union of past light cones of points in the time slice.⁸ What is the time slice's spatial extent? That will depend on context. If you are in a cosmology seminar, the 'now' you have in mind might conceivably be a time slice of the entire universe. In most other contexts you will probably intend to pick out something spatially restricted—a time slice of *you*, perhaps extended to include your immediate environment, or to embrace some relevant community (e.g. consider conversational contexts in which you and your interlocutors are implicitly committed to a shared referent for 'now'). And of course you may wish to associate chances with time slices of various other systems—radium nuclei, coin-tossing apparatuses etc.

Based on the discussion so far, it is tempting to say that all and only oriented time slices—whether of people, communities, radium nuclei, spacetimes or whatever—are potential chance bearers. However, a time slice in one coordinate system need not be a time slice in another, and it would be a disappointment if chances turned out to be tied to particular coordinate systems. A better bet is to exploit the natural coordinate-free generalization of the time slice concept. A *space-like surface* is

⁶ Might it also be the *strongest* contender? Of course I would like to think so, but I will not attempt to make that case here.

⁷ In other words, there is never any chance of a law being broken. In imposing this requirement I am following Lewis ([1994], p. 480).

⁸ I shall adopt the convention whereby a point is included in its own past light cone.

a surface of up to three dimensions such that no point on the surface lies in a light cone of any other point on the surface.⁹ Since they have no temporal ‘thickness’, we can think of space-like surfaces as *snapshots* of portions of the universe. That is a coordinate-free observation, but if we do introduce a time coordinate, the associated surfaces of simultaneity will be space-like. (For example, anything that deserves the label ‘person time slice’ will be a space-like surface.) My chance bearers will be *oriented space-like surfaces*. An oriented space-like surface is a logical construct comprising (i) a space-like surface and (ii) a formal tag designating one temporal direction as past.¹⁰ Non-deterministic chance can then be characterised as follows: if at oriented space-like surface s an outcome is determined by history (the maximally specific intrinsic property of s ’s past light cone union) and law, then the chance of that outcome at s must be 1. Call this the *nomic necessity principle*.

The principle is structurally similar to the time-symmetric Newtonian version above. Indeed, the only difference is that the relativistic chance bearer category can accommodate a geometrically richer range of past regions. One consequence of that difference is worth flagging. In a (time-symmetrically) deterministic Newtonian spacetime, at any oriented time t , history and law determine everything; thus, there can be no non-trivial non-deterministic chances; thus, the hallmark of *deterministic* chance will be ‘objective chance values other than 0 and 1, in a deterministic world’ (Schaffer [2007], p. 113). In relativistic spacetimes, things are not so simple. Let s_0 be your present person time slice, oriented to align with your psychological arrow of time. (Notice that s_0 is spatially localised: it is an oriented time slice of *you*, not of the universe.) Suppose that s_0 ’s history fails nomically to determine whether you will be eaten by bears aged 90. The corresponding chance, at s_0 , may then be non-trivial. Now ask: does our supposition rule out determinism-compatible stories about your future death? The answer is that it does not. It may yet be that your future death has a complete set of determining causes on every space-like section through *its* past light cone (‘past’ here is still the temporal direction designated by s_0). Indeed, it may yet be that your future death has a complete set of determining causes on every space-like section through either of its light cones, and the same may go for every other event in the spacetime. In other words, the nomic necessity principle allows for the possibility of non-trivial non-deterministic chances in *deterministic* spacetimes. Readers who think they detect a whiff of absurdity here need to divest themselves of Newtonian intuitions. In particular, the Newtonian principle that deterministic spacetimes leave no room for non-determining *histories* does not generalise to the relativistic case; and where there are non-determining histories there may be non-trivial non-deterministic chances. ‘But surely a scientific theory might imply indeterminism by postulating chances of a certain kind?’ Yes indeed—and the kind in question is a subspecies of non-trivial non-deterministic chance. Its definition is straightforward but technical, so I relegate it to a footnote.¹¹

⁹ Light cones extend in both temporal directions, so this definition is time-symmetric as well as frame-invariant.

¹⁰ The tagging is possible only if the surface is *orientable*, that is, only if a consistent past/future designation can be made for the surface as a whole. Exotic spacetimes may contain space-like surfaces that are not orientable. My analysis will imply that such surfaces cannot be associated with chances.

¹¹ Let s be any oriented space-like surface. A point p lies in s ’s *future domain of dependence* $D^+(s)$ iff every causal curve through p that is maximally extended in the direction of s ’s past intersects s . Put less formally, the condition is that p ’s past light cone has a cross-section that lies wholly within s . Processes indeterministic in at least one temporal direction exist if, for some s , there is a non-trivial chance at s that $D^+(s)$ has some intrinsic property. My definition of $D^+(s)$ is largely standard (see, e.g., Torretti [1996], p. 124). Notice, however, that since s designates its own ‘past’ and ‘future’ directions, my

From now on, I will use ‘chance’ without qualification to denote the analysandum described above. Thus, to repeat, chances attach to oriented space-like surfaces, and they respect the nomic necessity principle. There are certainly other concepts to which the label ‘chance’ could reasonably be applied; these concepts are not my topic.

Reflexive attributions of present chance (‘my present chance of getting cancer is 43%’; ‘the chance here, now, of that nucleus decaying in the next hour is 12%’) are ground-zero for discussions of the chance–credence connection (§6). The chance bearer in any such case will be a time slice of the relevant agent, oriented to align with the agent’s psychological arrow of time. From now on, I will speak simply of *person slices*, and I will make the following simplifying assumption: the context of a reflexive attribution—agent’s intentions included—precisely determines the relevant person slice. I do not assume that agents always (or indeed ever) *know*, precisely, the geometry of these slices. Readers may nevertheless object that my assumption is unrealistic. In lieu of a defence, I offer an excuse: vagueness worries about self-location indexicals are not *especially* problematic for finite frequentists. I shall set them aside.

3 Sketch and Skirmish

If chances are actual frequencies, and chance bearers are oriented space-like surfaces, then the chance at s of some particular outcome is the frequency with which some particular *property of oriented space-like surfaces* is instantiated in s ’s reference class. For example, the hypothesis that you are going to be eaten by bears aged 90 can be re-described as the hypothesis that s_0 is a person slice of someone who is eaten by bears aged 90. The frequency with which that property (person slice of person who is eaten by bears aged 90) is instantiated in s_0 ’s reference class is then the chance, at s_0 , of the unfortunate outcome. Similarly, the frequency with which the property *such that seven is prime* is instantiated in s_0 ’s reference class—one, on any sensible definition of ‘frequency’—is the chance, at s_0 , that seven is prime.

The finite frequentist’s central task is to explain how chance bearers pick out reference classes. I trailed my proposal in §1, but before we return to it, let us step back and consider more carefully the anatomy of the reference class problem. Good analyses of chance need to respect certain platitudes. One platitude connects chance with rational credence. Let Q be the proposition that you will live to be 100. Let X be the proposition that your present chance of living to be 100 is 0.15. Whatever your evidence, your present credence function C must, on pain of irrationality, satisfy $C(Q|X) = 0.15$.¹² However, if some version of finite frequentism is correct, X reports a mere *actuarial statistic*: 15% of some class of actual person slices are slices of people who live to be 100. It is no news that actuarial statistics might influence a rational person’s confidence in his own prospects. What is harder to understand, at least at first sight, is how a *single* such statistic could screen out all other rational-credence-influencing factors. Suppose you spend your days exercising, eating salad, and pursuing fulfilling but stressless projects. Might you not reasonably expect to beat the actuary? Won’t you (probably) outlive your reference classmates? The finite frequentist must deny this. It is clear what his strategy will be: he will say that your

definition is time symmetric.

¹² There is some small print: the constraint must be satisfied *unless* $C(X)=0$, in which case $C(Q|X)$ is not defined.

reference classmates share your healthy habits. Indeed, he will insist that they resemble you in whatever respects are necessary to ensure that your conditional credence $C(Q|X)$ cannot rationally depart from the actuarial odds. But what are these rational-credence-determining respects? The reference class problem would appear to be *epistemologically* challenging. Another platitude is encoded in the nomic necessity principle. If physical laws and facts about history together entail that you will live to be 100, then your present chance of reaching 100—and thus the frequency of eventual-centenarian slices in your reference class—had better be 1. So the reference class problem is *metaphysically* constrained too.

Still, if the finite frequentist's reference class problem were *simply* the problem of accommodating the chance–credence and nomic platitudes, it would rather obviously invite the solution trailed earlier. To wit: the reference class that defines your present chance of reaching 100 is the class of person slices whose past light cone unions are qualitatively identical to (i.e., are perfect intrinsic duplicates of) yours. Call these person slices your *cone counterparts*. Your present chance of reaching 100 would, on this suggestion, be the frequency of eventual-centenarian slices among your cone counterparts. More generally, the chance that x has property A would be the frequency of A s among x 's cone counterparts. This analysis will need some polishing (§4), but we can see already that it promises to respect the chance–credence and nomic platitudes. Your cone counterparts resemble you in *every* respect that can possibly be relevant to your appraisal, now, of the hypothesis that you will live to be 100. It does not seem unreasonable to suppose that an actuarial statistic relative to this very narrow reference class might indeed screen out all other rational-credence-influencing factors (for a full discussion, see §6). And whatever else laws of nature are, they are at least true universal generalizations. Thus, if law and facts about your history entail that you will live to be 100, it follows that *all* your cone counterparts are slices of people who will live to be 100, and hence that the chance—so analysed—is unity.¹³

Unfortunately, the reference class problem is not simply the problem of accommodating the chance–credence and nomic platitudes. We must also accommodate the intuition that chances can take non-trivial values. Here we face a difficulty: it is natural to worry that cone counterpart classes will always be singletons, and hence that the associated frequencies will be limited to the trivial values 0 and 1. Now, it will (I hope) be agreed on all sides that larger cone counterpart classes are *possible*. No metaphysical principle prohibits them. What critics will claim is that hypotheses postulating their *actual* existence are incredible, and, therefore, that my analysis has the ridiculous implication that hypotheses postulating non-trivial chances should be dismissed out of hand. My aim in the next few paragraphs is to show that this charge is unfair. I begin by arguing—rather briskly—that there are no compelling a priori grounds for denying the existence of non-trivial cone counterpart classes. Then I argue that there are no compelling empirical grounds, either. Along the way, I anticipate and block a couple of other, more modest, critical strategies.

The a priori argument for denying the existence of non-trivial cone counterpart classes goes like this: for the actual world to contain qualitatively identical but numerically distinct light cones would require an *extraordinary coincidence*; it follows that one ought to assign negligible credence to the hypothesis that such replicas exist. Both the premise and the inference are problematic. The inference depends upon an implicit, and manifestly unsafe, assumption about the size of the actual world. After all, given a big enough world, it would be extraordinary if extraordinary coincidences

¹³ Here I take for granted that the frequency of A s in a class *all* of whose members are A s is 1, irrespective of cardinality. I give a formal definition of 'frequency' in §4.

didn't occur. The premise would be false—replicas would not be coincidences—if the laws of nature imposed upon the actual world a ‘Grand Pattern’ with light-cone shaped (or at any rate very large) motifs. It is difficult to see how such laws could be ruled out on a priori grounds. (I will have more to say about laws in §8.)

Before we can properly assess the empirical case against non-trivial cone counterpart classes, we need to consider a more modest a priori argument. This argument purports to show that the geometry of spacetime puts strict limits on the range of possible cone counterpart classes. It starts from the superficially plausible premise that, necessarily, any two light cones overlap unless they point in opposite temporal directions (Figs. 1 and 2). Call this the *overlap thesis*. An immediate consequence is that *replica* light cones must either point in opposite temporal directions or each satisfy very restrictive internal symmetry constraints. (Notice that a ‘generic’ light cone—one that does not have strong internal symmetries—cannot overlap with a replica of itself.) If I conceded this claim but stuck to my analysis, I would have to say that non-trivial chances require *either* strongly symmetric pasts *or* massive violations of the familiar arrow of time; furthermore, since we have only two temporal directions to play with, I would have to say that chances other than 0, $\frac{1}{2}$ or 1 require strongly symmetric pasts, period. That is an unattractive position. Luckily, I do not need to defend it, because the overlap thesis is *only* superficially plausible. To begin with what some will regard as a cheap shot, it is false if worlds need not be spatiotemporally unified: a world consisting of many disjoint spacetimes will contain many mutually disjoint light cones. Whether such worlds are possible is a controversial philosophical question. I would like to believe that they are (more on this later), but for the sake of argument, let us suppose that they are not. We should still reject the overlap thesis, because disjoint co-oriented light cones occur in several possible spacetimes. Fig. 3 is a cartoon example; grown-ups may prefer the Friedmann-Lemaître-Robertson-Walker models of open universes found in standard cosmology textbooks (e.g. Ellis et al. [2012], p. 174-6).

The a priori arguments are unthreatening as they stand. Nonetheless, we should ask whether the second, more modest, argument can be given an empirical makeover. Does the scientific evidence suggest that co-oriented light cones in *our* spacetime always overlap? As it happens, the inflationary models presently favoured by cosmologists imply that past light cones of points in the *observable* universe—i.e., in our past light cone—do always overlap (ibid., pp. 240–1). (Future light cones of points in our past light cone overlap by definition: we fall in all of them.) However, the word ‘observable’ is crucial here: these models do not imply the *unrestricted* claim that co-oriented light cones in our spacetime inevitably overlap. Our spacetime may, for all we know, be much larger than our past light cone; indeed, recent work in the philosophy of cosmology has stressed the extent to which the global structure of our spacetime is underdetermined by the observations we can (even in principle) make.¹⁴ Thus, global cosmology—insofar as it is empirically grounded—seems unlikely to be a stumbling block.

We are now ready to assess the most obvious empirical argument against the existence of non-trivial cone counterpart classes. It goes like this: astronomical surveys have detected no hints of replica light cone structures; therefore, they probably don't exist. It faces the (by now familiar) objection that the actual world may be much larger than its presently surveyable portion. However, some may think the objection smacks of special pleading. ‘Surely’, they will say, ‘our failure to observe *unicorns* counts

¹⁴ See, e.g., (Ellis [2007], pp. 1203–6, 1230–2; Manchak [2009]; Butterfield [forthcoming]).

strongly against the hypothesis that unicorns exist. Why should the same not go for replica light cones? I have two responses. First, our failure to observe unicorns does not count *strongly* against the hypothesis that they exist. We have no idea how many planets support life. If the number were sufficiently vast, it would not be surprising if white, single-horned, horse-like creatures existed on some of them. Second, the empirical case against replica light cones is much weaker than the empirical case against unicorns, because it is vulnerable to a sampling bias objection. Light cone overlap worries boomerang here to my advantage. If the inflation cosmologists are right that co-oriented light cones of points *in our past light cone* inevitably overlap, then our past light cone is an especially unpromising place to look for traces of non-trivial cone counterpart classes.¹⁵ Astronomers' failure to detect replica light cones, then, is less telling than zoologists' failure to detect unicorns. To the question 'what would our world look like if most oriented space-like surfaces had non-trivial cone counterparts?', an entirely plausible answer is 'just as it in fact looks'.¹⁶

I conclude that my proposed solution to the reference class problem does not trivialise chance. That claim would have been much harder to defend had I insisted on working within a Newtonian framework. (Newtonian pasts, being spatially unrestricted, *must* overlap unless they are oppositely oriented or belong to disjoint spacetimes.) Luckily, we have excellent grounds for believing that we live in a relativistic world, so my analytical framework is realistic as well as convenient.

Critics prepared to concede that I have dealt with the triviality worry may wish to press two follow-up objections.¹⁷ *First objection*: any adequate analysis of chance ought to accommodate without undue strain a wide variety of chance values. In particular, it ought to accommodate fractions with irreducibly large denominators, e.g. 46.347% (46347/100000) or 3.63×10^{-10} (363/10¹²). If my analysis is correct, such chances require large cone counterpart classes. That looks like bad news: even the most broadminded reader may balk at the suggestion that she has *trillions* of past-light-cone-matching doppelgangers. *Second objection*: suppose we grant—just for the sake of argument—that past-light-cone matching doppelgangers are two-a-penny. We thereby grant that the *actual* world is friendly to my analysis. So what? There remains the intuition that finely-graded chances are metaphysically possible in small worlds, and my analysis cannot save this intuition. Let me take the objections in turn. The first boils down to a hunch about the size of the actual world—to wit, that cone counterpart

¹⁵ Suppose instead that the inflation cosmologists are wrong. It remains the case that we could detect replicas of our own past light cone only if our past light cone overlapped with replicas of itself. Thus, there is a weaker version of the sampling bias objection whose fortunes are not tied to those of any particular cosmological theory.

¹⁶ So far I have not mentioned the 'many-worlds' interpretation of quantum mechanics (MWI). This may have struck some readers as odd. MWI is, after all, a scientifically respectable view according to which the actual world is very much larger than its presently surveyable portion. (The many 'worlds' of MWI are not distinct *possible* worlds in a metaphysician's sense, but parts of the actual world.) However, I cannot plausibly argue from MWI to my analytical claim about chance. There are two obstacles. First, MWI is unfriendly to any attempt to count reference classmates, because its individuation of 'worlds' is to some extent arbitrary (Greaves [2007], pp. 120–121). Second, certain quantum mechanical chance ascriptions are incompatible with my analysis. I think I can shrug this fact off (§5.4), but it would be pushing my luck to claim that quantum theory as it stands, on any standard interpretation, provides me with positive support. Never mind! Making clear sense of quantum mechanics is notoriously difficult. I respect its authority to the extent that I take seriously the hypothesis that non-trivial chances (a) exist and (b) sometimes take values that can be calculated, to high degrees of accuracy, by physicists. However, I do not think I am under any great pressure to adopt a quantum mechanical worldview. (That said, I would of course be delighted if some future iteration of or successor to quantum theory turned out to be more friendly to my analysis.)

¹⁷ I am grateful to an anonymous referee for pressing both.

class cardinalities might perhaps run to three, or fifty, or a thousand, *but no higher*. Precisely where the line is being drawn is not clear, and it is very difficult to see what could motivate any particular choice. (We saw some apparently principled reasons for drawing the line at cardinality one above, but those reasons turned out to be unconvincing.) It seems to me that we have little reason to trust such hunches. I should acknowledge, though, that my analysis *would* be unattractive if the hunches could be supported by convincing arguments.¹⁸ The second objection is more damaging. There surely is a pre-theoretical intuition to the effect that finely-graded chances are possible in small worlds. Since this is an intuition about chance, and not about the contingent structure of the actual world, philosophical analysts must take it seriously; and I cannot save it. That is a real cost, but I do not think it is a deal-breaker. *Any* Humean analysis of chance will conflict with the intuition that there could be non-trivial chances without the allegedly necessary chance-making patterns. Consequently, any Humean analysis will face some version of the small world objection. To see this, consider an almost-empty world, containing one or two radium nuclei and nothing else. Pre-theoretical intuition rules that arbitrary decay chances are possible. Humeans must demur: given the limited supply of raw material in this world, few (if any) chance-making patterns can be realised. Of course, my chance-making patterns demand much more raw material—more spacetime and more stuff—than the patterns proposed by, say, David Lewis. Nevertheless, the small world intuitions we have to resist are recognisably intuitions of the same type, and it would be against the spirit of the Humean project to regard either as sacrosanct.

4 Articulation

In §1–3 I sketched a finite frequentist analysis of chance and, I hope, offered just enough argumentative support to motivate a more thorough discussion. In this section, I complete the articulation of the analysis.

Oriented space-like surfaces serve as the bearers of chance and the constituents of reference classes. To save ink, I shall from now on refer to them as *objects*.

Definition. An *object* is an oriented space-like surface.

Objects designate their own past temporal directions. Thus, each object picks out a past light cone union—the union of past light cones of points in the object. To define reference classes, we must specify an equivalence relation over objects. The previous section’s sketch suggests that the relation should be *qualitative identity of past light cone unions*. That is very nearly adequate, but it has undesirable implications in certain special cases. I will pause to explain the problem, and then fix it; but the issue is a narrowly technical one, and readers who skim this paragraph will be able to follow the rest of the paper without difficulty. Consider a spacetime bound between qualitatively distinct initial and final Cauchy surfaces (maximally extended space-like surfaces). Attach orientations to these bounding surfaces so that each lies in the other’s designated past. We now have two objects, each of which picks out the same past light cone union—to wit, the whole spacetime. Since numerical identity implies qualitative

¹⁸ Although I do not believe that the requisite arguments will be forthcoming, they are conceivable. Perhaps future metaphysicians will prove (or will accept that David Lewis has proved) that the actual world is spatiotemporally unified, and future cosmologists will—*pace* Ellis, Manchak and Butterfield—put an upper bound on the size of our own spacetime. Taken together, these advances could rule out large cone counterpart classes.

identity, our first-draft equivalence relation counts the oriented Cauchy surfaces as reference classmates. That is very much against the spirit of the analysis; after all, the *surfaces* are qualitatively distinct. The problem here is that an object's past light cone union, conceived simply as a set of spacetime points, need not encode orientation information. So, once again, I will appeal to a logical construct:

Definition. An object's *cone* is the oriented union of the past light cones of points in the object. The orientation is inherited from the object.¹⁹

The point of this definition is to support another:

Definition. Two objects are *cone counterparts* iff they have qualitatively identical cones.

Our pathological case is now blocked. The spacetime we imagined had qualitatively distinct initial and final Cauchy surfaces. Call them *blue* and *red*. The oriented Cauchy surfaces defined above pick out identical light cone unions, but they pick out (numerically and qualitatively) distinct cones: one is oriented red-to-blue, the other blue-to-red.

With that epicycle in place, I can formally define my reference classes:

Definition. An object's reference classmates are, precisely, its cone counterparts.

It remains to define 'frequency'. In a finite reference class, the frequency of a property is just the fraction of objects in the class that have the property in question. I shall extend the concept to embrace two trivial infinite cases: the frequency of a property is zero if *no* object in the (finite or infinite) reference class has the property, and unity if they *all* do. Let us tendentiously denote the frequency of *As* in *x*'s cone counterpart class $Ch_x(A)$:

$$Ch_x(A) \stackrel{\text{def}}{=} \begin{cases} \frac{\text{Number of } A\text{s in } x\text{'s cone counterpart class}}{\text{Total number of } x\text{'s cone counterparts}} & \text{(if } x\text{'s cone counterpart class is finite)} \\ 0 & \text{(if no cone counterpart of } x \text{ is an } A) \\ 1 & \text{(if every cone counterpart of } x \text{ is an } A) \end{cases}$$

My proposal is that this quantity is chance. $Ch_x(A)$ is the chance that object *x* has property *A*; furthermore, any chance is the chance that some object has some property. When it is necessary to distinguish analysans and analysandum, I shall refer to $Ch_x(A)$ as a *cone counterpart frequency*.

If *x*'s cone counterpart class is finite, then $Ch_x(A)$ is sure to be defined. If *x*'s cone counterpart class is infinite, then $Ch_x(A)$ will be defined only in trivial cases.²⁰ Since non-trivial infinite cone counterpart classes are metaphysically possible, the analysis commits me to defending the metaphysical possibility of chance 'gaps': <object, property> pairs that do not pick out any chance. That is not a commitment I

¹⁹ Notice that a union of past light cones of points *in an object* is necessarily orientable.

²⁰ Recall that a fraction cannot have an infinite denominator.

welcome, but it is surely not a reductio. The philosophically urgent task is to explain how there *could* be well-defined chances, not—I take it—to demonstrate that there could never fail to be.²¹ (I admit that my analysis would cease to be appealing if we discovered that all or most cone counterpart classes were *actually* infinite. However, it is very difficult to see how we could make such a discovery.)

Ned Hall once complained that ‘no reductionist has in fact ever provided an exact recipe that would show how categorical facts fix the facts about objective chance’ (Hall [2004], p. 111). I think I have provided just such a recipe.²²

5 Standard Worries

Finite frequentism, Alan Hájek tells us, ‘runs afoul of many important intuitions’ ([1997], p. 213); that is why it is ‘about as close to being refuted as a serious philosophical position ever gets’ (ibid., p. 226). Most versions of finite frequentism will, of course, run afoul of the chance–credence and nomic platitudes. I have already indicated how my particular choice of reference class allows me to respect them. I need to say a good deal more about rational credence, but before turning to that crucial topic, I want to deal with the other supposedly threatening intuitions.

5.1 Frequency tolerance

Let us stipulate that a *fair* coin’s chance of landing heads when tossed lies, always, in a narrow interval around 0.5. (It is implausible that real coins are chancy. Never mind!) Now, suppose you toss a fair coin 1,000 times. You would expect to get roughly equal numbers of heads and tails. However, it is of the essence of chance that you *might* not. Indeed, the coin might land heads every time. Generalizing, chances would appear to be *frequency tolerant*. Although a rational agent will allow the putative chances to guide his expectations, chances do not put *logical* constraints on frequencies.²³ That, at any rate, is the standard story, and if it’s right, finite frequentism stands refuted: tinkering with reference classes won’t help.

The standard story can be resisted. It begins well: the intuition about the coin tossing experiment is very strong. However, the generalization—the claim that chances put no logical constraints on frequencies *of any kind*—is problematic. Why are we inclined to accept it? The main reason, I suggest, is that it looks like an innocent summary of our intuitions about particular cases. It isn’t. To demonstrate this, it will be sufficient to show that my version of finite frequentism respects the coin tossing intuition, and any similar intuition about any imaginable experimental set-up. That my analysis nevertheless clashes with the generalization is a worry only to the extent that

²¹ Certainly the gappiness raises no mathematical difficulty. Associated with any object x is a probability space (Ω, F, P) . Ω is x ’s cone counterpart class, F is a set of subsets of Ω , and P is the probability function that assigns to any member S of F the cone counterpart frequency $Ch_x(A_S)$, where A_S is any property whose extension within Ω is S . F takes one of two forms. If Ω is finite, then F is the power set of Ω : the set of all Ω ’s subsets. If Ω is infinite, F comprises just the empty set and Ω itself, and we get a very boring probability space. But either way, all the Kolmogorov axioms are satisfied, including (though only vacuously) countable additivity.

²² *Lewisian* reductionists can be criticised for failing fully to define simplicity, strength, fit and the exchange rates governing trade-offs between these virtues. Lewisian analyses of chance do not provide ‘exact recipes’ in Hall’s sense.

²³ Except, perhaps, in trivial cases (chances of 0 and 1), but even that is controversial. If there are continuous chance distributions, then outcomes with zero chance are logically possible. My analysis does not allow for such distributions.

there is a *direct* intuition in favour of the generalization: an intuition that cannot be appeased by reassuring stories about coins and dice and radium nuclei. Perhaps there is such an intuition, but if so it strikes me as (a) weak and (b) just the sort of intuition one should be willing to give up in exchange for an otherwise satisfying analysis.

I want to show that my analysis allows for the possibility that a fair coin happens to land heads every time it is tossed. As usual, we face the difficulty of finding objects to stand in for ‘times’; as usual, the appropriate choice will depend on details of the case and the context. Since we are not presently concerned with the credence functions of rational agents, let us take the objects to be (suitably oriented) pre-toss time slices of some particular coin placed in a coin-tossing apparatus. Suppose each of these objects belongs to a cone counterpart class roughly half of whose members instantiate the ‘next toss lands heads’ property. (If my analysis is correct, we are thereby supposing that the coin is fair.) Crucially, each pre-toss time slice will belong to a different cone counterpart class. The stable chance of the coin landing heads puts strong constraints on the statistics of these classes, but it puts no constraint whatsoever on the individual coin’s *sequence* statistics. In particular, the coin might land heads every time. I hope it is obvious that I will be able to tell a similar story about any other repeatable trial. If reference classes are cone counterpart classes, then reference class statistics and sequence statistics are logically independent.²⁴

Some readers will complain that I have secured sequence-level frequency tolerance at ruinous cost. Prima facie, the ‘stability’ of my ‘fair’ coin’s chance of landing heads is just a coincidence. That the cone counterpart frequencies have hovered near 0.5 so far (though how would we know?) would appear to provide no reason whatsoever for expecting them to remain near 0.5 in the future. I will address the important issues raised by this objection in §7 and §8. Here, in the meantime, is a brisk preview. Nobody ought to insist that it is *metaphysically* necessary that any particular coin’s chance of landing heads remains stable over time. The attractive claim is that stability is sometimes *nomic* necessary. For example, one might conjecture that coin-tossing set-ups with certain locally-ascertainable properties are always fair, and that this is a consequence of the laws of nature. My analysis is undoubtedly compatible with the existence of chance regularities; the crucial question is whether it is compatible with their being law-like. Some readers will think it intuitively obvious that the answer is ‘no’. I will argue that there are principled grounds for distrusting this intuition. Come what may, though, I remain immune to one of the best-known objections to finite frequentism. ‘If the frequentist has his way’, Hájek complains, ‘we can’t say that the chance of the coin landing heads really was 1/2, but that there was an unusually high proportion of tails in the actual sequence of tosses. And yet that could be a very natural thing to say’ ([1997], p. 219). That could indeed be a natural thing to say, and I can say it.

Closely related to frequency tolerance worries are worries about *undermining futures*. A future is said to be undermining if it presently has a non-zero chance of occurring, but is such that its *actual* occurrence would require the present chance to be different. Good analyses of chance ought not to allow for such absurdities. They ought to respect Bigelow, Collins and Pargetter’s Basic Chance Principle (BCP): any outcome whose chance exceeds zero is compossible with that chance and with history to date (Bigelow et al [1993], p. 459). Reassuringly, my analysis entails the BCP. Suppose $Ch_x(A)=p>0$. An immediate consequence is that at least one of x ’s cone counterparts is an A . Thus, being an A is compossible with x ’s history to date and x ’s

²⁴ Again, except in trivial cases. My analysis implies that a coin whose chance of landing heads is always zero will *never* land heads.

cone counterpart frequencies. In fact, I can weaken the BCP's antecedent and strengthen its consequent. To guarantee the presence of an A among x 's cone counterparts, it is sufficient that $Ch_x(A)$ exceed zero *or be undefined*; and since cone counterparts are worldmates, x 's cone counterparts are subject to the same laws as x . Thus, we have the following principle: if $\neg(Ch_x(A)=0)$, then being an A is compossible with x 's history, x 's cone counterpart frequencies, and the laws that hold in x 's world.²⁵ (I am treating ' $Ch_x(A)$ ' as a Russellian definite description, so that the principle's antecedent is true when $Ch_x(A)$ is undefined.)

5.2 Leibniz's dictum

Leibniz thought that probability was a kind of 'graded possibility'.²⁶ We are interested specifically in chance, but the thought remains appealing. If a coin has a 10% chance of landing heads, then it will possibly land heads. If it has a 70% chance, one is tempted to say that the possibility is *greater*. Of course, neither metaphysical nor nomic possibility comes in degrees, so the temptation is problematic. Nonetheless, one might hope that a good analysis of chance would concede something to Leibniz. Finite frequentism seems destined to disappoint: one who identifies chances with *actual* frequencies would appear to be rejecting both the letter and the spirit of Leibniz's dictum.

On reflection, things are not quite so clear. Consider—in fact, for a page or so please *entertain*—David Lewis's analysis of metaphysical possibility *de re*: it is *de re* possible for an individual x to be an F iff x has, in some world or other, a *counterpart* that is an F (Lewis [1986a], p. 8). Counterparthood is a relation of qualitative similarity. Precisely which relation is something Lewis leaves vague: different precisifications will be appropriate in different contexts. Strictly speaking, *cone* counterparthood cannot be among the admissible precisifications. That is because Lewis's counterpart relation is defined over concrete entities, while cone counterparts are logical constructs (thanks to their orientation tags). This categorical mismatch is a purely technical inconvenience. We can rid ourselves of it by tweaking Lewis's formalism. Let us stipulate that oriented space-like surfaces are to be counted among the *res*, i.e., counted among the bearers of *de re* possibility. It is then not much of a stretch to describe the following as a Lewisian formula:

It is *de re* possible *given history* for object x to be an A iff x has (in some world or other) a cone counterpart that is an A .

Modifying the formula to pick out more restricted *de re* modalities is straightforward. It is *de re* possible *given history and law* for object x to be an A iff x has, in some *nominally accessible* world, a cone counterpart that is an A . It is *de re* possible *given history, law and chance* for object x to be an A iff x has, in some *nominally accessible* world, a *chance-matching* cone counterpart that is an A . And so on. A *de re* formulation of the BCP is now analytic:

for any object x , if $\neg(Ch_x(A)=0)$, then x is possibly A ,

where the possibility in question is possibility given history, cone counterpart

²⁵ Schaffer ([2007]) argues persuasively that the BCP's consequent indeed ought to be strengthened to hold chances and laws fixed.

²⁶ 'Probabilitas est gradus possibilitas' (cited in Hacking [1971], p. 345).

frequencies and law.

That is all very well, but possibility given history, cone counterpart frequencies and law cannot be the possibility that is *graded* by $Ch_x(A)$. $Ch_x(A)$ surveys only x 's worldmates; the possibility it grades must be similarly restricted. Let us call the restricted possibility in question L-possibility ('L' for 'Leibniz'). x is L-possibly A iff x has a cone counterpart *in its own world* that is an A . L-possibility entails possibility given history, law and cone counterpart frequencies (whence the BCP); the converse, of course, does not hold. 'But L-possibility does not deserve its name! Whether actual objects are L-"possibly" A is settled by how things *actually* are!' The observation is correct, but the complaint is unfair. If the Lewisian picture is right, then L-possibility really is a species of metaphysical possibility *de re*.²⁷ Moreover, L-possibility is *non-trivial*, in the sense that a non- A object can be L-possibly A .

Let us take stock. My analysis of chance, when conjoined with a (lightly customized) Lewisian account of metaphysical possibility *de re*, does a good job of saving certain modal intuitions. In particular, $Ch_x(A)$ becomes, literally, a graded possibility, with the extreme values 0 and 1 corresponding to L-impossibility and L-necessity respectively. Readers inclined to accept Lewis's modal realism will naturally regard this as good news. However, many readers won't be inclined to accept modal realism, and I do not wish to endorse it myself. (In fact, for reasons that will soon be clear, I hope it is false.) Do *I* have any good news to celebrate? I think I do. There might be non-Lewisian ways to make an honest metaphysical modality of L-possibility, but that is a topic for another occasion.²⁸ There is certainly a viable fudge: even if $Ch_x(A)$ is not, strictly speaking, a graded possibility, we can appeal to the virtues of modal realism—virtues that fall short of truth—to argue that $Ch_x(A)$ is rather *like* a graded possibility. I shall leave this as a sketch, because the details don't much matter. If you are a non-believing admirer of modal realism, then you will acknowledge its allure. If you acknowledge its allure, you should agree that we can appease Leibniz (up to a point) by appealing to cone counterparts in lieu of genuine possibilia.

5.3 Single case chance

Finite frequentism entails that chance-bearers belonging to singleton reference classes can bear only trivial chances: 0 or 1. This is standardly seen as counterintuitive. After all, '[m]any experiments are most naturally regarded as being unrepeatable—a football game, a horse race, a presidential election, a war, a death, certain chancy events in the very early history of the universe. Nonetheless, it seems natural to think of non-extreme probabilities attaching to some of them' (Hájek [1997], p. 221). These examples pose no special difficulty for me. To accommodate non-trivial chances of any sort, I have to posit replica light cones; and replica light cones may host replica horse races etc.

But there is a more troublesome version of the single case problem. Let's hear from Hájek again:

Certain statements are 'single case' in virtue of their very logical form: for example, universal generalizations and existential claims. Some people think

²⁷ Assuming, with the later Lewis, that numerically distinct counterparts may be worldmates (Lewis [1986a], p. 232, n. 22).

²⁸ L-possibility at least has the right *formal* properties: if x is A , it is L-possibly A ; if x is L-necessarily A , it is A ; x is L-possibly A iff x is not L-necessarily not- A .

that non-trivial (objective) probabilities attach to such statements—as it might be, ‘the probability that all ravens are black is 0.9’, or ‘the probability that there exist tachyons is 0.1’. If there is sense to be made of such probabilities, then it is not the frequentist who can make it, for such statements only get one opportunity to be true or false. (Ibid., p. 221.)

The worry now is not that certain reference classes must be singletons, but rather (if I may put it in my terms) that certain properties will be instantiated either by all objects or by none. Consider the property *being such that tachyons exist*. If tachyons exist, every object has that property; if tachyons do not exist, none do. Thus, the chance that tachyons exist must—it seems—be either 1 or 0. Bullet-biting may be an option here: as Hájek implicitly acknowledges, the intuition at stake is not overwhelmingly strong. However, my preferred strategy is to meet the intuition half way by appealing to restricted quantifiers. Roughly: the sentence ‘tachyons exist’ should, at least in some contexts, be taken to express the indexical proposition that tachyons exist *in this cosmic vicinity*. That is vague as it stands, but if readers are willing to grant the assumption that the world might not be spatiotemporally unified, there is a natural precisification: identify cosmic vicinities with *complete spacetimes*.²⁹ On this proposal, a cosmologist who claims that there is a 10% chance that tachyons exist speaks truly iff 10% of her cone counterparts live in tachyon-containing spacetimes.³⁰ And, mutatis mutandis, I could handle other ‘single case’ examples—chancy universal generalisations, chancy Big Crunches—in the same way. Going this route would force me to reject David Lewis’s identification of possible worlds with maximal individuals (individuals spatiotemporally related to all and only their own parts), for Lewis’s view makes it a necessary truth that the world is spatiotemporally unified.³¹ The dialectical pressure I face here is nothing to celebrate, but nor is it cause for great alarm: Lewis’s view is not widely accepted, and among the reasons for its unpopularity is the very implication I find inconvenient.³²

5.4 Missing values

If any version of finite frequentism is true, then all irrational chances ‘go missing’ (Hájek [1997], p. 224): it will be logically impossible for chances to take values such as $1/\sqrt{2}$. Is this a problem? There is perhaps an a priori intuition that chances can take irrational values. (The Kolmogorov axioms permit it.) On the other hand, those who remember schoolroom probability exercises may be tempted by the thought that chances are necessarily fractions. These are brute hunches; probably neither should be given much weight. However, I must also contend with the fact that irrational chances are postulated by quantum mechanics, one of our best-confirmed physical theories. This undoubtedly counts against finite frequentism, but the blow is not fatal. Quantum mechanics is plagued by notorious conceptual difficulties. We do not know—and

²⁹ See (Bricker [2001]) for a defence of the suggestion that there might, for all we know, be many actual spacetimes. (Bricker speaks of *many actual worlds*. I prefer to reserve ‘actual world’ for the totality of everything actual.)

³⁰ Of course, in typical scientific contexts, one who utters the *sentence* ‘there is a 10% chance that tachyons exist’ will not be making a chance claim at all; rather, he will be recommending a particular credence assignment. This is a useful reminder that chancy existentials are a rather peculiar case.

³¹ Some readers may wonder whether I have another option: embrace modal realism, but drop the requirement that cone counterparts be worldmates. I cannot sensibly go this route. Lewis’s pluriverse is (wildly) infinite, and infinite cone counterpart classes rule out non-trivial cone counterpart frequencies.

³² See, e.g., (Bigelow and Pargetter [1990], pp. 189–92).

arguably should not believe—that it is true as it stands. What we must acknowledge is that it has an outstanding empirical track record. However, since irrational numbers can be approximated arbitrarily closely by rationals, inferences from the empirical track record of quantum mechanics to the existence of irrational chances are not very compelling.

5.5 Explanation

Hájek takes finite frequentism to be an operationalist doctrine: it defines chance ‘in terms of ... the results of trials of the relevant sort’. But chances are also supposed to *explain* the results of such trials, and ‘you can’t explain something by reference to itself’ (Hájek [1997], p. 219, original emphasis).³³ My version of finite frequentism is so distant from operationalism that readers may worry that it makes chances unknowable. I will tackle that objection in §7. What matters here is that explanations of experimental data that appeal to cone counterpart frequencies are not circular: explanans and explanandum are manifestly distinct.

Non-circularity is a low bar. Can a positive case be made for the explanatory relevance of cone counterpart frequencies? I don’t expect to be able to satisfy everyone: the concept of explanation is vague, and we have nothing approaching an uncontroversial analysis.³⁴ That said, a *prima facie* case is implicit in the chance–credence story I sketched in §3. Consider the hypothesis that 95% of your cone counterparts are person slices of people who will observe certain patterns in their radium decay data. If the hypothesis is true, you should expect to observe those patterns. If you do go on to observe them, then the hypothesis provides an explanation in the following sense: your observation is *unsurprising* given the hypothesis.

Hempelians will not be convinced. ‘The essence of scientific explanation’, they will insist, is ‘*nomio* expectability—that is, expectability on the basis of lawful connections’ (Salmon [1989], p. 57, original emphasis). Since my radium decay ‘explanation’ does not invoke any laws, they will judge it to be—at best—incomplete. If they also assume that cone counterpart frequencies are never law-governed, they will draw a stronger conclusion: that such frequencies are explanatorily impotent in principle. Even non-Hempelians may feel that there is something importantly right about this line of attack. For example, one might think that unification of diverse phenomena is the only *essential* feature of explanation, but nevertheless suspect that unification will, in practice, almost always involve appeal to lawful regularities.

In §8 I will argue that, superficial appearances to the contrary notwithstanding, my analysis does allow for chance laws. If that argument is successful, I can stand firm on the explanatory front. The chances that we have any hope of *knowing* about are the law-governed ones; and law-governed chances can be cited in explanations without falling foul of Hempelian strictures.

5.6 Counterfactual chances

There is no radium-224 nucleus in your atom trap. But *if there were*, its chance of decaying in the next four days would be slightly more than 50%. Can a finite frequentist make sense of this fact? If he insists on appealing only to actual reference classes, then there is a worry: the reference class might be empty (Hájek [1997], p.

³³ I am grateful to an anonymous referee for drawing my attention to this objection.

³⁴ Even the claim that chance has an explanatory role is somewhat controversial: see, e.g., (Woodward [1989]). If chances are credence guiding but *not* explanatory, then the objection considered here does not get started.

220). In particular—to focus on my version of finite frequentism—it might be that no cone, anywhere, is qualitatively identical to the cone an oriented time slice of a radium-224 nucleus would have were it in your trap now. On the other hand, perhaps there *are* such cones. Readers might expect me to stress the latter point, and to plead once again the virtues of a ‘big actual world’ hypothesis.³⁵ Not this time! Counterfactual chances are chances in *non-actual* worlds, and the chances in a non-actual world w are—according to my analysis—cone counterpart frequencies in w .³⁶ The story I need to tell about the atom trap is as follows: *if there were* an oriented time slice of a radium-224 nucleus in your trap, then it *would* have many cone counterparts, of which slightly more than 50% would have the decays-within-four-days property. Furthermore, if there were an oriented time slice of an unobtainium nucleus in your trap, then it too would have many cone counterparts, and their decay statistics would be whatever they needed to be to make the half-life come out right.³⁷ Some will object that the counterfactual dependencies envisaged here are intolerably mysterious. I reply that we should understand them nomically: the laws of nature permit certain structures (e.g., time slices of radium or unobtainium nuclei) to occur *only* within large cone counterpart classes, and they impose strict constraints on the statistical characteristics of those classes. The proposed laws are, of course, radically non-local. Once again, I must defer further discussion to §8.

6 Chance and Credence

A reminder: if Q is the proposition that you will live to be 100, and X is the proposition that your present chance of living to be 100 is 0.15, then—whatever your evidence—it is a requirement of rationality that your present credence function C be such that $C(Q|X)$, if defined at all, equals 0.15. This is a platitude. It ought to follow more or less naturally from the correct analysis of chance. That is not to say that it must be an *entailment* of the correct analysis. In arguing from the analysis to the credence constraint, we are—I shall assume—entitled to appeal to a priori intuitions about rationality. Unless hypotheses about chance are taken to have normative *content*, it is hard to see how such appeal can be avoided. In any case, beggars can’t be choosers. No existing analysis of chance scores well on the chance–credence front.³⁸ The particular objection faced by finite frequentism is that hypotheses about reference class frequencies will not, in general, screen out all case-specific evidence; but the objection loses its force if reference classes are defined narrowly enough. This thought, noted but left undeveloped in §3, provided much of the initial motivation for my analysis. It is time to spell it out.

6.1 An infeasible indifference intuition

Let X' be the hypothesis that a fraction 0.15 of your cone counterparts are slices of people who live to be 100. Suppose $C(X') \neq 0$. I claim it is then a requirement of rationality that $C(Q|X') = 0.15$. A rough-and-ready argument goes as follows. Suppose

³⁵ I owe this suggestion to an anonymous referee (who does not endorse it).

³⁶ Finite frequentism can reasonably be characterized as the view that ‘chances are actual frequencies’, but—as noted in §1—‘actual’ should be read non-rigidly.

³⁷ Here I am imagining (with Lewis [1994]) that there are no unobtainium nuclei in the actual world, but that there are nevertheless facts about what their decay chances would be if they existed.

³⁸ David Lewis prided himself on scoring least badly: he could—he claimed—see ‘dimly but well enough’ how his own analysis might ground a Principal-Principle-shaped constraint (Lewis [1994], p. 484). I shall be making rather stronger claims for my analysis.

God, who you know always tells the truth, (i) tells you that you have a finite number of cone counterparts, (ii) gives you names for all of them, (iii) tells you which are (and which are not) person slices of people who reach 100, and (iv) does not tell you which cone counterpart is *your* present person slice.³⁹ Reading through God's list, you discover that 15% of your cone counterparts are centenarian slices. Since precisely one of the cone counterparts is your present person slice, and since each is an equally plausible candidate for that distinction, there is a very strong intuition that you should be 15% confident that you will live to be 100. There is an equally strong intuition that the percentage is all that matters: so long as this is held constant, your confidence that you will live to be 100 should be invariant under permutations, expansions or contractions of God's list. But if God simply tells you that X' is true, he's given you the percentage. (He has also thereby revealed that you have a finite number of cone counterparts.) Thus, $C(Q|X')$ ought to be 0.15.

The chain of thought just sketched depends crucially on a self-location indifference intuition. I think the intuition is compelling as it stands; furthermore, I doubt that a non-question-begging *proof* of its correctness is possible. Those inclined to dismiss indifference intuitions across the board should reflect on the uncontroversial fact—admitted even by Hájek—that '[rational] opinion is guided by finite frequency information' (Hájek [2011]). Why should this be? It is difficult to imagine an explanation that did not appeal to indifference intuitions of *some* sort.⁴⁰ Generically, of course, these intuitions will be defeasible. If you have good evidence that you have just stepped off a 200 meter cliff, then you need no longer take account of your insurance provider's actuarial data when forming an opinion about your likely life span. (You are falling towards the sea; your actuarial reference classmates are not.) One great virtue of cone counterpart frequencies is that they support *indefeasible* indifference intuitions. Even as you fall towards the sea, your conditional credence $C(Q|X')$ should—if defined—be 0.15. After all, your cone counterparts are in qualitatively identical predicaments, yet, if X' is true, 15% of them are slices of people who will see their hundredth birthdays. Strong evidence of imminent catastrophe is evidence against X' , but it is no reason for you to adjust your *conditional* credence. In the remainder of this section, I want to articulate the indefeasible indifference intuition as a general credence constraint, to defend that constraint against objections, and, finally, to plead its virtues as a chance–credence bridge principle.

Before I attempt to generalise, I ought to deal with a technical worry. $C(Q|X')$ is, by definition, $C(Q)/C(X')$. If $C(X')=0$, then $C(Q|X')$ is undefined, and the indifference constraint—that $C(Q|X')$ should be 0.15 *if it is defined*—is vacuously satisfied. Now, some readers may worry that $C(X')$ *has* to be zero, on pain of irrationality, and that the indifference constraint is therefore toothless. Probably they have something like the following argument in mind:

Any apparent grounds for favouring hypothesis X' will, on closer inspection, turn out to be grounds for favouring—and favouring roughly equally—members of a *family* of mutually incompatible hypotheses {a fraction p of my

³⁹ By definition, your cone counterparts receive qualitatively identical deliverances. Thus, unless all or none of your cone counterparts are centenarian slices, God *cannot* reveal your fate.

⁴⁰ Witness Jonathan Schaffer ([2003], p. 32): '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'. Though he does not say so, it seems clear that Schaffer is appealing to an indifference intuition.

cone counterparts live to be $100 \mid 0.15 - \delta \leq p \leq 0.15 + \delta$, where $\delta > 0$. But there are infinitely many hypotheses in that family. The only way to favour them ‘roughly equally’ is to assign each zero credence.

The argument’s premise is dubious: conceivably, symmetry or other theoretical considerations might favour one particular fraction over arbitrarily close approximations. Still, I can—and should—immunise myself against this line of attack by switching to interval-based credence constraints. Let I be a real interval (open, half-open or closed) and let X_I be the hypothesis that there exists $p \in I$ such that a fraction p of your cone counterparts live to be 100. I claim that your credence function must, on pain of irrationality, respect the following constraint:

If $C(Q|X_I)$ is defined, then $C(Q|X_I) \in I$.

The driving intuition remains the same: indefeasible indifference.⁴¹ But since no one will want to argue that $C(X_I)$ is required to be zero *for all* I , the interval constraint is certainly not toothless.

In the following subsection, I will work my way, in stages, towards a general credence principle. Until the very final stage, I will suppress the interval formalism and work with hypotheses ascribing precise cone counterpart frequencies. This aids readability and makes no substantive difference to the argument. (Readers who want to formulate interval versions of the provisional principles should have no difficulty doing so.)

6.2 Articulating the credence principle

I shall need to introduce some new notation. First up is ‘*’. ‘*’ is a reflexive indexical that picks out an agent’s present person slice.⁴² So, for example, if you believe A^* , you believe that your present person slice has property A . Where there is no danger of confusion, I shall sometimes use ‘*’ to refer to the owner of the credence function we are discussing. We can now write down what looks at first sight like an innocent generalisation of the lifespan example:

Rational credence functions must satisfy $C(A^*|Ch_*(A)=p)=p$ for any A, p for which the conditional credence is defined. (CP1)

I shall assume that properties are individuated finely enough that we don’t need to keep separate track of modes of presentation.

Unfortunately, CP1 is *not* an innocent generalisation of our lifespan example. It faces counterexamples. I will describe one, sketch a general framework for thinking about similar cases, and then present an amended credence constraint. Let B_1 be the property *is a person slice of a person who one day meets Barack Obama*. Now consider the hypothesis $Ch_*(B_1)=0.01$. The hypothesis implies that at least one of your cone counterparts has property B_1 . Can you appeal to an indifference intuition here? Surely you cannot. Spatiotemporally *distant* cone counterparts are most unlikely to be B_1 s. Obama, after all, lives on Earth in the twentieth and twenty-first centuries. You

⁴¹ For example: conditional on the hypothesis that between 15% and 30% of your cone counterparts live to be 100, your confidence that *you* will live to be 100 should not be less than 15% or more than 30%.

⁴² Reminder: a *person slice* is a time slice of a person oriented to match the person’s psychological arrow of time.

might meet him, but short-lived creatures born on the other side of the universe haven't a prayer. Furthermore, you have excellent evidence that you have no non-trivial cone counterparts in your spatiotemporal vicinity. It follows that rationally permissible values of $C(B_1^*|Ch_*(B_1)=0.01)$ are near unity. At any rate, it certainly isn't a requirement of rationality that $C(B_1^*|Ch_*(B_1)=0.01)$, if defined, must be 0.01.

The way to block such counterexamples is to restrict the range of properties to which the credence constraint applies. Encouragingly, the hypothesis that $Ch_*(B_1)=0.01$ is a rather peculiar one. It does not correspond to any thought we can articulate without explicitly referring to cone counterparts. Perhaps you believe that your present chance of meeting Obama is 1%. If so, you surely do not believe that 1% of your cone counterparts meet the man *you* refer to as 'Barack Obama'. What you might believe—were you to accept the letter and spirit of my analysis—is that 1% of your cone counterparts meet the man *they* refer to as 'Barack Obama'. Let us say that y has property B_2 iff y is a person slice of a person who one day meets the man y refers to as 'Barack Obama'.⁴³ Good news: it plausibly *is* a requirement of rationality that $C(B_2^*|Ch_*(B_2)=0.01)$, if defined, is 0.01. What is the difference that makes the difference here? It seems to have something to do with the fact that in picking out property B_1 , but not in picking out property B_2 , you enjoy—and know that you enjoy—a *semantic privilege* that is not shared with other members of your cone counterpart class. Coming to grips with this phenomenon will enable us to patch up CP1.

Suppose you (currently) have non-trivial cone counterparts; let x be one such. Suppose further that you entertain the proposition B_1^* . Some sentence tokened by you serves, given your context, to pick out that proposition. The same sentence is tokened by x , where it serves, given x 's context, to pick out a different proposition. Furthermore, some *open* sentence tokened by you serves, given your context, to pick out property B_1 . Its extension is the set of person slices of people who one day meet Barack Obama. The same open sentence is tokened by x , where it serves, given x 's context, to pick out a different property. *Its* extension is the set of person slices of people who one day meet a particular *duplicate* of Barack Obama. I shall say that x entertains a *counterpart proposition* and self-attributes a *counterpart property*. I shall denote the former $x(B_1^*)$ and the latter $x(B_1)$. Formally, $x(\cdot)$ is a function that takes as its argument a proposition or property picked out by $*$ and returns, as its value, the counterpart proposition or property picked out by x . If we let H and A range over (respectively) propositions and properties picked out by $*$, we can define the function as follows:

$x(H)$ is the proposition picked out for x by any sentence that picks out H for $*$
 $x(A)$ is the property picked out for x by any open sentence that picks out A for $*$

Implicit in this definition is an assumption: if two sentences are propositionally equivalent for $*$, then they are propositionally equivalent for x ; similarly, if two open sentences pick out the same property for $*$, then they pick the same property for x . That assumption perhaps requires some defence. Let's begin with the propositional case. Propositions serve, for us, as arguments of rational credence functions.⁴⁴ That means we must individuate them rather finely. For example, the sentence 'water is H_2O ' does not have the same propositional content—for speakers of standard

⁴³ Read the definite description à la Russell: if 'Barack Obama' does not pick out anyone for y , then y lacks property B_2 .

⁴⁴ I claim that this is *a* role of propositions, not necessarily the sole or defining role. I want to remain as neutral as possible on the metaphysics of propositions.

English—as ‘water is water’. (Proof: you can legitimately doubt that water is H₂O, but not that water is water.) Given such a fine individuation, there is no temptation to suppose that the propositional equivalence or inequivalence of a pair of sentences could depend, for person slice y , on events lying outside y ’s cone. If someone claims that two sentences are propositionally inequivalent for y *on the supposition* that y ’s future (or absolute elsewhere) takes a certain form, we will retort that the sentences are propositionally inequivalent for y *period*; and if our interlocutor denies this, we will argue that his individuation of propositions is inappropriately coarse-grained.⁴⁵ The upshot is that we can plausibly maintain that the propositional equivalence or inequivalence for y of sentences S_1 and S_2 supervenes on y ’s history. A matching supervenience principle for open sentences and properties (to wit: whether open sentences O_1 and O_2 pick out the same property for y supervenes on y ’s history) is an immediate consequence, so long as the individuation of properties conforms—in the following sense—with that of propositions: for any given person slice, open sentences O_1 and O_2 pick out distinct properties iff O_1^* and O_2^* pick out distinct propositions. ‘Proposition’ and ‘property’ are of course terms of art that can be construed in various ways. Ultimately, I must *stipulate* that they be construed in ways that respect our supervenience principles. I have argued that the stipulation is a natural one given the central role rational credence functions play in this paper, but strictly speaking I need only insist that it is coherent. Readers who will grant that—and I can think of no grounds for doubting it—should grant that $x(\cdot)$ is well defined.

We can now repair CP1. I shall state my proposal and then defend it. An instance of CP1 is a requirement of rationality if $*$ knows that the extension of $x(A)$ does not depend on x (where x now ranges over *all* $*$ ’s cone counterparts, $*$ included). If $*$ knows that the extension of $x(A)$ is constant in this sense, then he can deduce the following: (1) the truth value of $x(Ch_*(A)=p)$ is constant; (2) $Ch_*(A)=p$ implies that $x(A^*)$ is true for a fraction p of $*$ ’s cone counterparts. Both are important. (1) ensures that the condition hypothesis in CP1 does not itself break the indifference symmetry. (2) ensures that the condition reports an epistemic success rate. Together, (1) and (2) clear the way for the indefeasible indifference intuition. Here, then, is a revised credence principle:

Rational credence functions must satisfy $C(A^*|Ch_*(A)=p)=p$ for any A, p such that

- (i) the conditional credence is defined, and
- (ii) $C(\text{the extension of } x(A) \text{ is constant})=1$, where x ranges over $*$ ’s cone counterparts. (CP2)

Notice that if a *rational* credence function satisfies clause (ii), the certainty involved must be rational: $*$ must be *entitled* to be sure that the extension of $x(A)$ is constant. Luckily, this requirement is rather easily met. For example, you should not doubt that the extension of $x(B_2)$ is constant. You pick out property B_2 using the open sentence ‘ y is a person slice of a person who one day meets the man y refers to as “Barack

⁴⁵ An example might help. Here are two sentences whose propositional equivalence or inequivalence for an English speaker *might* be claimed to turn on events that lie outside the speaker’s cone: ‘iron has atomic number 26’; ‘the element most abundant by mass in the next meteor to strike my planet has atomic number 26’. Nevertheless, it is obvious that, whatever my future actually holds, I could sincerely and rationally assert the first sentence while denying the second. Thus, *if propositions are to serve as arguments of rational credence functions*, the two sentences are propositionally inequivalent for me, whatever the referent of ‘the element most abundant by mass in the next meteor to strike my planet’ turns out to be.

Obama”’. It is obvious that an object will satisfy *your* tokening of this open sentence if and only if it satisfies each of your cone counterparts’ tokenings. To understand this, you do not require exotic knowledge of cosmically far-flung regions of spacetime; what you need is a limited appreciation of the role your context plays in fixing the semantic content of your utterances.

We began with a simple credence principle that turned out to be incorrect. We have fixed it, at the cost of some small-print: clause (ii) of CP2. Does the need for this clause reflect badly on my analysis? That depends: if the clause prevents me from saving standard chance–credence intuitions, it is bad news; otherwise, it is neutral. There is no hint of bad news in the discussion above. In particular, the B_1 case is unalarming, because p -substitution instances of $Ch_*(B_1)=p$ —unlike p -substitution instances of $Ch_*(B_2)=p$ —were never in the running to serve as analyses of pre-theoretically intelligible chance ascriptions. In fact, the B_1/B_2 story is reassuring. It suggests that *prima facie plausible* proposals for cone counterpart frequency analyses of present chance ascriptions will not fall foul of clause (ii). At any rate, the kinds of context sensitivity that block our indifference intuition seem easy enough to avoid.

CP2 governs conditionalisation on present cone counterpart frequency ascriptions. It says nothing about conditionalisation on stronger hypotheses that happen to *entail* such ascriptions. This is a significant limitation, because we will in due course want to consider how scientific theories postulating cone counterpart frequency regularities might be confirmed. Luckily, if there is a compelling indifference intuition in favour of CP2, there is an equally compelling intuition in favour of a somewhat stronger principle:

Rational credence functions must satisfy $C(A^*|Ch_*(A)=p \wedge T)=p$ for any A, p, T such that

- (i) the conditional credence is defined,
- (ii) $C(\text{the extension of } x(A) \text{ is a constant function of } x)=1$, and
- (iii) $C(\text{the truth value of } x(T) \text{ is a constant function of } x)=1$, where x ranges over $*$ ’s cone counterparts. (CP3)

Clause (iii) ensures that conditioning on T does not break the indifference symmetry. To illustrate, let’s begin from a familiar instance of CP2:

$$C(Ch_*(B_2)=0.01) \neq 0 \rightarrow C(B_2^*|Ch_*(B_2)=0.01)=0.01. \quad (1)$$

(I have absorbed clause (i) into the antecedent of a material conditional.) Conditional on the hypothesis that 1% of your cone counterparts meet the man they refer to as ‘Barack Obama’, you should be 1% confident that you will meet the man you refer to as ‘Barack Obama’, so long as the conditional credence in question is defined. Our stronger principle, CP3, yields

$$C(Ch_*(B_2)=0.01 \wedge T) \neq 0 \rightarrow C(B_2^*|Ch_*(B_2)=0.01 \wedge T)=0.01, \quad (2)$$

for any proposition T that satisfies clause (iii). *Everyone meets at least one famous politician* is one such proposition: if you entertain the proposition, so do all your cone counterparts. Substituting it for T in (2), we get an intuitively compelling credence constraint. The reason it is intuitively compelling—just as compelling as (1)—is that the second conjunct in the condition does nothing to break the indifference symmetry. The same is true if we substitute *I have met many of Barack Obama’s friends* for T .

Here, the counterpart propositions are not identical, but they are materially equivalent: if you have met many of Barack Obama’s friends, then for any x in your cone counterpart class, x has met many friends of the man x refers to as ‘Barack Obama’. Once again, the indifference symmetry is unthreatened. On the other hand, *I will meet many of Barack Obama’s friends* does not satisfy clause (iii), and if we substitute it for T in (2), the result manifestly is not a requirement of rationality.

Finally, let’s extend the credence principle to handle intervals.

Rational credence functions must satisfy $C(A^*|Ch_*(A)\in I \wedge T)\in I$ for any A , T and interval I such that

- (i) the conditional credence is defined,
- (ii) $C(\text{the extension of } x(A) \text{ is a constant function of } x)=1$, and
- (iii) $C(\text{the truth value of } x(T) \text{ is a constant function of } x)=1$, where x ranges over $*$ ’s cone counterparts. (CP4)

CP4 is a generalisation of CP3, and CP3 is a generalisation of CP2.⁴⁶ Nevertheless, the indefeasible indifference intuition that supports CP2 supports CP4 just as strongly. I will not need to refer to CP2 or CP3 again.

6.3 Appraising the credence principle

So far, I have described the indefeasible indifference intuition that I think supports CP4, and I have explained the motivation for clauses (i)–(iii). I have not attempted, and will not attempt, to *derive* CP4 from more basic and less controversial principles. Those inclined to look favourably on indifference intuitions will—I hope—agree that no such derivation is required. However, they will want to be reassured that CP4 can resist certain familiar objections.

Indifference principles that instruct us to spread our (conditional or unconditional) credences ‘uniformly’ across *infinite* sets of epistemic alternatives are notoriously problematic. Prima-facie-compelling examples have often turned out, on reflection, to be inconsistent, ambiguous, or arbitrary (van Fraassen [1989], pp. 303–5), or to require violations of countable additivity (Ross [2010]). I needn’t discuss these sorrows in detail, because CP4—to the extent that it is an indifference principle—is finitistic.⁴⁷ Its measure over cone counterparts (a simple fraction) is well defined and well motivated, and the principle certainly is not inconsistent. Still, readers might worry that it is hubristic. The rational appraisal of hypotheses is a complex matter. When we do it well, it is usually thanks to the hard-won wisdom of scientists or other topic specialists. A philosopher who proposes substantive, general constraints on the enterprise should not necessarily expect to be taken seriously. Notice, though, that CP4 is an indifference principle for *self-locating* belief. It governs, or purports to govern, the distribution of credence within (not between) sets of hypotheses whose members differ *only* in respect of $*$ ’s location in the world. It does not seek to arbitrate between rival descriptions of the world. That already draws much of the sting from the hubris objection. CP4 might yet have implications for scientific practice. However, the best reason for thinking that it might is that chances might be cone counterpart frequencies;

⁴⁶ To obtain CP3 from CP4, set I to be $[p, p]$, and notice that $Ch_*(A)\in[p, p]$ is logically equivalent to $Ch_*(A)=p$. To obtain CP2 from CP3, set T to be any tautology.

⁴⁷ Recall that the only well-defined cone counterpart frequencies that can arise in *infinite* cone counterpart classes are the trivial frequencies 0 and 1. The associated credence constraints are not in any interesting sense indifference constraints: they are requirements of logic. For example, $C(Ch_*(A)=1 \wedge T)\neq 0 \rightarrow C(A^*|Ch_*(A)=1 \wedge T)=1$: if *all* your cone counterparts are A s, then *you* are an A .

and if chances *are* cone counterpart frequencies, then CP4's implications—roughly, that the Principal Principle is true—are conservative.

Appeals to self-location indifference intuitions are not uncommon in the philosophical literature. Consider the much-discussed Sleeping Beauty case. Beauty is put to sleep, and depending on the outcome of a fair coin toss, she is woken once (heads) or twice (tails). Before any second waking, Beauty's memory of her first waking is erased; more generally, the experimental protocol ensures that wakings are, from Beauty's point of view, indistinguishable. Beauty knows all this. On waking, what should Beauty's credence be in the proposition *the coin landed heads*? Most authors think the answer is $\frac{1}{3}$; David Lewis argued for $\frac{1}{2}$. What matters for our purposes is the common ground: halvers and thirder agree that Beauty should assign equal credence to *the coin landed tails and this is my first waking* and *the coin landed tails and this is my second waking*. These two hypotheses differ only in respect of Beauty's (present person slice's) location in the world, and Beauty could have no good reason to favour one over the other. Adam Elga ([2004]) has attempted to parlay such relatively uncontroversial intuitions into a general self-location indifference principle. As Brian Weatherson ([2005]) notes, Elga runs into trouble on two fronts: his principle is not finitistic (cue countable additivity worries), and he appeals to a relation of subjective indistinguishability defined over agent predicaments. This relation is most naturally taken to be intransitive; *deserves the same credence as*, by contrast, is transitive, and the mismatch is fatal. My principle draws on similar intuitions, but is much less ambitious than Elga's. CP4's indifference demands are finitistic, and *is a cone counterpart of* is unquestionably an equivalence relation.

My case for CP4, then, boils down to this: it is an articulation of a compelling indifference intuition, and it is not vulnerable to the objections that have, historically, given indifference principles a bad name. Notice what the case does *not* require. It does not require that readers buy a 'package deal' account of natural law, chance and inductive confirmation (compare Lewis [1994], pp. 478–80). You should accept CP4 *whatever* your views on law, and *whatever* your views on chance. Your views on inductive confirmation will be an obstacle only if they prompt you to reject the self-location indifference intuition, and it is hard to see how they could.

If CP4 is true, how far does that strengthen the case for my analysis? I think it strengthens it very considerably. Conjoined with the analysis, CP4 becomes a chance–credence bridge principle. Many such principles have been discussed in the literature, all more or less in the ballpark of Lewis's original Principal Principle, but mine has some significant advantages. Most importantly, it involves no heroic abstractions: true or false, it is unquestionably intelligible. The same cannot be said for all its cousins. Some versions of the Principal Principle, Lewis's original formulation included, are constraints on *initial credence functions*.⁴⁸ Their authors assume—explicitly or implicitly—that a rational agent's current credence function is her initial credence function conditionalised on her *total evidence proposition*. The italicised concepts are manifestly problematic. CP4, by contrast, is a constraint on current credence functions. There is no appeal to total evidence propositions; in fact, there is no mention of evidence at all. This brings us to a second attractive feature: CP4 has—and requires—no admissibility clause.⁴⁹ Suppose you have a crystal ball with an impeccable track record. The ball assures you that you will live to be 100. You trust the prediction, and you are entitled to do so. The case is peculiar, but it poses no special difficulty for me.

⁴⁸ For example, (Lewis [1986b] and Schaffer [2003]). Meacham ([2010]) points out that some versions of the Principal Principle that are not so formulated ought (by their own lights) to be.

⁴⁹ NB that clause (iii) is not an *evidential* restriction.

Your *conditional* credence $C(Q|X)$ still ought to be 0.15.⁵⁰ After all, your cone counterparts have counterpart balls with equally impressive track records. Cone counterpart frequencies screen out the deliverances of crystal balls in precisely the same way that they screen out less exotic types of evidence.

The lack of an admissibility clause does not expose my bridge principle to counterexamples, but it does limit its scope. Suppose you know that a fair coin has been tossed (NB the past tense). What should your current credence be that it landed heads? The tempting answer is that your credence should be 0.5, so long as you have no inadmissible evidence. That answer is compatible with CP4, but it is not a *consequence* of CP4. Should I be worried? I think I should be relieved. The tempting answer is controversial, at least in some circumstances: witness the Sleeping Beauty case. Furthermore, and relatedly, the concept of inadmissible evidence is notoriously difficult to pin down. In CP4, I have a well-defined and well-motivated bridge principle that I can defend against all comers while preserving my freedom of manoeuvre on other philosophical fronts. That is a good position to be in. The situation would be very different if my account threatened to *undermine* retrospective chance-credence intuitions—but it does not. Suppose chances are cone counterpart frequencies, and suppose you know that a fair coin was tossed on a particular day in 1898. What object does the chance of 0.5 attach to? For this example, it doesn't much matter; let's take it to be an oriented time slice of the coin, shortly before it was tossed. What you know, then, is that the oriented time slice in question belongs to a non-trivial cone counterpart class, half of whose members have 'heads' futures, and half 'tails'. How confident should you be that the coin—the coin that *you* successfully refer to—landed heads? Of course it depends on your evidence. You might reasonably be very confident indeed: perhaps the result was recorded by a reliable witness, or perhaps you have good reason to believe that world history would have turned out very differently had the coin landed tails; 'thirders' might also want to imagine analogues of the Sleeping Beauty case. But if you have no evidence that directly bears on the result of the coin toss, and no reason to think that you are in anything like Sleeping Beauty's predicament, then there is a strong intuition that your credence that the coin landed heads ought to be 0.5. This is a *defeasible* (but, by hypothesis, undefeated) indifference intuition. The bottom line is that I have not made a mystery of our retrospective chance-credence intuitions. I leave them as I found them: messy, and unfriendly to exceptionless principles.

7 Confirmation

If the radium nuclei in your lab have non-trivial cone counterparts, then they are scattered widely throughout the actual world; they may even be spatiotemporally disconnected. You are in no position to survey them. Any knowledge you have of the decay statistics associated with such classes must be inferred with the help of a *scientific theory*—a theory telling you that objects with such-and-such locally ascertainable characteristics are associated with cone counterpart frequencies of a particular sort. But what grounds could you possibly have for believing the theory? Some readers will think it obvious that the answer is 'none'. They will complain that my analysis makes (non-trivial) chances unknowable. I shall defend a different

⁵⁰ Unless your confidence in the ball's prediction is total ($C(Q)=1$), in which case $C(Q|X)$, if defined, must be 1. CP4 then requires that $C(X)=0$, to ensure that $C(Q|X)$ is *not* defined. The result is intuitive: if you are certain you will reach 100, you must also be certain that there is no chance that you won't.

answer: you might have *strong confirming evidence*.

Let's begin by rehearsing a familiar confirmation-theoretic story. A scientific theory that postulates stable patterns of chance will inevitably ascribe high chances to certain outcomes. For example, if a coin always has a 50% chance of landing heads when tossed, then there is a high chance that the long-run proportion of heads in many thousands of observed tosses will lie close to 50%. Thanks to the chance–credence link, this qualifies as an empirical prediction: one can *test* the hypothesis that a coin is fair by tossing it many times. More generally, one can test a theory that postulates stable patterns of chance by measuring frequencies. Can anything like this story be made to work if chances are cone counterpart frequencies? Since my analysis leaves the chance–credence link intact, the prospects look good. Suppose theory T_0 , when conjoined with some routine auxiliary propositions justifiably believed by you, entails that 99.9% of your cone counterparts will observe a particular statistical pattern in their radium decay data. Then, on the supposition that T_0 is true, *you* should expect to observe the pattern. So T_0 is testable.

That was rather brisk. Let's fill in some formal details. Suppose theory T_0 entails $(\forall x)(Bx \rightarrow Ch_x(A_0) \in I_0)$, where quantification is over objects, A_0 is such that you can reasonably expect one of your future person slices to learn, retrospectively, whether your *present* person slice has property A_0 , I_0 is a narrow interval near (but not including) unity, and your present credence function C satisfies $C(B^*) \approx C(B^*|T_0) \approx 1$, because (say) you have good empirical evidence that B^* .⁵¹ Suppose also that you are rational. It follows that $C(Ch_*(A_0) \in I_0 | T_0) \approx 1$. I now want to show that $C(A_0^*|T_0)$ must be high. Notice first that the following inequality is a requirement of probabilistic coherence:

$$C(A_0^*|T_0) \geq C(A_0^* | Ch_*(A_0) \in I_0 \wedge T_0) C(Ch_*(A_0) \in I_0 | T_0).$$

Thus, since $C(Ch_*(A_0) \in I_0 | T_0) \approx 1$, $C(A_0^*|T_0)$ cannot be much less than $C(A_0^* | Ch_*(A_0) \in I_0 \wedge T_0)$. But if conditions (i) through (iii) of CP4 are satisfied, then the latter must lie within I_0 —in which case $C(A_0^*|T_0)$ cannot be much less than unity, and we can reasonably describe T_0 as making a testable prediction: on the supposition that T_0 is true, you should expect A_0^* to be true.

We have just seen that cone counterpart frequency regularity hypotheses might in principle have empirical content. However, it does not follow that the epistemology of chance is unproblematic for me. I anticipate that readers will have worries on two fronts. First: can the confirmation-theoretic template I have sketched actually be applied? Can I, for example, describe plausible candidates for the A_0 , B and T_0 roles in the radium decay case? Second: wouldn't any non-trivial cone counterpart frequency regularity qualify as a *cosmic coincidence*? If so, then theories postulating such regularities deserve little credence, however impressive their empirical track records.⁵² I shall focus on the first worry in this section, and present my case for resisting the 'cosmic coincidence' intuition in the next.

Suppose you have a sample of radium on your bench. A scientific theory of radium decay should tell you what sort of decay statistics to expect; my task is to explain how a theory postulating cone counterpart frequency regularities might do that.

⁵¹ The supposition that I_0 does not include unity is not essential to the argument in this paragraph, but I want to focus attention on non-trivial cone counterpart frequencies.

⁵² This claim is not in tension with my confirmation-theoretic sketch. If T_0 postulates a cosmic coincidence, then rationally permissible values of $C(T_0|A_0^*)$ will be small, even if $C(A_0^*|T_0)$ is required to be near unity.

Empirical predictions will have the logical form A_0^* . As noted in §2, we have some flexibility in how we choose to interpret ‘*’. I shall suppose that its referent—your present person slice—extends far enough into your environment to encompass the radium sample. It is then natural to take A_0 to be the property picked out, for you, by an open sentence of the following form:

The proportion of radium nuclei intersected by y that decay within $t(y)$ lies in interval I_1 ,

where y ranges over objects, and $t(y)$ denotes some specific portion of y ’s future-pointing cone, defined relative to y (e.g. *the next hour* or *between six and eight hours later*).⁵³ Notice that A_0^* is, as required, an empirically contentful hypothesis: it says that the proportion of radium nuclei in your sample that decay within $t(*)$ lies in interval I_1 . Notice also that you ought to be certain that the extension of $x(A_0)$ is the same for every x in your cone counterpart class. (Proof: when your cone counterparts say ‘radium’, they pick out the same natural kind you do. Furthermore, they have the same spatiotemporal concepts, the same concepts of *decay* and *interval*, and so on.) That means we can appeal to CP4 without falling foul of clause (ii). So far so good; we now need to consider the theory, T_0 , that is supposed to support this prediction.

We want T_0 to entail $(\forall x)(Bx \rightarrow Ch_x(A_0) \in I_0)$, and we want the latter proposition to say that radium samples of a particular kind have a high chance of manifesting decay statistics in the range specified by A_0 . B is a property of objects. What will it look like? Being a B will *at least* involve being a space-like surface that encompasses a sufficiently large radium sample.⁵⁴ Perhaps there will be further intrinsic restrictions (on nuclear excitation states and so on). However, we should not expect B to be an intrinsic property of objects. One reason is that radioactive decay overwhelmingly favours one temporal direction. That means it is sensitive to some *physical* arrow of time; but the direction of a physical arrow at a given space-like surface will not, in general, be settled by intrinsic properties of the surface.⁵⁵ There is also a second reason. Consider the prediction that roughly half of a sample of radium-228 nuclei will undergo nuclear decay within 5.75 years. The prediction fails if the sample gets annihilated by a blob of antimatter after only 10% of its nuclei have decayed. If you make the prediction on the strength of your knowledge of radium-228’s half-life, then presumably you believe that environmental conditions in the sample’s vicinity will remain within certain bounds over the time period in question: at the very least, you believe that the sample won’t get bombarded by antimatter. Now, we are imagining that T_0 is a theory about *radium decay*, not a theory about the trajectories of antimatter blobs. We should not expect it to tell us that radium samples with certain intrinsic characteristics are fated never to meet with antimatter blobs; rather, we should expect it to tell us that radium samples with certain intrinsic characteristics *which happen not to encounter antimatter blobs* will (probably) display certain decay patterns. So, to recap, we should expect open formula By to be satisfied by all and only objects that (i) are space-like surfaces through radium samples with appropriate intrinsic characteristics, (ii) are aligned with the physical arrow of time

⁵³ y ’s future-pointing cone is the union of *future* light cones of points in y , where the future direction is the temporal direction y does not designate as its past.

⁵⁴ The minimum size of the sample will depend on the details of the statistical pattern encoded by A_0 .

⁵⁵ This is perhaps most obvious if one considers a *small* space-like surface. A snapshot of a microscopic portion of, say, a hot cup of tea might provide few clues about the direction of the thermodynamic arrow.

implicated in radium decay, and (iii) have future cones whose initial segments are free from certain specified kinds of interference. The spatiotemporal region constrained by (iii) will be related to the timeframe set by A_0 (roughly speaking, it will be the smallest initial segment of y 's future cone that contains $t(y)$). Since radioactive decay chances are known to be rather insensitive to local conditions, we should not expect the intrinsic or the extrinsic constraints to be very exacting. Where does this leave the knowability of B^* ? Clearly there is no question of B^* being *infallible* knowledge, for it is (or at any rate entails) a substantive hypothesis about $*$'s future. All we require, however, is that $*$ could reasonably assign B^* high credence, given suitable evidence. That requirement is not very demanding. If you think scientists are ever in a position to make justified predictions about decay statistics, you will grant that they (sometimes) have justified beliefs about certain intrinsic characteristics of their samples, about the relevant arrow of time, and about the absence now *and in the future* of certain disrupting environmental conditions.⁵⁶

My initial sketch envisaged a single universal generalisation of the form

$$(\forall x)(Bx \rightarrow Ch_x(A_0) \in I_0). \quad (3)$$

However, if T_0 is to be a general theory of radium decay, then it had better entail infinitely many instances of the form, covering radium samples of all sizes, and decay statistics over all time frames and at every level of probability. I therefore now relax the requirement that I_0 be an interval *near unity*, and stipulate instead that for a regularity to qualify as an instance of (3), I_0 must exclude the extreme values 0 and 1. (Without that stipulation, the claim that some instances of (3) are testable would be trivial: consider the intervals $[0,0]$ and $[1,1]$.) I do not propose to pursue the question of how T_0 might best be axiomatised. That issue will arise for anyone who, seeking relativistic generality, wishes to associate chances with space-like surfaces; *qua frequentist* I am under no special obligation to tackle it. An issue I am obliged to tackle is that of smoothness. Scientifically-tutored common sense tells us that the chance that a given radium nucleus decays within t seconds should—given some not-very-restrictive environmental assumptions—be a smoothly increasing function of t . By choosing an appropriate object (e.g. an oriented time slice of a radium nucleus in a suitable environment), I can express that intuition within my relativistic framework; what I cannot do—since cone counterpart frequencies are ‘quantised’ in units of the reciprocal of cone counterpart class cardinality—is reconcile it with my analysis. A fortiori, T_0 won’t save it. However, for any $\delta > 0$, T_0 *could* imply that the decay-within- t cone counterpart frequency lies within the interval $[p(t)-\delta, p(t)+\delta]$, where $p(t)$ is a smoothly increasing function of t . Thus, smoothness to any degree of precision short of perfection is in principle within reach. Let us suppose that T_0 does a good job on the approximate smoothness front. The supposition has a significant (but by now familiar) consequence: T_0 will imply that certain types of object belong to very large cone counterpart classes.

Finally, if T_0 is to have empirical content, the conditions of CP4 must be satisfied. We have already noted that A_0 satisfies clause (ii). Clause (iii) does not look any more threatening. So long as T_0 is a general scientific theory, the sentence you use

⁵⁶ Recall that I am assuming that, in any given context, ‘ $*$ ’ picks out one particular oriented space-like surface, but *not* that the tokening agent knows, with precision, what that oriented space-like surface is (§2). For B^* to be knowable, it is important that the truth value of B^* be relatively insensitive to small perturbations in the referent of ‘ $*$ ’. This requirement seems likely to be fulfilled, for—as already noted—we should not expect B 's instantiation conditions to be very exacting.

to pick it out will be free of proper names and self-location indexicals such as ‘here’, ‘now’ and ‘*’. The type of indexicality we should expect is that associated with natural kind terms (‘water’ refers—for us—to the natural kind in *our* oceans), but this won’t be a problem, because cone counterparts’ natural kind terms are coreferential.⁵⁷ The truth value of $x(T_0)$, then, will be the same for any x in your cone counterpart class. That leaves clause (i). Perhaps theories postulating non-trivial cone counterpart frequency regularities are, ex ante, very implausible. If that is right, then $C(Ch^*(A_0) \in I_0 \wedge T_0)$ ought to be very small for an agent who has yet to gather much relevant evidence—but it is hard to imagine a convincing argument for the claim that it must be exactly zero. Absent such an argument, there is no good reason to doubt that $C(A_0^* | Ch^*(A_0) \in I_0 \wedge T_0)$ could be defined.

I conclude that a theory postulating cone counterpart frequency regularities might indeed make testable predictions.

8 Laws

I have already asked readers to take seriously the suggestion that the actual world might accommodate a rich range of cone counterpart frequencies. My final task is to persuade them to take seriously the further suggestion that some cone counterpart frequency regularities might be *nomic*ly necessary. Which regularities, exactly? Whichever are required to save the nomic facts about chance. Cone counterparts may be arbitrarily widely dispersed, so the laws I have in mind will not be local laws. For example, whether a particular radium sample respects or violates a hypothesised chancy decay law over some period will not, on my account, supervene on the intrinsic properties of the region of spacetime occupied by the sample over the period, nor indeed on the intrinsic properties of the spacetime neighbourhood, however generously it is defined (unless we are *maximally* generous, and count as the neighbourhood the whole of the actual world). Laws of this sort would be irreducibly global constraints on global statistics. Could there really be such laws? I cannot prove that the answer is ‘yes’. Instead, I shall argue that it is not obvious that the answer is ‘no’.

The stakes here are considerable. If I were not permitted to appeal to *lawful* cone counterpart frequency regularities, then I would struggle to account for the role of chance-ascribing hypotheses in explanations (§5.5) and in counterfactual conditionals (§5.6). Worse, I would struggle to avoid the conclusion that chances are unknowable. I hinted at this last point in the previous section; let me now spell it out. If there are no lawful cone counterpart frequency regularities, then the regularities entailed by T_0 (the radium-decay theory envisaged above) are *coincidental* regularities. It is sometimes rational to believe hypotheses that postulate coincidences. However, it is never rational to believe hypotheses that postulate bigger coincidences than are required to explain one’s evidence. For example, it might in certain circumstances be rational to believe that a fair coin happened—by coincidence—to land heads ten times in a row, but it would not be rational to believe this on the strength of five observed tosses. Now, the regularities postulated by T_0 are *universal* regularities; and if T_0 is non-vacuously true, then the world is a very big place. It could never be rational to believe, on the strength of local evidence, that such universal regularities *coincidentally happen* to hold.

⁵⁷ One can perhaps imagine freak cases where this will not be true. Let ‘futurium’ denote the element predominant by mass in the next meteor to hit our planet. I am inclined to deny that ‘futurium’ counts as a natural kind term. But if it *does* count, then cone counterparts’ natural kind terms need not always be coreferential. For clause (iii) to be satisfied, the referents of any natural kind terms that appear in the sentence used to pick out T_0 must be settled by *’s history. This is not a very demanding constraint.

Notice, however, that T_0 does not postulate that the regularities are coincidences: it just postulates that the regularities do, in fact, hold. The claim that they would be coincidences is a substantive philosophical thesis. My task in the remainder of this section is to show that we have good grounds for doubting it.

The coincidence thesis seems to be supported by the following observation: if the truth of T_0 were nomically necessary—if, for various values of B , A_0 and I_0 , being a B nomically necessitated belonging to a cone counterpart class whose proportion of A_0 s lay in interval I_0 —then physical law would be non-local in a radical and hitherto unsuspected way. I don't wish to challenge the observation. On reflection, though, it is not clear that it provides much support for the thesis. Why shouldn't there be laws of the envisaged sort? Is there a *metaphysical* prohibition? There might be, but I have no idea how one would prove it. Those who wish to make the attempt ought to be demoralised by the following observations. First, if anything remotely like Lewis's 'best systems' account of laws is right, then instances of (3) plausibly *will* qualify as laws in radium-containing worlds at which T_0 is true: they are simple, and they are very informative. Second, the best-known *non*-Lewisian account of law, that due to Armstrong ([1983]), Dretske ([1977]) and Tooley ([1977])—henceforth 'ADT'—is unthreatening. ADT hold that it is a law that all F s are G s just in case a relation of *necessitation* holds between corresponding universals, F_{ness} and G_{ness} . Are there worlds in which our generalisations qualify? ADT do not appear to be under any pressure to answer affirmatively, but nor, qua defenders of their theory, are they obliged to answer negatively. Their *theory* does not tell them what universals there (possibly) are, and it is a standard complaint that it doesn't say much about the necessitation relation either. The most we can say is that ADT leave room for an independent intuition to the effect that no instance of (3) is a law in any possible world. I suspect many readers will feel that intuition; I confess to feeling it (weakly) myself. We shouldn't trust it. It can be explained away too easily as a spurious surrogate for a vaguer, but also safer, intuition: laws à la (3) would be *spooky*. They surely would be, but metaphysical possibility is a broad church. If talking donkeys are possible, why not spooky laws?

I hope I have convinced readers to take seriously the possibility that instances of (3) might be laws in *some* worlds. But even—and perhaps especially—if I have, I foresee the following protest. 'Damn your talking donkeys and your spooky laws. Physics teaches us what *actual* laws look like, and they don't look like (3).' The protester concedes, if only for the sake of argument, that instances of (3) are laws in some worlds. He presumably also admits that some instances might, for all we know, be actual truths. What he insists upon is that no instance is an actual law. Fleshed out, I imagine his case would look something like this:

If it is nomically required that the proportion of A_0 s in any B 's cone counterpart class lies within I_0 , then (since laws support counterfactuals) radium nuclei are caught up in webs of non-local counterfactual dependency. For example, decays of the radium nuclei in your lab would be prevented or compelled by (respectively) the decays or non-decays of sufficient numbers of distant doppelgangers. But this is absurd. Physical intuition tells us that the decays of widely separated radium nuclei are counterfactually independent—and we should trust physical intuition, because it is a distillation of lots of diverse empirical evidence.

One might dispute the strength of the intuition appealed to here. After all, quantum

mechanics provides well-known examples of counterfactual dependency that (on the standard view) are not underwritten by local mechanisms. Readers familiar with that physics might be very slightly less scandalised by my mooted decay laws. Still, I don't want to deny the datum: physical intuition surely is hostile to the suggestion that instances of (3) are actual laws. My claim is that physical intuition is untrustworthy in this case, because it has no legitimate empirical basis. To think otherwise is to think that our experiences (and the intuitions they shape) would somehow be different if instances of (3) *were* laws. But how would our experiences be different? Obviously we would not notice the counterfactual dependencies mentioned above; cone counterparts are too remote. Perhaps a case could be made for the following subjunctive conditional: if the laws of physics included spooky constraints on cone counterpart frequencies, then they would include spooky constraints on various *measurable* frequencies too (where a frequency is measurable if we can obtain statistically significant samples from the relevant reference class). I have no idea whether this conditional is true, but fortunately I don't need to resist it. Robustly regular measurable frequencies are two-a-penny. The hypothesis that some of these regularities are *nominally required* may be suspect in all sorts of ways, but it can hardly be said to clash with experience.⁵⁸

Intuitions are stubborn things. Some readers will, I expect, remain confident that any non-trivial cone counterpart frequency regularities would be coincidences. They may even be right. What I hope I have shown in this section is that they are not *obviously* right.

9 Summing Up

The analysis of chance I have defended is reductive, simple, and metaphysically undemanding. It makes no appeal to modal concepts, or to objective measures of simplicity or informativeness. It commits us to no particular analysis of laws. So far these are generic strengths: almost any species of finite frequentism could claim them. My version, however, has some special advantages.

- Cone counterpart classes are well defined. If you buy the analysis you get a complete solution to the reference class problem.
- The analysis makes possible a simple and satisfying explanation of the chance–credence connection.
- The analysis implies that the past is not chancy. Add the not-very-contentious assumption that laws of nature hold universally, and it also implies the nomic necessity principle.
- The analysis saves the intuition that a fair coin could land heads every time it is tossed.
- The analysis saves something of the spirit of Leibniz's dictum. If cone counterpart frequency is not *strictly speaking* a graded possibility, it is at least a promising ersatz.

⁵⁸ Roberts [unpublished] argues that there could be laws—and, for all we know, *actual* laws—that constrain 'the frequency with which *F*s are *G*s' to lie in a particular interval. The radium decay laws I envisaged above are not quite instances of this form (they say that *all B*s belong to cone counterpart classes whose *A*₀ frequencies lie in a particular interval; this is compatible with there being an infinite number of such classes, and hence with the non-existence of any overall frequency of *A*₀s within the *B*s). Nevertheless, the 'spookiness' worry, if it is a worry at all, is a worry for both of us—and I think Roberts allays it rather successfully.

There are also some disadvantages. No version of finite frequentism can accommodate irrational chances. Thus, the analysis is incompatible with the strict truth of certain quantum mechanical chance ascriptions. It is compatible with their *approximate* truth. However, if actual chances are roughly as quantum mechanics presents them, then the analysis requires that the actual world is vastly larger than its presently surveyable portion. These implications are unwelcome, but they are not absurd. The conceptual foundations of quantum mechanics are, after all, problematic; and there is no known upper bound on the size of the actual world.

Two recalcitrant intuitions remain: first, that non-trivial chances should at least be *possible* in small worlds; second, that non-trivial cone counterpart frequency regularities would be mere coincidences. Both are damaging, but I do not think either is fatal. I sketched grounds for doubting the second intuition in the previous section. By contrast, I have made no direct attempt to discredit the first. Instead, my suggestion is that the analysis's advantages—listed above—are attractive enough that we should consider giving it up. (As I noted in §3, no Humean account of chance will save every version of the small world intuition. ‘So much the worse for Humeanism’, some will say; but there are plenty of Humeans all the same.)

The arguments in this paper do not prove that chance is cone counterpart frequency, but I think they do show that cone counterpart frequency can be taken seriously as a contender for the chance role. Finite frequentism's reputation as a non-starter is undeserved.

Funding

Millennium Research Fund, NUI Galway.

Acknowledgements

Thanks to Harjit Bhogal, Paul Dicken, Daniel Elstein, John Norton, Jack Ritchie, John Roberts, Charles Young, and (especially) two anonymous referees. Thanks also to audiences in Ankara, Cambridge, Dublin, Galway, Leeds, and Stirling, and to the Irish Philosophical Club, for feedback on talks ancestral to this paper.

*School of Humanities
National University of Ireland, Galway
Republic of Ireland
nick.tosh@nuigalway.ie*

References

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

Bigelow, J. and Pargetter, R. [1990]: *Science and Necessity*, Cambridge: Cambridge University Press.

Bigelow, J., Collins, J. and Pargetter, R. [1993]: ‘The Big Bad Bug: What are the Humean's Chances?’, *British Journal for the Philosophy of Science*, **54**, pp. 171–9.

- Bricker, P. [2001]: ‘Island Universes and the Analysis of Modality’, in G. Preyer and F. Siebelt (eds), *Reality and Humean Supervenience: Essays on the Philosophy of David Lewis*, Lanham: Rowman and Littlefield, pp. 27–55.
- Butterfield, J. [forthcoming]: ‘On Under-determination in Cosmology’, *Studies in History and Philosophy of Modern Physics*.
- Dretske, F. [1977]: ‘Laws of Nature’, *Philosophy of Science*, **44**, pp. 248–68.
- Elga, A. [2004]: ‘Defeating Dr. Evil with Self-Locating Belief’, *Philosophy and Phenomenological Research*, **69**, pp. 383–96.
- Ellis, G. [2007]: ‘Issues in the Philosophy of Cosmology’, in J. Butterfield and J. Earman (eds), *Philosophy of Physics, Part B*, Amsterdam: Elsevier, pp. 1183–285.
- Ellis, G., Maartens, R. and MacCallum, M. [2012]: *Relativistic Cosmology*, Cambridge: Cambridge University Press.
- Gillies, D. [2000]: *Philosophical Theories of Probability*, London: Routledge.
- Greaves, H. [2007]: ‘Probability in the Everett Interpretation’, *Philosophy Compass*, **2**, pp. 109–28.
- Hacking, I. [1971]: ‘Equipossibility Theories of Probability’, *British Journal for the Philosophy of Science*, **22**, pp. 339–55.
- Hájek, A. [1997]: ‘“Mises Redux”-Redux: Fifteen Argument Against Finite Frequentism’, *Erkenntnis*, **45**, pp. 209–27.
- Hájek, A. [2011]: ‘Interpretations of Probability’, in E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*,
[<http://plato.stanford.edu/archives/win2012/entries/probability-interpret/>](http://plato.stanford.edu/archives/win2012/entries/probability-interpret/).
- Hall, N. [2004]: ‘Two Mistakes About Credence and Chance’, *Australasian Journal of Philosophy*, **82**, pp. 93–111.
- Handfield, T. [2012]: *A Philosophical Guide to Chance*, Cambridge: Cambridge University Press.
- Lewis, D. [1986a]: *On the Plurality of Worlds*, Oxford: Blackwell.
- Lewis, D. [1986b]: ‘A Subjectivist’s Guide to Objective Chance’, in *Philosophical Papers, Vol. 2*, Oxford: Oxford University Press.
- Lewis, D. [1994]: ‘Humean Supervenience Debugged’, *Mind*, **103**, pp. 473–490.
- Manchak, J. [2009]: ‘Can We Know the Global Structure of Spacetime?’, *Studies in History and Philosophy of Modern Physics*, **40**, pp. 53–56.
- Meacham, C. [2010]: ‘Two Mistakes Regarding the Principal Principle’, *British*

Journal for the Philosophy of Science, **61**, 407–431.

Roberts, J. [unpublished]: ‘Laws About Frequencies’, <philsci-archive.pitt.edu/id/eprint/5058>.

Ross, J. [2010]: ‘Sleeping Beauty, Countable Additivity, and Rational Dilemmas’, *Philosophical Review*, **119**, pp. 411–47.

Salmon, W. [1989]: *Four Decades of Scientific Explanation*, Minneapolis: University of Minnesota Press.

Schaffer, J. [2003]: ‘Principled Chances’, *British Journal for the Philosophy of Science*, **54**, 27–41.

Schaffer, J. [2007]: ‘Deterministic Chance?’, *British Journal for the Philosophy of Science*, **58**, pp. 113–140.

Tooley, M. [1977]: ‘The Nature of Laws’, *Canadian Journal of Philosophy*, **7**, pp. 667–98.

Torretti, R. [1996]: *Relativity and Geometry*, Mineola: Dover.

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

Weatherson, B. [2005]: ‘Should We Respond to Evil with Indifference?’, *Philosophy and Phenomenological Research*, **70**, pp. 613–35.

Woodward, J. [1989]: ‘The Causal Mechanical Model of Explanation’, in W. Salmon and P. Kitcher (eds), *Scientific Explanation*, Minneapolis: University of Minnesota Press, pp. 357–83.

Figures




Fig. 1 Overlapping, co-oriented light cones




Fig. 2 Disjoint, oppositely-oriented light cones



Fig. 3 Disjoint, co-oriented light cones in a past-finite universe (wavy line marks boundary of spacetime)