

The Third Way on Objective Probability: A Skeptic's Guide to Objective Chance

Carl Hoefer

1/29/07

Publication version for *MIND*

Abstract

The goal of this paper is to sketch and defend a new interpretation or ‘theory’ of *objective chance*, one that lets us be sure such chances exist and shows how they can play the roles we traditionally grant them. The account is ‘Humean’ in claiming that objective chances supervene on the totality of actual events, but does not imply or presuppose a Humean approach to other metaphysical issues such as laws or causation. Like Lewis (1994) I take the Principal Principle to be the key to understanding objective chance. After describing the main features of Humean objective chance (HOC), I deduce the validity of PP for Humean chances, and end by exploring the limitations of Humean chance.

1. Introduction

The goal of this paper is to sketch and defend a new interpretation or ‘theory’ of *objective chance*, one that lets us be sure such chances exist and shows how they can play the roles we traditionally grant them.

My subtitle obviously emulates the title of Lewis’ seminal 1980 paper ‘A Subjectivist’s Guide to Objective Chance’ — while indicating an important difference in perspective. The view developed below shares two major tenets with Lewis’ last (1994) account of objective chance:

- (1) The Principal Principle tells us most of what we know about objective chance;
- (2) Objective chances are not primitive modal facts, propensities, or powers, but rather facts entailed by the overall pattern of events and processes in the *actual* world.

But it differs from Lewis’ account in most other respects.

Another subtitle I considered was ‘A Humean Guide ...’ But while the account of chance below is compatible with any stripe of Humeanism (Lewis’, Hume’s, and others’), it *presupposes* no general Humean philosophy. Only a skeptical attitude about probability itself is presupposed (as in point (2) above); what we should say about causality, laws, modality and so on is left a separate question. Still, I will label the account to be developed ‘Humean

objective chance?.

2. Why a new theory of objective chance?

Why have a *philosophical theory* of objective chance at all, for that matter? It certainly seems that the vast majority of scientists using non-subjective probabilities overtly or covertly in their research feel little need to spell out what they take objective probabilities to be. It would seem that one can get by leaving the notion undefined, or at most making brief allusions to long-run frequencies. The case is reminiscent of quantum mechanics, which the physics community uses all the time, apparently successfully, without having to worry about the measurement problem, or what — in the world — quantum states actually *represent*. Perhaps a theory is not needed; perhaps we can think of objective probability as a theoretical concept whose only possible definition is merely implicit. Sober (2004) advocates this *no-theory theory* of objective probabilities.

I find this position unsatisfactory. To the extent that we are serious in thinking that certain probabilities are objectively *correct*, or ‘out there in the world’, to the extent that we intend to use objective probabilities in explanations or predictions, we owe ourselves an account of what it is about the world that makes the imputation and use of certain probabilities correct. Philosophers are entitled to want a clear account of what objective probabilities are, just as they are entitled to look for solutions to the quantum measurement problem.¹

The two dominant types of interpretation of objective probability in recent years are *propensity* interpretations and *hypothetical or long-run frequency* interpretations. Propensity interpretations come in a wide range of flavors (as Gillies (2000) shows), and not all of them involve deep modal/causal/metaphysical implications. For example, some philosophers who advocate the theoretical term/implicit definition approach may be happy to characterize the probabilities we find in science, in some cases at least, as propensities. For the purposes of this paper, I will however restrict the term ‘propensity’ to the metaphysically robust, causally efficacious dispositional sort of property postulated by some philosophers’ accounts of objective chance.

The difficulties of propensity and long-run frequency views are well enough known

¹ Sober (2004) advocates his no-theory theory on grounds of the severe shortcomings of the traditional views. About these shortcomings we are in full agreement; but I hope to provide, below, an alternative with none of those shortcomings.

not to require much rehearsal here.² My own view of these problems is that the hypothetical frequency interpretation is metaphysically and epistemologically hopeless unless it includes some account of *what grounds the facts about hypothetical frequencies*. (Such an account tends to end up turning the interpretation into one of the other standard views: actual frequency, subjective degree of belief, or propensity.) And propensity views, while still actively pursued by many philosophers, add a very peculiar new sort of entity, property, or type of causation to the world.³ One can argue at length about whether or not this makes propensities metaphysically suspect. I think it more clear that propensities are epistemologically hopeless (i.e., one can only claim that statistics are a reliable guide to propensities *via* arguments that are all, in the end, invalid — usually, circular). In section 5.2 a closely related problem for propensity views of chance will be discussed: their inability to justify Lewis' Principal Principle. For now I will just register my dissatisfaction with both hypothetical frequency and propensity views of chance; those who share at least some of my worries will hopefully agree with me that a third way obviating at least some of their problems would be worth spelling out.

Of course, a third way not suffering from *any* of the problems alluded to above is already available: the actual frequency interpretation (sometimes called 'finite frequency'). The defects of this view are usually vastly overestimated, and its virtues underappreciated. Indeed, the actual frequency interpretation is the only natural starting point for an empiricist — or skeptical — approach to objective chance. Both Lewis' (1994) theory and the theory sketched below are in a sense 'sophistications' of the actual frequency approach. They try to fit better with common sense, with certain uses of probability in sciences such as quantum mechanics and statistical mechanics, and with classical gambling devices. But the grounding of all objective chance in matters of actual (non-modal, non-mysterious) fact is shared by all three approaches.

The goal of this paper is thus to develop and defend a 'third way' (different from Lewis' and from standard actual frequentism) among 'third way' approaches (neither propensity- nor hypothetical frequency-based). The chances to be described here *exist* — whether or not determinism is true, and whether or not there exist such things as primitive propensities or probabilistic causal capacities in nature. The interpretation can thus be

² See Hájek (2003).

³ D. H. Mellor (1994), *The Facts of Causation*, contains an extended and thorough exposition and defense of a theory of causation based on a propensity view of objective chance.

defended without making any contentious metaphysical assumptions. The positive arguments for the view will turn on two points: first, its coherence with the main uses of the notion of objective chance, both in science and in other contexts; and second, its ability to justify the Principal Principle.

3. Correcting the Subjectivist's Guide: Lewis' program, 1980 - 1994

Because Lewis' approach to objective chance is well-known, it is perhaps best to introduce his view, and work toward the proper skeptical/Humean view by correcting Lewis' at several important places.⁴

3.1 PP

As noted above, one of the two shared fundamentals of Lewis' interpretation and mine is the claim that the Principal Principle (PP) tells us most of what we know about objective chance. PP can be written:

$$Cr(A|XE) = x \quad (PP)$$

Here '*Cr*' stands for 'credence', i.e., a subjective probability or degree of belief function. *A* is any proposition you like, in the domain of the objective chance or objective probability function *Pr*. *X* is the proposition stating that the objective chance of *A* being the case is *x*, i.e., *X* = '*Pr*(*A*) = *x*'. Finally, *E* is any 'admissible' evidence or knowledge held by the agent whose subjective probability is *Cr*.⁵ The idea contained in PP, an utterly compelling idea, is this: if all you know about whether *A* will occur or not is that *A* has some objective probability *x*, you ought to set your own degree of belief in *A*'s occurrence to *x*. Whatever else we may say about objective chance, it *has to* be able to play the PP role: PP captures, in essence, what objective chances *are for*, why we want to know them.

Crucial to the reasonableness of PP is the limitation of *E* to 'admissible' information. What makes a proposition admissible or non-admissible? Lewis defined admissibility completely and correctly in 1980 — though he considered this merely a vague, first-approximation definition:

⁴ For a clear recent exposition and defense of Lewis' approach, see Loewer (2004).

⁵ Throughout I will follow Lewis in taking chance as a probability measure over a sub-algebra of the space of all propositions. Intuitively speaking, the propositions say that a certain outcome occurs in a certain chance setup. Unlike [what many assume about] rational credence, the probability measure should *not* be assumed to extend over all, or even most, of this whole proposition space. Here we need only assume that the domain of *Cr* includes at least the domain of *Pr* and enough other stuff to serve as suitable *X*'s and *E*'s.

Admissible propositions are the sort of information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes.

(Lewis 1980, p. 92)

This is almost exactly right. When is it rational to make one's subjective credence in A exactly equal to (what one takes to be) the objective chance of A ? When one simply has *no information* tending to make it reasonable to think A true or false, except by way of making it reasonable to think that the objective chance of A has a certain value. If E has any such indirect information about A , i.e., information relevant to the objective chance of A , such information is cancelled out by X , since X gives A 's objective chance outright. Here is a slightly more precise definition of

Admissibility: Propositions that are admissible with respect to outcome-specifying propositions A_i contain only the sort of information whose impact on reasonable credence about outcomes A_i , if any, comes entirely by way of impact on credence about the chances of those outcomes.

This definition of admissibility is clearly consonant with PP's expression of what chance is for, namely guiding credence when no better guide is available. The admissibility clause in PP is there precisely to exclude the presence of any such 'better guide'.

Notice that in this definition of admissibility, there is no mention of past or future, complete histories of the world at a given time, or any of the other apparatus developed in Lewis (1980) to substitute a precise-looking definition of admissibility in place of the correct one. We will look at some of that apparatus below, but it is important to stress here that none of it is needed to understand admissibility completely. Lewis' substitution of a precise 'working characterization' of admissibility in place of the correct definition seems to be behind two important aspects of his view of objective chance that I will reject below: first, the alleged 'time-dependence' of objective chance; second, the alleged incompatibility of chance and determinism.⁶

3.2 Time and Chance.

Lewis claims, as do most propensity theorists, that the past is 'no longer chancy'. If A

⁶ It also created mischief in other ways. For example, in the context of his "reformulated" PP, which we will see below, it caused Lewis to believe for a long time that the true objective chances in a world had to be *necessary*, i.e., never to have had any chance of not being the case. This misconception helped delay his achievement of his final view by well over a decade.

is the proposition that the coin I flipped at noon yesterday lands heads, then the objective chance of A is *now* either zero or one — depending on how the coin landed. (It landed tails.) Unless one is committed to the ‘moving now’ conception of time, and the associated view that the past is ‘fixed’ whereas the future is ‘open’ (as propensity theorists, I argue elsewhere, seem to be)⁷, there is little reason to make chance a time-dependent fact in this way. I prefer the following way of speaking: my coin flip at noon yesterday was an instance of a chance setup with two possible outcomes, each having a definite objective chance. It was a chance event. The chance of heads was $\frac{1}{2}$. So $\frac{1}{2}$ is the objective chance of A . It still is; the coin flip is and always was a chance event. Being to the past of me-now does not alter that fact, though as it happens I now know A is false.

PP, with admissibility properly understood, is perfectly compatible with taking chance as not time-dependant. It seems at first incompatible, because of the ‘working characterization’ of admissibility Lewis gives, which says that at any given time t , any historical proposition — that is, any proposition about matters of fact at or before t — is admissible. *Now*, a day after the flip, that would make $\neg A$ itself admissible; and of course $Cr(A|\neg AE)$ had better be zero (see (1980), p. 98). But clearly this violates the correct definition of admissibility. $\neg A$ carries maximal information as to A ’s truth, and not by way of any information about A ’s objective chance; so it is inadmissible. My credence about A is now unrelated to its objective chance, because I know that A is false. But as Ned Hall (1994) notes, this has nothing intrinsically to do with time. If I had a reliable crystal ball, my credences about some *future* chance events might similarly be disconnected from what I take their chances to be. (Suppose my crystal ball shows me that the next flip of my lucky coin will land ‘heads’. Then my credence in the proposition that it lands ‘tails’ will of course be zero, or close to it.)

Why did Lewis not stick with his loose, initial definition of admissibility? Why did he instead offer a complicated ‘working definition’ of admissibility in its place? One reason, I think, is that Lewis (1980) was trying to offer an account of objective chance that mimics the way we think of chance when we think of them as *propensities*, *making* things happen (or *unfold*) in certain ways. If we think of a coin-flipping setup as having a propensity (of strength $\frac{1}{2}$) to make events unfold a certain way (coin-lands-heads), then once that propensity

⁷ “*Time and Chance Propensities*” manuscript. Oddly, Lewis rather explicitly embraces a moving-now and branching-future picture in “A Subjectivist’s Guide”. He never, to my knowledge, discusses how such a picture can be reconciled with relativistic physics.

has done its work, it's all over. The past is fixed, inert, and free of propensities (now that they've all 'sprung' and done their work, so to speak). These metaphors are part and parcel of the notion of chance as a propensity, and oddly enough they seem to have a grip on Lewis too, despite his blunt rejection of propensities (particularly in (1994)). We will see further evidence of this below.

There is a real asymmetry in the amount and quality of information we have about the past, versus the future. We tend to have lots of inadmissible information about past chance events, very little (if any) inadmissible information about future chance events. But there need be nothing asymmetric or time-dependent in the chance events themselves.⁸ Taking PP as the guide to objective chance illustrates this nicely. Suppose you want to wager with me, and I propose we wager about yesterday's coin toss, which I did myself, recording the outcome on a slip of paper. I tell you the coin was fair, and you believe me. Then your credences should be $\frac{1}{2}$ for both A and $\neg A$, and it's perfectly rational for you to bet either way. (It would only be irrational for you to let *me* choose which way the bet goes.) The point is just this: if you have no inadmissible information about whether or not A , but you do know A 's objective chance, then your credence should be equal to that chance — whether A is a past or future event. Lewis (1980) derives the same conclusions about what you should believe, using the Principal Principle on his time-dependent chances in a roundabout way. I simply suggest we avoid the detour.⁹

3.3 *The Best System Analysis of laws and chance.*

David Lewis applies his Humeanism about all things modal across the board: counterfactuals, causality, laws, and chance all are analyzed as results of the vast pattern of actual events in the world. This program goes under the name 'Humean Supervenience', HS for short. Fortunately we can set aside Lewis' treatments of causation and counterfactuals here. But his analysis of laws of nature must be briefly described, as he explicitly derives objective chances and laws together as part of a single 'package deal'.

⁸ There *need be* no time asymmetry to objective chances, but often there *is* a presupposed time-directedness. Typically chance setups involve a temporal asymmetry, the "outcome" occurring after the "setup" conditions are instantiated. But in no case do the categories of past, present or future (as opposed to before/after) need to be specified.

⁹ By avoiding the detour, we also avoid potential pitfalls with backward-looking chances, such as are utilized in Humphrey's objection to propensity theories of chance (see Humphreys (2004)).

Take all deductive systems whose theorems are true. Some are simpler, better systematized than others. Some are stronger, more informative, than others. These virtues compete: an uninformative system can be very simple, an unsystematized compendium of miscellaneous information can be very informative. The best system is the one that strikes as good a balance as truth will allow between simplicity and strength. ... A regularity is a law *iff* it is a theorem of the best system. (1994, p. 478)

Lewis modifies this BSA account of laws so as to make it able to incorporate probabilistic laws:

. . . we modify the best-system analysis to make it deliver the chances and the laws that govern them in one package deal. Consider deductive systems that pertain not only to what happens in history, but also to what the chances are of various outcomes in various situations — for instance, the decay probabilities for atoms of various isotopes. Require these systems to be true in what they say about history. We cannot yet require them to be true in what they say about chance, because we have yet to say what chance means; our systems are as yet not fully interpreted. ...

As before, some systems will be simpler than others. Almost as before, some will be stronger than others: some will say either what will happen or what the chances will be when situations of a certain kind arise, whereas others will fall silent both about the outcomes and about the chances. And further, some will fit the actual course of history better than others. That is, the chance of that course of history will be higher according to some systems than according to others. ...

The virtues of simplicity, strength and fit trade off. The best system is the system that gets the best balance of all three. As before, the laws are those regularities that are theorems of the best system. But now some of the laws are probabilistic. So now we can analyze chance: the chances are what the probabilistic laws of the best system say they are.' (1994, p. 480)

A crucial point of this approach, which makes it different from actual frequentism, is that considerations of symmetry, simplicity, and so on can make it the case that (a) there are objective chances for events that occur seldom, or even never; and (b) the objective chances may sometimes diverge from the actual frequencies even when the actual 'reference class'

concerned is fairly numerous, for reasons of simplicity, fit of the chance law with other laws of the System, and so on. My account will preserve this aspect of Lewis' Best Systems approach. Law facts and other sorts of facts, whether supervenient on Lewis' HS-basis or not, may, together with some aspects of the HS-basis 'pattern' in the events of the world, make it the case that certain objective chances exist, even if those chances are not grounded in that pattern alone. Examples of this will be discussed in section 4.¹⁰

Analyzing laws and chance together as Lewis does has at least one very unpleasant consequence. If this is the right account of objective chances, then *there are* objective chances only if the best system for our world says there are. But we are in no position to know whether this is in fact the case, or not; and it's not clear that further progress in science will substantially improve our epistemic position on this point. Just to take one reason for this, to be discussed further below: the Lewisian best system in our world, for all we now know, may well be deterministic, and hence (at first blush) need no probabilistic laws at all.¹¹ If that is the case, then on Lewis' view, contrary to what we think, there aren't any objective chances in the world at all.

This is a disastrous feature of Lewis' account, for obvious reasons. Objective probabilities *do* exist; they exist in lotteries, in gambling devices and card games, and possibly even in my rate of success at catching the 9:37 train to work every weekday. In science, they occur in the statistical data generated in many physical experiments, in radioactive decay, and perhaps in thermodynamic approaches to equilibrium (e.g. the ice melting in your cocktail). Any view of chance that implies that there may or may not be such a thing after all — it depends on what the laws of nature turn out to be — must be mistaken.¹² Or put another way: the notion of 'objective chance' described by the view is not the notion at work in actual science and in everyday life.

It is understandable that some philosophers who favor a propensity view should hold this view that we don't know, and may never know, whether there are such things as

¹⁰ See the discussion of 'stochastic nomological machines', section 4.1.

¹¹ Lewis points to the success of quantum mechanics as some reason to think that probabilistic laws are likely to hold in our world. But a fully deterministic version of quantum mechanics exists and is growing steadily more popular, namely Bohmian mechanics. Suppes (1993) offers general arguments for the conclusion that we may never be able to determine whether nature follows deterministic or stochastic laws.

¹² Notice that almost no philosophers today would be willing to make a parallel assertion about *causation*, namely that it may or may not be "real" in the world, depending on what view of laws is ultimately right.

objective chances (though it is, I think, equally disastrous for them). It is less clear why Lewis does so. On the face of it, it *is* a consequence of his ‘package deal’ strategy: chances are whatever the BSA laws governing chance say, which is something we may never be able to know. But if we, as I urge, set aside the question of the nature of laws, and think of the core point of Lewis’ Humean approach to chance, it is just this: objective chances are simply facts following from the vast pattern of events that comprise the history of this world. *Some* of the chances to be discerned in this pattern may in fact be consequences of natural laws; but why should *all* of them be?

Thinking of the phenomena we take as representative of objective chance, the following path suggests itself. There may be some probabilistic laws of nature; we may even have discovered some already. But there are also *other* kinds of objective chances, that arguably do not follow from laws of nature (BSA or otherwise): probabilities of drawing to an inside straight, getting lung cancer if one smokes heavily, being struck by lightning in Florida, and so on. Only a very strong reductionist would think that such probabilities must somehow be *derivable* from the true physical laws of our world, if they are to be genuinely objective probabilities; so only a strong reductionist bias could lead us to reject such chances if they *cannot* be so derived. And why not accept them? The overall pattern of actual events in the world can be such as to make these chances exist, whether or not they deserve to be written in the Book of Laws, and whether or not they logically follow *from* the Book. As we will see below in sections 4 and 5, they are there because they are capable of playing the objective-chance role given to us in the Principal Principle.

Suppose we do accept such objective chances not (necessarily) derivable from natural laws. That is, we accept non-lawlike, but still objective, chances, because they simply are *there* to be discerned in the mosaic of actual events (as, for Lewis, the laws of nature themselves are). Let’s suppose then that Lewis could accept these further non-lawlike chances *alongside* the chances (if any) dictated by the Best System’s probabilistic laws. Now we can turn to the question of whether objective chances exist if determinism is true.

3.4 Chance and Determinism. Lewis considers determinism and the existence of non-trivial objective chances to be incompatible. I believe this is a mistake.

In 1986 Lewis discussed this issue, responding to Isaac Levi’s charge (with which I am, of course, in sympathy) that it is a pressing issue to say how to reconcile determinism

with objective chances.¹³ In his discussion of this issue ((1986), pp. 117 - 121) Lewis does not *prove* this incompatibility. Rather he seems to take it as obvious that, if determinism is true, then all propositions about event outcomes have probability zero or one, which then excludes nontrivial chances. How might the argument go? We need to use Lewis' working definition of admissibility and his revised formulation of PP,

$$(PP2) \quad C(A|H_{tw}T_w) = x = Pr(A)$$

in which H_{tw} represents the complete history of the world w up to time t , and T_w represents the 'complete theory of chance for world w '. T_w is a vast collection of 'history to chance conditionals'. A history-to-chance conditional has as antecedent a proposition like H_{tw} , specifying the history of world w up to time t ; and as consequent, a proposition like X , stating what the objective chance of some proposition A is. The entire collection of the true history-to-chance conditionals is T_w , and is what Lewis calls the 'theory of chance' for world w .

Suppose that L_w are the laws of world w , and that we take them to be admissible. Now we can derive the incompatibility of chances with determinism from this application of PP2:

$$PP2 \quad C(A|H_{tw}T_wL_w) = x = Pr(A)$$

Determinism is precisely the determination of the whole future of the world from its past up to a given time (H_{tw}) and the laws of nature (L_w). But if H_{tw} and L_w together *entail* A , then by the axioms, $Cr(A|H_{tw}T_wL_w)$ must be equal to 1 (and *mutatis mutandis*, zero if they entail $\neg A$). A contradiction can only be avoided if all propositions A have chances of zero or one. Thus PP2 seems to tell us that non-trivial chances are incompatible with deterministic laws of nature.

But this derivation is spurious; there is a violation of the correct understanding of admissibility going on here. For if $H_{tw}L_w$ entails A , then it has a *big* (maximal) amount of information pertinent as to whether A , and not by containing information about A 's objective chance!¹⁴ So $H_{tw}L_w$, so understood, must be held inadmissible, and the derivation of a

¹³ Levi (1983).

¹⁴ $H_{tw}L_wT_w$ may entail that A has chance 1. That's beside the point; if it's a case of normal deterministic entailment, $H_{tw}L_wT_w$ *also* entail A itself. And that is carrying information relevant to the truth of A other than by carrying information about A 's objective chance.

contradiction fails.

PP, properly understood, does not tell us that chance and determinism are incompatible. But there is another way we might explain Lewis' assumption that they are incompatible. It has to do with the 'package deal' about laws. Lewis may have thought that deterministic laws are automatically as strong as strong can be; hence if there is a deterministic best system, it can't possibly have any probabilistic laws in its mix. For they would only detract from the system's simplicity without adding to its already maxed-out strength.

If this is the reason Lewis maintained the incompatibility, then again I think it is a mistake. Deterministic laws may not after all be the last word in strength — it depends how strength is defined in detail. Deterministic laws say, in one sense, almost nothing about what actually happens in the world. They need initial and boundary conditions in order to entail anything about actual events. But are such conditions to form part of Lewis' axiomatic systems? If they can count as part of the axioms, do they increase the complexity of the system infinitely, or by just one 'proposition', or some amount in between? Lewis' explication does not answer these questions, and intuition does not seem to supply a ready answer either. What I urge is this: it is not at all obvious that the strength of a deterministic system is intrinsically maximal and hence cannot be increased by the addition of further probabilistic laws. If this is allowed, then determinism and non-trivial objective chances are not, after all, incompatible in Lewis' system.¹⁵ Nor, of course, are they incompatible on the account I develop below.

3.5 Chance and credence. Lewis (1980) claims to prove that objective chance is a species of probability, that is, obeys the axioms of probability theory, in virtue of the fact that PP equates chances with certain ideal subjective credences, and it is known that such ideal credences obey the axioms of probability.

'A reasonable initial credence function is, among other things, a probability distribution: a non-negative, normalized, finitely additive measure. It obeys the laws of mathematical probability theory. . . . Whatever comes by conditionalizing from a

¹⁵ Loewer (2001, 2004) has refined and advocated a Lewisian Best Systems account of chance, and he comes to the same conclusion: determinism can coexist with nontrivial chances.

probability distribution is itself a probability distribution. Therefore a chance distribution is a probability distribution.’ (1980, p. 98).

This is one of the main claims of the earlier paper motivating the title ‘A *subjectivist’s* guide ...’ But it seems to me that this claim must be treated carefully. Ideal rational degrees of belief are shown to obey the probability calculus only by the Dutch book argument, and this argument seems to me only sufficient to establish a ‘*ceteris paribus*’ or ‘*prima facie*’ constraint on rational degrees of belief. The Dutch book argument shows that an ideal rational agent *with no reasons to have degrees of belief violating the axioms* (and hence, no reason not to accept any wagers valued in accord with these credences) is irrational if he/she nevertheless does have credences that violate the axioms. By no means does it show that there can never be a reason for an ideal agent to have credences violating the axioms. Much less does it show that finite, non-ideal agents such as ourselves can have no reasons for credences violating the axioms. Given this weak reading of the force of the Dutch book argument, then, it looks like a slender basis on which to base the requirement that objective probabilities should satisfy the axioms.

Lewis’ chances obey the axioms of probability just in case T_w makes them do so. It’s true that, given the role chances are supposed to play in determining credences *via* PP, they ought *prima facie* to obey the axioms. But there are other reasons for them to do so as well. Here is one: the chances have, in most cases, to be close to the actual frequencies (again, in order to be able to play the PP role), and actual frequencies are guaranteed to obey the axioms of probability.¹⁶ So while it is true in a broad sense that objective chances must obey the axioms of probability because of their intrinsic connection with subjective credences, it is an oversimplification to say simply that objective chances must obey the axioms because PP equates them with (certain sorts of) ideal credences, and ideal credences must obey the axioms.

Secondly, on either Lewis’ or my approach to chance, it’s not really the case that objective chances are ‘objectified subjective credences’ as Lewis (1980) claims. This phrase makes it sound as though one starts with subjective credences, does something to them to remove the subjectivity (according to Lewis: conditionalizing on $H_{tw}T_w$), and what is left then plays the role of objective chance. In his reformulation of the PP, Lewis presents the principle as if it were a universal generalization over all reasonable initial credence functions

¹⁶ Setting aside worries that may arise when the actual outcome classes are infinite.

(RICs):

‘Let C be any reasonable initial credence function. Then for any time t , world w , and proposition A in the domain of P_{tw}

$$[\text{PP2:}] \quad P_{tw}(A) = C(A|H_{tw}T_w).$$

In words: the chance distribution at a time and a world comes from any reasonable initial credence function by conditionalizing on the complete history of the world up to the time, together with the complete theory of chance for the world.’ (1980/6, pp. 97-8).

Read literally, as a universal generalization, this claim is just false. There are some RICs for which the equation given holds, and some for which it does not, and that is that. It is no part of Lewis’ earlier definition of what it is for an initial credence function to be reasonable, that it must respect PP! But, clearly, any RIC that does not conform to PP will fail to set credences in accordance with the equation above.

PP is of course meant to be a principle of rationality, and so perhaps we *should* build conformity to it into our definition of the ‘reasonable’ in RIC. This may well be what Lewis had in mind (see (1980), pp. 110 –11). Then the quote from Lewis above becomes true by definition. Nevertheless the impression it conveys, that somehow the *source* of objective chances is to be found in RICs, remains misleading.

Humean objective chances are simply a result of the overall pattern of events in the world, an aspect of that pattern guaranteed, as we will see, to be useful to rational agents in the way embodied in PP. But they do not start out as credences; they *determine* what may count as ‘reasonable’ credences, via PP. In Lewis’ later treatments this is especially clear. The overall history of the world gives rise to one true ‘theory of chance’ T_w for the world, and this theory says what the objective chances are wherever they exist.

4. What Humean objective chance is.

So far I have been laying out my Humean view of chance indirectly, by correcting a series of (what I see as) errors in Lewis’ treatment. Now let me give a preliminary, but direct, statement of the interpretation I advocate. This approach has much in common with Lewis’ as amended above — but without the implied reductionism to the microphysical.

4.1 The basic features

Chances are in the first instance probabilities of outcomes conditional on the instantiation of a proper chance setup, and additionally such probabilities as can be derived from the basic chances with the help of logic and the probability axioms. I follow Alan Hájek (2003b) in considering conditional chance as the more basic notion; the ‘definition’ of conditional probability,

$$\Pr(A|B) = \Pr(A \wedge B)/\Pr(B)$$

is a constraint to be respected, where the unconditional probabilities are well-defined, but it is no complete analysis of the relationship. As Hájek reminds us, the probability that I get heads given that I flip a fair coin is 1/2; but the probability that I flip the coin? Typically, that does not exist. It would be better to write the above constraint like this:

$$\Pr(A|B \wedge C) = \frac{\Pr(A \wedge B|C)}{\Pr(B|C)}$$

to remind ourselves that objective chances must always be conditional on a chance setup. But where no misunderstanding will arise, the conditionalization on the (instantiation of the) chance setup may be omitted for brevity, as it was in sections 1 - 3 above.

The domain over which the $\Pr(_|_)$ function ranges may be quite limited, and is determined by what the Humean mosaic in our world is like.

Chances are constituted by the existence of patterns in the mosaic of events in the world. The patterns have nothing (directly) to do with time or the past/future distinction, and nothing to do with the nature of laws or determinism. Therefore, neither does objective chance.

From now on, I will call this kind of chance that I am advocating ‘Humean objective chance’ (or HOC for short). But it should be kept in mind that the Humeanism only covers chance itself; not laws, causation, minds, epistemology, or anything else.

These patterns are such as to make the adoption of credences identical to the chances rational in the absence of better information, in a sense to be explored below. Sometimes the chances are just finite/actual frequencies; sometimes they are an idealization or model that ‘fits’ the pattern, but which may not make the chances strictly equal to the actual frequencies. (This idea of ‘fit’ will be explored through examples, below and in section 5).

It appears to be a fact about actual events in our world that, at many levels of scale (but especially micro-scale), events look ‘stochastic’ or ‘random’, with a certain stable distribution over time; this fact is crucial to the grounding of many objective chances. I call

this the Stochasticity Postulate, SP. We rely on the truth of SP in medicine, engineering, and especially in physics. The point of saying that events ‘look stochastic’ or ‘look random’, rather than saying they *are* stochastic or random, is dual. First, I want to make clear that I am referring here to ‘product’ randomness, not ‘process’ randomness (in Earman’s (1986) terminology). Sequences of outcomes, numbers and so on can look random even though they are generated by (say) a random-number generating computer program. For the purposes of our Humean approach to chance, looking random is what matters. Second, randomness in the sense intended is a notion that has resisted perfect analysis, and is especially difficult when one deals with finite sequences of outcomes. Nevertheless, we all know roughly how to distinguish a random-looking from a non-random-looking sequence, if the number of elements is high enough. Our concern at root, of course, is with the applicability of PP. Sets or sequences of events that are random-looking with a stable distribution will be such that, if forced to make predictions or bets about as-yet-unobserved parts of them (e.g., the next ten tosses of a fair coin), we can do no better than adjust our expectations in accord with the objective chance distribution.

Some stable, macroscopic chances that supervene on the overall pattern are explicable as regularities guaranteed by the structure of the assumed chance setup. These cases will be dubbed *Stochastic Nomological Machines* (SNM’s), in an extension of Nancy Cartwright’s (1999) notion of a nomological machine. A nomological machine is a stable mechanism that *generates a regularity*. A SNM will be a stable chance setup or mechanism that generates a probability (or distribution). The best examples of SNM’s, unsurprisingly, are classical gambling devices: dice on craps tables, roulette wheels, fair coin tossers, etc. For these and many other kinds of chance setup, we can, in a partial sense, deduce their chancy behavior from their setup’s structure and the correctness of the Stochasticity Postulate at the level of initial conditions and external influences. Not all genuine objective chances have to be derivable in this way, however. We will consider examples of objective chances that are simply *there*, to be discerned, in the patterns of events.

Nevertheless, any objective chance should be thought of as tied to a well-defined chance *setup* (or reference class, as it is sometimes appropriate to say). The patterns in the mosaic that constitute Humean chances are regularities, and regularities of course link one sort of thing with another. In the case of chance, the linkage is between the well-defined chance setup and the possible outcomes for which there are objective probabilities.

Some linking of objective probabilities to a setup, or a reference class, is clearly

needed. Just as a Humean about laws sees ordinary laws as, in the first instance, patterns or regularities — in the mosaic, whenever F, then G — so the Humean about chances sees them as patterns or regularities in the mosaic also, albeit patterns of a different and more complicated kind: whenever S, $\Pr(A) = x$.

Two further comments on the notion of ‘chance setup’ are needed. First, ‘well-defined’ does not necessarily mean non-vague. ‘A fair coin is flipped decently well and allowed to land undisturbed’ may be vague, but nevertheless a well-defined chance setup in the sense that matters for us (it excludes lots of events quite clearly, and includes many others equally clearly). Second, my use of the term ‘chance setup’, which is historically linked to views best thought of as propensity accounts of chance (e.g. Popper, Giere, Hacking) should not be taken as an indication that my goal is to offer a Humean theory that mimics the features of propensity theories as closely as possible. Rather, making chances conditional on the instantiation of a well-defined setup is necessary once we reject Lewis’ time-indexed approach. For Lewis, a non-trivial time-indexed objective probability $Pr_t(A)$ is, in effect, the chance of A occurring given the instantiation of a big setup: the entire history of the world up to time t . Since I reject Lewis’ picture of the world unfolding in time under in accordance with chancy laws, I don’t have his big implicit setup. So I need to make my chances explicitly linked to the appropriate (typically small, local) setup. Again, and unlike propensity theorists, I do not insist that the chance-bearing ‘outcome’ must come after the ‘setup’ (though almost all chances we care about have this feature).

To understand the notion of patterns in the mosaic, an analogy from photography may be helpful. A black and white photo of a gray wall will be composed of a myriad of grains, each of which is either white or black. Each grain is like a particular ‘outcome’ of a chance process. If the gray is fairly uniform, then it will be true that, if one takes any given patch of the photo, above a certain size, there will be a certain ratio of white to black grains (say 40%), and this will be true (within a certain tolerance) of every similar-sized patch you care to select. If you select patches of smaller size, there will be more fluctuation. In a given patch of only 12 grains, for example, you might find 8 white grains; in another, only 2; and so on. But if you take a non-specially-selected collection of 30 patches of 12 grains, there will again be close to 40% whites among the 360 total grains. The mosaic of grains in the photo is exactly analogous to the mosaic of events in the real world that found an objective chance such as, e.g., the chance drawing a spade in a well-shuffled deck. In neither case does one

have to postulate a propensity, or give any kind of explanation of exactly how each event (black, white; spade, non-spade) came to be, for the chance (the grayness) to be objective and real.

Of course, like photos, patterns in the mosaic of real world outcomes can be much more complex than this. There can be patterns more complex and interesting than mere uniform frequencies made from black and white grains (not to speak of colored grains). There may be repeated variations in shading, shapes, regularities in frequency of one sort of shape or shade following another (in a given direction), and so on. (Think of these as analogies for the various types of probability distributions found to be useful in the sciences.)

There may be regularities that can only be discerned from a very far-back perspective on a photograph (e.g., a page of a high school yearbook containing row after row of photos of 18 year olds, in alphabetical order — so that, in the large, there is a stable ratio of girl photos to boy photos on each page, say 25 girls to 23 boys). This regularity may be associated with an SNM — it depends on the details of the case — but in any case, the regularity about boys and girls on pages is objectively there, and makes it reasonable to bet ‘girl’ if offered an even-money wager on the sex of a person whose photo will be chosen at random on a randomly selected page.

4.2 Examples

Not every actual frequency, even in a clearly defined reference class, is an objective chance. Conversely, not every chance setup with a definite HOC need correspond to a large reference class with frequencies matching the chances. I will illustrate the main features of Humean objective chances through a few examples, and then extract the salient general features.

1. Chance of 00 on a roulette wheel.

I begin with an example of a classic gambling device, to illustrate several key aspects of HOC. The objective chance of 00 is, naturally, $x = 1/[\text{the number of slots}]$. What considerations lead to this conclusion? (We will assume, here and throughout unless otherwise specified, that the future events (and past events outside our knowledge) in our world are roughly what we would expect based on past experience). First of all, presumably there is the actual frequency, very close to x . But that is just one factor, arguably not the most important. (There has perhaps never been a roulette wheel with 43 slots; but we believe that if we made one, the chance of 00 landing on it would be $1/43$.)

Consider the type of chance setup a roulette wheel exemplifies. First we have spatial symmetry, each slot on the wheel having the same shape and size as every other. Second, we have (at least) four elements of *randomization* in the functioning of the wheel/toss: first, the spinning (together with facts about human perception and lack of concern) gives us randomness of the initial entry-point of the ball, i.e., the place where it first touches. The initial trajectory and velocity of the ball is also fairly random, within a spread of possibilities. The mechanism itself is a good approximation to a classical chaotic system — that is, it embodies sensitive dependence on initial conditions. Finally, the whole system is not isolated from external perturbations (gravitational, air currents, vibrations of the table from footfalls and bumps, etc.), and these perturbations also can be seen as a further randomizing factor. The dynamics of the roulette wheel are fairly Newtonian, and it is therefore natural to expect that the results of spins with so many randomizing factors, both in the initial conditions and in the external influences, will be distributed stochastically but fairly uniformly over the possible outcomes (number slots). And this expectation is amply confirmed by the actual outcome events, of course.

The alert reader may be concerned at my use of ‘randomness’ and ‘randomizing’, when these notions may appear bound up with the notion of chance itself (and maybe, worse, a propensity understanding of chance). But recall, for the Humean about chance, all randomness is product randomness.¹⁷ Randomness of initial conditions is thus nothing more than stochastic-lookingness of the distribution of initial (and/or boundary) conditions, displaying a definite and stable distribution at the appropriate level of coarse-graining. The randomness adverted to earlier in my description of the roulette wheel is just this, a Humean-compatible aspect of the patterns of events at more-microscopic levels. Here we see the Stochasticity Postulate in action: it grounds our justified expectation that roulette wheels will be unpredictable and will generate appropriate statistics in the outcomes.

2. *Good coin flips*. Not every flip of a coin is an instantiation of the kind of stochastic nomological machine we implicitly assume is responsible for the fair 50/50 odds of getting heads or tails when we flip coins for certain purposes. Young children’s flips often turn the coin only one time; flips where the coin lands on a grooved floor frequently fail to yield

¹⁷ The term ‘product’ randomness is unfortunate, since it seems to imply that the randomness involved has been produced by some process. HOC prescind from any such assumption.

either heads or tails; and so on. Yet there is a wide range of circumstances that do instantiate the SNM of a fair coin flip, and we might characterize the machine roughly as follows:

1. The coin is given a goodly upward impulse, so that it travels at least a foot upward and at least a foot downward before being caught or bouncing;
2. The coin rotates while in the air, at a decent rate and a goodly number of times;
3. The coin is a reasonable approximation to a perfect disc, with reasonably uniform density and uniform magnetic properties (if any);
4. The coin is either caught by someone not trying to achieve any particular outcome, or is allowed to bounce and come to rest on a fairly flat surface without interference
5. If multiple flips are undertaken, the initial impulses should be distributed randomly over a decent range of values so that both the height achieved and the rate of spin do not cluster tightly around any particular value.

Two points about this SNM deserve brief comment. First, this characterization is obviously vague. That is not a defect. If you try to characterize what is an *automobile*, you will generate a description with similar vagueness at many points. This does not mean that there are no automobiles in reality. Second, here too the ‘randomness’ adverted to is meant only as random-lookingness, and implies nothing about the *processes* at work. For example, we might instantiate our SNM with a very tightly calibrated flipping machine that chooses (a) the size of the initial impulse, and (b) the distance and angle off-center of the impulse, by selecting the values from a pseudo-random number generating algorithm. In ‘the wild’, of course, the reliability of nicely randomly-distributed initial conditions for coin flips is, again, an aspect of the Stochasticity Postulate.¹⁸

¹⁸ Sober (2004) discusses a coin-flipping setup of the sort described here, following earlier analyses by Keller and Diaconis based on Newtonian physics. Sober comes to the same conclusion: if the distribution of initial conditions is appropriately random-looking (and in particular, distributed approximately equally between IC’s leading to heads and IC’s leading to tails), then the overall system is one with an objective chance of 0.5 for heads.




Diagram 1: Diaconis' Newtonian coin-flip model.

The diagram above from Diaconis (1998) illustrates this, on a Newtonian-physics model of coin flipping. Initial conditions with ω (angular velocity) and V/g (vertical velocity) falling in a black area land heads; those in white areas land tails. (The coins are all flipped starting heads-up.) From the SP we expect the initial angular velocities and vertical velocities to be scattered in a random-looking distribution over the square (not an *even* distribution, but rather random-looking in the sense of not having any correlation with the black and white bands). When this is the case, the frequency of heads (black bands) and tails (white) will be approximately 50/50.

3. *The biased coin flipper.* The coin flip SNM just described adds little to the roulette wheel case, other than a healthy dose of vagueness (due to the wide variety of coin flippers in the world). But the remarks about a coin-flipping machine point us toward the following, more interesting SNM. Suppose we take the tightly-calibrated coin flipper (and 'fair' coin) and: make sure that the coins land on a very flat and smooth, but very mushy surface (so that they never, or almost never, bounce); try various inputs for the initial impulses until we find one that regularly has the coin landing heads when started heads-up, as long as nothing disturbs the machine; and finally, shield the machine from outside disturbances. Such a machine can no doubt be built (probably has been built, I would guess), and with enough engineering sweat can be made to yield as close to chance = 1.0 of heads as we wish.

This is just as good an SNM as the ordinary coin flipper, albeit harder to achieve in practice. Both yield a regularity, namely a determinate objective probability of the outcome heads. But it is interesting to note the differences in the kinds of ‘shielding’ required in the two cases. In the first, what we need is shielding from conditions that bias the results (intentional or not). Conditions i, ii, iv and v are all, in part at least, shielding conditions. But in the biased coin flipper the shielding we need is of the more prosaic sort that many of our finely tuned and sensitive machines need: protection from bumps, wind, vibration, etc. Yet, unless we are aiming at a chance of heads of precisely 1.0, we cannot shield out these micro-stochastic influences completely! This machine makes use of the micro-stochasticity of events, but a more delicate and refined use. We can confidently predict that the machine would be harder to make and keep stable, than an ordinary 50/50 -generating machine. There would be a tendency of the frequencies to slide towards 1.0 (if the shielding works too well), or back toward 0.5 (if it lets in too much from outside).

4. The radium atom decay. Nothing much needs to be said here, as current scientific theory says that this is an SNM with no moving parts and no need of shielding. In this respect it is an unusual SNM, and some will wish for an explanation of the reliability of the machine. Whether we can have one or not remains to be seen (though Bohmians think they already have it — and their explanation invokes the SP with respect to particle initial position distributions). Other philosophers will want to try to reduce *all* objective chances to this sort. Whether they can have their way will be the subject of section 6.2

Many Humean objective chances — especially the paradigm cases — will be associable with an SNM whose structure we can lay out more or less clearly. But we should expect that many other Humean chances will not have such a structure. If they exist, out there in the wild, they exist because of the existence of the appropriate sort of pattern in actual events. But the class of ‘appropriate’ patterns will not be rigorously definable. There will be no clear-cut, non-arbitrary line that we can draw, to divide genuine objective probabilities on one side, from mere frequencies on the other.

5. The 09:37 train. Let’s assume that there is no SNM that produces a chance regularity (if there is one) in the arrival time of my morning train. Is there nevertheless an objective chance of the 09:37 train arriving within ± 3 minutes of scheduled time? Perhaps — it depends on what the pattern/distribution of arrival times looks like. Is it nicely random-

looking while overall fitting (say) a nice Gaussian distribution, over many months? Is the distribution stable over time, or if it shifts, is there a nice way to capture how it slowly changes (say, over several years)? If so, it makes perfect sense to speak of the objective chance of the train being on time. On the other hand, suppose that the pattern of arrivals failed to be random-looking in two significant ways: it depends on day of the week (almost always late on Friday, almost always on time on Monday, ...), and (aside from the previous two generalizations), the pattern fails pretty badly to be stable across time, by whatever measure of stability we find appropriate. In this case, it probably does not make sense to say there is an objective chance of the train being on time — even though, taking all the arrivals in world history together, we can of course come up with an overall frequency.

When we discuss the deduction of the Principal Principle, we will see why such mere frequencies do not deserve to be called objective chances. Stability is the crucial notion, even if it is somewhat vague. When there is a pattern of stable frequencies among outcomes in clearly-defined setups (which, again, are often called *reference classes*), then guiding one's expectations by the Humean chances (which will either be close to, or identical with, the frequencies) will be a strategy demonstrably better than the relevant alternatives.

4.3 *The Best System of chances*

I think Lewis was right to suppose that a Humean approach to objective chance should involve the notion of a Best System of chances — though not a Best System of laws + chances together. Now it is time to say a bit more about this idea.

Those who favor the BSA account of laws are welcome to keep it; my approach to chance does not require rejecting it. All I ask, as noted earlier in section 3, is that we allow that, in addition to whatever chances the BSA laws may provide, we can recognize other Humean chances as well, without insisting that they be part of (or follow from) the chancy laws. Then, in addition to a Best System of laws, there will also be a Humean Best System of chances, which I will now characterize.

Lewis was able to offer what appeared to be a fairly clear characterization of his Best Systems with his criteria of strength, simplicity and fit. By contrast, my characterization of chance best systems may appear less tidy from the outset. But there is a good justification for the untidiness. First, the 'best' in Best System means best *for us* (and for creatures relevantly similar). The system covers the sorts of events we can observe and catalogue, and uses the full panoply of natural kind terms that we use in science and in daily life. Pattern regularities

about coins and trains may be found in the Best System, not only regularities about quarks and leptons. Since we are not trying to vindicate fundamentalism or reductionism with our account of chance, but rather make sense of real-world uses of the concept, there is no reason for us to follow Lewis in hypothesizing a privileged physical natural-kinds vocabulary.

In any case, closer inspection of Lewis' theory destroys the initial impression of tidiness. Simplicity and strength are meant to be timeless, objective notions unrelated to our species or our scientific history. But one suspects that if BSA advocates aim to have their account mesh with scientific practice, these notions will have to be rather pragmatically defined. Moreover, simplicity and strength are simply not clearly characterized by Lewis or his followers. We don't know whether initial conditions, giving the state of the world at a time, should count as one proposition or as infinitely many (nor how to weigh the reduction of simplicity, whatever answer we give); we don't know whether deterministic laws are automatically as strong as can be, or whether instead some added chance-laws may increase strength at an acceptable price; if the latter, we don't know how to weigh the increase in strength so purchased. Finally, as Elga (2004) has noted, the notion of 'fit' certainly cannot be the one Lewis proposed, in worlds with infinite numbers of chance events (according to one or more competitor System), since every such system would have the same fit: zero.

I propose to retain the three criteria of simplicity, strength and fit — understanding fit as Elga suggests, as 'typicality' — but now applied to systems of chances alone, not laws + chances. These three notions can be grasped for chance systems alone at least as clearly as they can be grasped for laws + chances systems. For Lewis, *strength* was supposed to measure the amount of the overall Humean mosaic 'captured' by a system. Although vague, this notion of strength appears to be unduly aimed at the capturing of petty details (e.g. the *precise* shape, mass, and constitution of every grain of sand on every beach . . .).¹⁹ When considering chance systems, the capturing of quite particular detail is not necessarily either desirable or achievable, and strength is instead most naturally understood in terms of how many different *types* of phenomena the system covers. Strength should be determined by the net domain of the system's objective probability functions. So if system 1's domain includes everything from system 2's domain, plus more, then 1 beats 2 on strength. Where two systems' domains fail to overlap, it may be difficult to decide which is stronger, since that may require adjudicating which builds strength more: covering apples, or covering oranges

¹⁹ Maudlin has criticized Lewisian approaches to laws for this apparent emphasis on the trivial (personal conversation).

(so to speak). But fortunately, since our systems' chances are not constrained to have the kind of simplicity that scientists and philosophers tend to hope that the true fundamental laws have, this difficulty is easily overcome: a third system that takes (e.g.) system 1 and adds the chances-about-oranges found in system 2, will beat both in terms of strength.

What about simplicity? The value of simplicity is to be understood not in terms of extreme brevity or compact expression, but rather in terms of (a) elegant unification (where possible), and (b) user-friendliness for beings such as ourselves. But elegance is not such an overriding virtue that we should consider it as trumping even a modest increase in strength bought at the expense of increased untidiness. In fact, I tend to see the value of elegant unification as really derivative from the value of user-friendliness. User-friendliness is a combination of two factors: *utility* for epistemically- and ability-limited agents such as ourselves, and *confirmability* (which, *ceteris paribus*, elegant unification tends to boost).

Objective chances are a 'guide to life', and one that ideally we can get our hands on by observation, induction and experimentation. Lewis tried to distance himself from such agent-centered values, in describing his criteria of simplicity and strength, because he wanted his account to mimic, as closely as possible, the physicist's notion of fundamental laws. But for an account of chances alone, there is no need to insist on this kind of Platonic objectivity. There may be such deep laws in our world, and some of them may even be probabilistic. But there are also lots of other chances, dealing with mundane things like coins, cabbages, trains and diseases.

HOC thus offers an empiricist account of objective chance, but one more in the mould of Mach than of Lewis. This may seem like a disadvantage, since the anti-metaphysical positivism of Mach is almost universally rejected, and rightly so. But here I endorse none of Mach's philosophy, not even his view of scientific laws as economical summaries of experience. I remind the reader that this is a *skeptic's* guide to objective chance. If we are convinced that propensity theories and hypothetical frequentist theories of objective probability are inadequate, as we should be, then can objective probabilities be salvaged at all? HOC offers a way to do so, but inevitably its objective chances will appear more agent-centric and less Platonically objective than those postulated by the remaining non-skeptics.

What chances are there in the Best System, how much of the overall mosaic they 'cover' and how well they admit systematization, are all questions that depend on the contingent specifics of the universe's history. And while we have come to know (we think) a lot about that history, there is still much that we have yet to learn. Despite our relative

ignorance, there are some aspects of a Best System of chances for our world that can be described with some confidence.

Earlier we discussed roulette wheels and I mentioned that for any well-made wheel with N slots (within a certain range of natural numbers N), each slot's number has a probability of $1/N$ of winning on each spin. This is an example of the kind of higher-level chance fact that we should expect to be captured by the Best System for our world. It goes well beyond frequentism, since it applies to roulette wheels with few or zero actual trials, and it 'smooths off' the actual frequencies to make them fall into line with the symmetries of the wheels. But still, this chance regularity is just a regularity about roulette wheels. We can speculate as to whether the Best System for our world is able to capture this regularity as an instance of a still higher-level regularity: a regularity about symmetrical devices that amplify small differences in initial conditions and/or external influences to produce (given the SP) a reliable symmetric and random-looking distribution of outcomes over long sequences of trials. Given what we know about the reliability of certain kinds of mechanisms, and the reliability of the stochasticity of the input/boundary conditions for many such mechanisms, this seems like a solid speculation. I would not want, however, to try to articulate a full definition of such SNMs, which have as sub-classes roulette wheels, craps tables, lottery ball drums, and so forth. But we do not have to be able to specify clearly all of the domains of objective chance, in order to have confidence in the existence of some of them.

4.4 More about what the Best System contains

The full domain of chance includes more than just gambling devices, however, even at the macro-level. There may or may not be an objective chance of the 09:37 train being on time, but there certainly is (due to the biological processes of sexual reproduction) an objective chance of a given human couple having a blue-eyed child if they have a baby, and there may well be an objective chance of developing breast cancer (in the course of a year), for adult women of a given ethnicity aged 39 in the United States. I say 'may well', because it is not automatically clear that in specifying the reference class in this way, I have indeed described a proper chance setup that has the requisite stability of distribution, synchronically and diachronically. The problem is well-known: if there is a significant causal factor left out of this description, that varies significantly over time or place, or an irrelevant factor left in, then the required stability may not be found in the actual patterns of events (remember: over all

history).²⁰ If the required stability is present, though, then there is a perfectly good objective chance here, associated with the ‘setup’ described.²¹

It may not be the *only* good objective chance in the neighborhood, however. Perhaps there are different, but equally good and stable statistics for the onset of breast cancer among women aged 39 *who have children and breast-feed them for at least 6 months*. There is a tendency among philosophers to suppose that if this objective chance exists, then it cancels out the first one, rendering it non-objective at best. But this is a mistake. The first probability is perfectly objective, and correct to use in circumstances where one needs to make predictions about breast cancer rates and *either* (a) one does not know about the existence of the second objective probability, or (b) one has no information concerning child-bearing and breast-feeding for the relevant group. There *is* a sense in which the second probability can ‘dominate’ over the first, however, if neither (a) nor (b) is the case. PP, with admissibility correctly understood, shows us this. Suppose we are concerned to set our credence in

A: Mrs. K, a randomly selected woman from the New Brunswick area aged 39, will develop breast cancer within a year.

and we know these probabilities:

X1: $Pr(B. cancer | woman 39, \dots) = x_1$

X2: $Pr(B. cancer | woman 39 \& has breast-fed, \dots) = x_2$, and

X3: $Pr(B. cancer | woman 39 \& does not breast-feed, \dots) = x_3$

and for the population, we have all the facts about which women have had children and breast-fed them. With all of this packed into our evidence *E*, we *cannot* use PP in this way:

$Cr(A | X1 \& E) = x_1$

Why not? Since our evidence *E* contains *X2*, *X3*, and the facts about which women have breast-fed children (including Mrs. K), our evidence contains information relevant to the truth of *A*, which is not information that comes by way of *A*’s objective chance *in the X1*

²⁰ As I argue in “Humean Effective Strategies”, the requisite stability may also fail just due to statistical “bad luck” - one should not think that presence/absence of causal factors will explain everything about the actual statistical patterns.

²¹ One way for frequencies such as these to have the requisite stability, of course, is for there to be a well-defined but time-varying chance function. Given the trend of increasing cancer rates in Western countries over the 20th century, this is probably the only way that an objective chance of the kind under discussion can exist. I will generally ignore the possibility of time- and/or space-variable chance functions, but it is good to keep in mind that the Best System will probably include many such chances.

setup (the one whose invocation we are considering). So this information is all jointly inadmissible. By contrast, we *can* apply PP using X_2 , because all our evidence is admissible with respect to that more refined chance. Knowing X_2 , we also know that X_1 is not relevant for the truth of A for cases where we know whether a woman has breast-fed or not.²²

What I have done here is reminiscent of the advice that empiricist/frequentists have sometimes given, to set the (relevant) objective probability equal to the frequencies *in the smallest homogeneous reference class for which there are 'good' statistics*. But we have, hopefully, a clearer understanding of what 'homogeneous' means here, in terms of the chance setups that make it into the best system's domain; we see that we may be able to apply the objective chance even if the relevant events form a reference class too small to have good statistics - namely, if the Humean chance is underwritten by a higher-level pattern that gives coverage to the setup we are considering; and finally, we see why this advice does not automatically undercut the claim to objectivity of probabilities for larger, less-homogeneous reference classes.²³

When considering the chances of unpleasant outcomes like cancer, of course, we would typically like something even better than one of these objective probabilities for setups with many people in their domain. We would like to know *our own, personal* chance of a certain type of cancer, *starting now*. The problem is that we can't have what we want. Since a given person's history does not suffice to ground a chance-making pattern for cancer, for such a chance to exist it would have to be grounded at a different level, perhaps by reduction to micro-level probabilities. But even if this reductionist objective chance exists as a consequence of the best system — and I think one can reasonably doubt this — we are never going to be able to know its value (not being omniscient Laplacean demons). So for the

²² How do we know this, the irrelevance of the X_1 probability once X_2 becomes available? It follows directly from the fact that the X_2 reference class is a subset of the X_1 class, given the justification of PP (see sec. 5). To anticipate: while the argument justifies setting credences concerning a (medium-large) number of instances of the X_2 setup using the X_2 chance, it would fail if presented as a justification for using the X_1 chance to set credences for a similar number of instances of the X_2 setup.

²³ It might seem that the Best System aspect of our Humean account of chance should rule out such overlapping chances: the system should choose the *best* of the two competing chances. In this case, that would perhaps mean jettisoning the chance X_1 and retaining the X_2 chances. In some cases that may be correct, but not in general. Considerations of discoverability and utility (applicability) will often be enough to demand the retention of more-general chances alongside more-specific ones, in the Best System.

purposes of science, of social policy, and of personal planning, such individualized objective chances may as well not exist; and a philosophical account of chance that hopes to be relevant to the uses of probability in these areas of human life needs to look elsewhere.

Does this mean that HOC denies the existence of single-case objective chances? Not exactly; rather, single-case chances exist wherever a situation's description fits into the domain of the Best System's chance functions. Every time *you*, a competent adult gambler, flip a fair coin, the objective chance of *heads* exists, and is $1/2$.²⁴ But, of course, nothing distinguishes that flip from your next one, for which the chance of heads will be the same. For some philosophers, this means that we are not really talking about genuine single-case probabilities here after all. For them, the specifics of your physical situation just prior to the flip may be quite relevant, and entail that your next flip has a very different $Pr(\textit{heads}|\textit{flip})$ than the one after that. But such philosophers are thinking of chances as (meta)physical propensities, and making themselves hostage to the fortunes of determinism in physics. If Bohm turns out to have been right about quantum mechanics, or if Diaconis is right in modeling coin flips as a Newtonian process, then on both of your next two flips, $Pr(\textit{heads}|\textit{flip})$ is either zero or one. This is an awkward consequence, since it entails that any probability for *heads* near $1/2$ can only be a subjective probability, not an objective probability.

The Humean about chance, without having to reject or endorse determinism, sets aside such propensities (if they exist) and defines objective chances differently, along the lines we've sketched in this section. They are intrinsically generic: whenever the setup conditions are instantiated, the objective chance is as the Best System specifies. But they are fully applicable to single cases, as long as HOC can rationalize the Principal Principle — the question to which we turn in section 5.

4.5 Relations to other accounts of objective probability

If an account of objective probability is going to square at all well with the way we understand paradigmatic cases and uses of the concept, inevitably it will have resemblances to the most widely discussed and defended earlier accounts. Here I will briefly lay out how I see some of these relations to earlier accounts, and how HOC remedies their defects.

²⁴ Or close to it. Diaconis, Holmes and Montgomery (forthcoming) have done empirical work suggesting that the chance of *heads* depends on whether the coin starts heads-up or not. But their figures take the chance of *heads* away from 0.5 by at most 0.01.

Classical probability. The early definition of probability, based on the principle of indifference and an *a priori* assignment of equal chance to the outcomes among which a rational person is ‘indifferent’ (due to some symmetry of the setup), works well when it comes to classical gambling games and devices. Its close cousins, the *logical probabilities* of Carnap and Keynes, continue to exert a fascination on the minds of some philosophers. HOC captures much of the attractive elements of classical probability, due to a few of its features. First, like classical probability, HOC accepts that the domain of objective chance may be restricted: the world may not be such as to constrain rational degrees of belief for *all* propositions. Second, via the SNMs associated with classical gambling devices and with the help of the SP, HOC shows why the symmetries that provoke our indifference are just the symmetries that explain objectively equiprobable outcomes. The Best System element of HOC lets us capture the classical intuition that on an N-slotted roulette wheel, the chance of 00 is $1/N$ — whatever the actual frequency may be! And finally, by rationalizing PP, HOC will capture the intuition that objective probabilities are the degrees of belief an ideal rational agent *should* have (when lacking further/better information). These successes come without the price of vulnerability to the Achilles’ heel of classical probability, the incoherence of applying the principle of indifference to systems with more than one symmetry. The Best System automatically singles out one symmetry — the ‘right’ one, given the actual Humean mosaic in our world — as the symmetry to associate with equiprobability of outcomes.

Propensities. In one way, HOC has nothing at all in common with propensity accounts of chance: as a skeptical, Humean view, it is dead-set against the invocation of precisely what the propensity theorist offers. But at the same time, under the assumption of determinism (or an effectively deterministic mechanics, for a given type of process — such as Diaconis offers us for coin flips), HOC again captures what was attractive in the propensity view: the idea that certain systems are correctly described as having a *tendency* to produce outcomes ‘at random’, but with certain stable long-run frequencies. Again, SNMs and the SP are the key here. Assuming the correctness of SP for (say) coin-flip initial/boundary conditions, it is perfectly correct to say: a coin-flipping mechanism has a tendency or propensity to produce the outcome *heads* half of the time, in a long enough sequence of flips. This behavioral disposition is not grounded in some new, mysterious property of the flipper, but rather in what follows *logically* from the SP and the mechanism’s structure (given the causal or natural laws governing the SNM itself). But unless universal determinism is assumed, HOC cannot lay claim to capturing all the propensity-chances that

their advocates typically postulate.

Frequentism. When it comes to *hypothetical* frequentism, again we see that HOC can capture what seemed most correct about the earlier view, in its paradigm applications, using the resources of SNMs and the SP. When you think about it, what is really meant in saying something like: ‘If we continued flipping this coin forever, generating an infinite sequence of outcomes, the limiting frequency of *heads* would be $\frac{1}{2}$, and no place-selection function would be able to select a subsequence with limiting frequency of heads different from $\frac{1}{2}$.’? A real coin would disintegrate to its component atoms in a finite time, and to keep it going with repairs, we would need an infinite supply of metals; and a universe not subject to heat death or a Big Crunch; Taking its counterfactuals literally, hypothetical frequentism is not an especially plausible or attractive view. Nor, I suggest, did its proponents really take their counterfactuals literally. But if we replace ‘infinite’ with ‘really long’, then SP makes very plausible that, for coin-flipping SNMs, ‘If we continued flipping the coin a really long time, the frequency of *heads* would be very near $\frac{1}{2}$, and the sequence would pass all typical statistical tests for randomness.’

Actual frequentism. In a clear sense, HOC is closer to actual frequentism than to any other earlier account of objective probabilities. In section 5 we will see how this is the key to its ability to justify PP. The ways in which HOC may be seen as a ‘sophistication’ of actual frequentism are precisely the features that leave it invulnerable to the most important standard objections to actual frequentism. Frequentism is often taxed with the ‘problem of the single case’. The probability of *heads* for this coin is supposed to be the actual frequency of *heads*; but this coin is flipped only once, then destroyed, so the probability is either 1 or 0. HOC overcomes *this* form of the objection easily: there is no setup for *this* coin alone, rather the more generic setup of ‘flips of a fair coin’, as described in sec. 4.2 — and in actual world history, they are numerous indeed, and the frequency of *heads* is very close to $\frac{1}{2}$. Fine; but the example can be reformulated with some gambling device that genuinely is only used once in world history (e.g. a 234-slotted roulette wheel). Still no problem: in the Best System, we have every reason to think, the chance of 00 on a 234-slot roulette wheel will be dictated by the symmetries of the device and the patterns in the Humean mosaic at many levels — from micro-initial-condition distributions on up (at least) to roulette-wheel-outcome patterns. HOC tells us that the chance is $\frac{1}{234}$, and not 0 or 1.

Similarly, frequentism is plagued by the ‘reference class problem’, but the Best

System feature of HOC mitigates the problem as much as it can be mitigated.²⁵ We saw an example of this in the examples of breast-cancer probabilities above. Considerations of stability and random-lookingness of outcome frequencies eliminate many potential reference classes (or chance setups), while others are eliminated by low frequencies and lack of support from patterns in the mosaic at higher or lower levels. Still, a given situation may instantiate more than one chance setup, and thus have more than one set of objective chances applicable to it. This is not a defect of HOC, because our understanding of admissibility and PP shows us how and when one of the chances dominates over the other, if both are known.

Finally, philosophers attracted to propensity views sometimes object to frequentism on the ground that it implausibly makes the probability of an outcome *here, now* dependent not just on the physical facts about the setup system here and now, but also on a myriad of outcomes elsewhere in space and time. This ‘non-locality’ of objective chance is not mitigated at all in HOC, but must simply be accepted.

Summing up: Chances are constituted by the existence of patterns in the mosaic of events in the world. These patterns are such as to make the adoption of credences identical to the chances rational in the absence of better information, if one is obliged to make guesses or bets concerning the outcomes of chance setups (as I will show in section 5). Some stable, macroscopic chances that supervene on the overall pattern are explicable as regularities guaranteed by the structure of the assumed chance setup (the SNM), together with our world’s micro-stochastic-lookingness (SP). These are as close as one can get to the propensity theorist’s single-case chances, within a Humean view. Not all genuine objective chances have to be derivable from the structure of the SNM and the correctness of the SP at the level of initial/boundary conditions, however. The right sort of stability and randomness of outcome-distribution, over all history, for a well-defined chance setup, is enough to constitute an objective chance. Moreover, setups with few actual outcomes, but the right sort of similarities to other setups having clear objective chances (e.g. symmetries, similar dynamics, etc.) can be ascribed objective chances also: these chances supervene on aspects of the Humean mosaic at a higher level of abstraction. The full set of objective chances in our world is thus a sort of Best System of many kinds of chances, at various levels of scale and with varying kinds of support in the Humean base. What unifies all the chances is their aptness to

²⁵ The reference class problem is not just a problem for frequentism and HOC; see Hájek (2006), “The Reference Class Problem is Your Problem Too”.

play the role of guiding credence, as codified in the Principal Principle.

5. Deduction of PP

In this section I will show how, if objective chances are as the HOC view specifies, the rationality of PP follows. Lewis claimed: ‘I think I see, dimly but well enough, how knowledge of frequencies and symmetries and best systems could constrain rational credence.’ (1994, p. 484) Recently, Michael Strevens and Ned Hall have claimed that Lewis was deluding himself, as either there is *no way at all* to justify PP, on any view of chance (Strevens 1999), or no way for a Humean to do the job (Hall 2004). I will try to prove these authors mistaken by direct example, offering a few comments along the way on how they went astray.

5.1 Deducing the reasonableness of PP

The key to demonstrating the validity of PP for Humean chances rests on the fact that the account is a ‘sophistication’ of actual frequentism. For the purposes of this section, we can think of HOC as modifying simple actual frequentism by:

- A. Requiring that outcomes not only have an actual frequency (or limit, if infinity is contemplated, given a reasonable time-ordering of setup instantiations), but also that the distribution of outcomes ‘look chancy’ in the appropriate way — stability of distribution over time and space, no great deviations from the distribution in medium-sized, naturally selected subsets of events, etc. This is something a smart frequentist would insist on, in any case.²⁶
- B. Allowing higher-level and lower-level regularities (patterns), symmetries, etc. to ‘extend’ objective chance to cover setups with few, or even zero, actual instances in the world’s history (often with the help of the SP and the notion of an SNM)
- C. Anchoring the notion of chance to our epistemic needs and capabilities through the Best Systems aspect of the account.
- D. Insisting that the proper domain of application of objective chances is intrinsically limited, as we will see in section 6.²⁷

²⁶ Von Mises (1928) insisted on something of this nature, but in order to make it mathematically tractable in the way he desired, he had to make the unfortunate leap to hypothetical infinite “collectives”.

²⁷ The limitation most familiar to readers will be the limitation imposed by the undermining phenomenon, which is this: objective chances cannot be used to guide credence

With these ideas in mind, we can sketch the basic argument that establishes the reasonableness of PP for Humean chance along what Strevens (1999) calls ‘consequentialist’ lines.

Let’s recall PP: $Cr(A|XE) = x$ In words: given that you believe (fully) that the objective chance of A being the case is x , and have no further information that is inadmissible with respect to A , then your subjective degree of belief in A should be x .²⁸ What we need to demonstrate is that there is an objective sense in which, if your belief about the objective chance is correct, then the recommended level of credence x is better than any other that you might adopt instead. Here our assumption is that we are dealing with a simple, time-constant objective chance of A in the setup S . Not all objective chances in the Best System need be like this, of course, but the argument will carry over to less-simple chance laws and distributions more or less directly and obviously.

Let’s suppose that this is a typical objectively chancy phenomenon, in the sense that S occurs very many times throughout history, and A as well. Then our Humean account of chance entails that the frequency of A in S will be approximately equal to x , and also that the distribution of A -outcomes throughout all cases of S will be fairly uniform over time (stable), and stochastic-looking. If the first (frequency) condition did not hold, x could not be the objective chance coughed up by the Best System for our world — world history would undermine the value x . If the stability condition did not hold, then either our Best System would not say that there exists an objective chance of A in S , or it would say that the chance is a certain function of time and/or place - contrary to our assumption. If the stochastic-lookingness condition did not hold, then again either the Best System would not include a chance for A in S , or it would include a more complicated, possibly time-variable or non-Markovian chance law — contrary to our assumption.²⁹

unrestrictedly, over events of arbitrary “size” within the overall mosaic; their proper use is restricted to relatively small events or sets of events, small compared to the global patterns. See Hoefer (1997) and section 6.1 below.

²⁸ More typically one does not have full belief in one particular value of the objective chance, instead spreading one’s degrees of belief over a range of possible chance-values. PP generalized to accommodate this fact reads: GPP: $Cr(A|E) = \sum_i Cr(X_i|E) \cdot x_i$, where the propositions X_i specify that the objective chance of A is x_i . The justification of this GPP depends on the prior justification of PP in a fairly obvious way, but I will not explore this issue here.

²⁹ For example, it might be that after any two consecutive flips turn up heads, the frequency of heads on the next flip of the same coin is 0.25, on average, throughout history. If this were

Therefore, we know the following: at *most* places and times in world history, if one has to guess at the next n outcomes of the S setup, then if n is reasonably large (but still ‘short run’ compared to the entire set of actual S instances), the proportion of A outcomes in those n ‘trials’ will be close to x . And sometimes the proportion will be greater than x , sometimes less than x ; there will be no discernible pattern to the distribution of greater-than- x results (over n trials) vs less-than- x results; and if this guessing is repeated at many different times, the average error (sum of the deviations of the proportions of A in each set of n trials from x) will almost always be close to zero. Notice that we have made use of the ordinary language quantifiers ‘most’ and ‘almost always’ here, and we have *not* said anything about getting things right ‘with high probability’ or ‘with certainty, in the limit’, or anything of that nature.

What we have seen so far shows that, if one has a practical need to guess the proportion of A -outcomes in a nontrivial number of instances of S , then guessing x is not a bad move.³⁰ Is there a better move out there, perhaps? Well, yes: it is even better to set one’s credences equal to the actual frequency of A in *each specific set of trials* - that way, one never goes wrong at all. This is basically the same as having a crystal ball, and we already knew that guiding one’s credences by a reliable crystal ball is better than any strategy based on probabilities. We can associate this idea with a competitor ‘theory’ of objective chance, which we will call ‘Next- n Frequentism’ (NNF). NNF-chances are clearly very apt indeed

so, then (depending on how other aspects of the pattern look) we might have a very different, non-Markovian chance law for coin flips; or a special sub-law just for cases where two flips have come up heads, and so forth.

³⁰ Those who like proofs are referred to Appendix B of Strevens (1999), in which he proves that following PP is a “winning strategy” in the [medium-] short run if the relative frequency f of A is close to the objective chance x , where “close” means that f is closer to x than are “the odds offered by Mother Nature”. Strevens has in mind a betting game in which Mother Nature proposes the stakes, and the chance-user gets to decide which side to bet (A , or $\neg A$).

Since Strevens maintains that no justification of PP is possible at all, under any theory, what does he make of this short-run argument for PP in his appendix B? He views it as a case of “close, but no cigar” because he thinks that there can be no *guarantee* that the antecedent conditions are satisfied, in a chancy world (pp. 258 - 259) (Like Hall, and most other authors, he thinks of objective chances as being, like propensity theorists maintain, compatible with any results whatsoever over a finite set of cases). Strevens does consider the possibility that satisfying these conditions might be built in to the definition of chance directly (as, in effect, is done in Humean chance), but rejects this on the grounds that we then make it much harder to establish that the objective chances exist at all. Strevens’ argument here is directed against a limiting-frequency account rather than a Best Systems account, so his worry does not in fact apply to the account on offer here. In effect, by taking objective chances too metaphysically seriously (i.e., like propensity theorists or hypothetical frequentists do), Strevens ignores the logical subspace of theories of chance occupied by our Humean best system account.

for playing the PP role. To the extent that they diverge from the Humean chances in value, someone who needs to predict the next n outcomes would be well advised to set their credences equal to the NNF chances.

Unfortunately, unlike Humean chances (which can be discovered inductively, with no more problem than what Hume's problem of induction poses for all knowledge), you really would need a crystal ball to come to know the NNF-chances. Since no reliable crystal balls seem to exist, and there is no other way to arrive at this superior knowledge, we can set aside NNF-chances as irrelevant. The question for our justification of PP is rather: is there any other fixed proportion x' that would be an even better guess than x ? And as long as the difference between x' and x is non-trivial, the answer to this question is going to be, clearly, 'No'. Any such x' will be such that it diverges from the actual proportion of A in a set of n trials in a certain direction (+ or -) more often, overall, than x does. And its average absolute error, over a decent number of sets of n trials, will *almost always* be greater than the average absolute error for guessing x .

'Almost always' is not 'always'. It is true that a person might set their credences to x' , guess accordingly, and do better in a bunch of sets of n trials than an x -guesser does. PP can't be expected to eliminate the possibility of bad luck. But the fact that *most of the time* the x -guesser accumulates less error than the x' -guesser is already enough — enough to show the reasonableness of PP. As we will see below, this is far more than any competitor theory of chance can establish.

There are further comments to be made, and loopholes to close off, but the basic argument is now out in the open. PP tells you to guess (set your credences) equal to the objective chance; in general, the Humean objective chance is close to the actual frequency; and setting one's credences equal to the actual frequency is better, at most times and places, than setting them equal to any other (significantly different) value. This basic argument raises many questions; I will address two important ones briefly now.

From n-case to single-case reasonability?

The basic argument looks at a situation in which we need to guess the outcomes of S setups over a decent-sized number of 'trials'. But sometimes, of course, we only need to make an educated guess as to the outcome of a single 'trial' — e.g., if you and I decide to bet on the next roll of the dice, and then quit. Does our argument show that application of PP in such

‘single-case’ circumstances is justified? Assuming, as always, that we have no inadmissible information regarding this upcoming single case, the answer is ‘yes, it does.’ Our argument shows that setting our level of credence in A to the objective chance is a ‘winning’ strategy most of the time when guessing many outcomes is undertaken. But setting our credence for outcome A equal to a *constant*, x , over a series of n trials is *the same thing as* setting our credence equal to x for each individual trial in the collection. It cannot be the case that the former is reasonable and justified, but the latter — its identical sub-components — are unreasonable or unjustified. Just as two wrongs do not make a right, concatenating unreasonable acts does not somehow make a composite act that *is* reasonable. So when we have to set our credence - make a bet, in other words - on a single case, the prescription of PP remains valid and justified, even though that justification is, so to speak, inherited from the fact that the PP-recommended credences are guaranteed to be winners most of the time, over medium-run sets of ‘trials’.

Suppose the contrary. That is, suppose that we think it is possible that over decent-sized numbers of trials using PP is justified, but nevertheless in certain *specific* single cases — say, *your very next* coin flip — it may in fact not be justified. What could make this be the case? Well, you might think that some specific local factors exist that make a higher-than-0.5 credence in Heads reasonable. Perhaps you have put the coin heads-up on your thumb, and as it happens, your coin flips tend to turn over an even number of times, so that flips starting heads-up (tails-up) land heads (tails) more often than 50% of the time. What this amounts to, of course, is postulating that *your* coin flips in fact comprise a different SNM than ordinary coin flipping devices (including persons) — condition 5 in our earlier description of the fair coin flipping SNM is violated by these assumptions. Nevertheless, let’s assume for the moment that the 50/50 chance of heads and tails on coin flippings is in fact a proper part of the Best System, and that your flips do fall under the 50/50 regularity (perhaps the Best System is content with a looser specification of the coin-flip setup). By our assumptions, your flips *also* satisfy the conditions for a less common SNM. Now there are two cases to consider: first, that we (who have to set our credences in *your next flip*) know about this further objective chance; second, that we don’t.

If we do know about it, then the case is parallel to the breast cancer example above. Knowing the higher-than-50% chance for your flips when they start heads-up, and seeing that indeed you are about to flip the coin starting from that position, we should set our credences to the higher level. The rules of admissibility tell us that the ordinary 50/50 chances are

inapplicable here, but the chances arising from the more constrained SNM that your flips instantiate are OK. This does nothing, of course, to undercut the reasonableness of PP in *ordinary* applications to a single case, when such trumping information is not available. Chance is a guide to life *when you can't find a better one*, and has never aspired to anything higher!

If we do not know about it, then the question becomes: is setting our credence to 0.5 as PP recommends reasonable? It might seem that the answer is negative, because the 'real' probability of heads on that next flip is in fact higher. But as Humeans about chance, we know that this way of thinking is a mistake. Both objective chances are 'real'; the fact that sometimes one trumps the other does not make it more real.³¹ An application of PP is not made unreasonable just because there is a *better* way of guiding credence out there; it only becomes unreasonable if that better way comes to the chance user's attention, and when that happens, the original chance can no longer be plugged into PP because of the violation of admissibility.

Few- and no-case chances

The Best System aspect of Humean chance takes us away from a (sophisticated) actual frequentism in two ways: by 'smoothing out and rounding off' the chances and chance laws, and by using higher-level and lower-level regularities to extend the domain of objective chances to cover setups that have few, or no, instances in actual history (like our 43-slot roulette wheel). The former aspect makes objective chances easier to discover and to work with than pure finite frequencies, which is all to the good in a concept, such as the concept of objective chance, whose nature is bound up with its utility to finite rational agents. Since it never drags the objective chance far away from the frequencies, when the numbers of actual instances of the setup in history are large, it is not problematic *vis a vis* the deduction of PP. But the latter aspect may well look problematic. For we know that if the number of actual instances of setup S in all history is relatively small, then the actual frequency is may well be quite far from the objective chance dictated by the higher-level pattern. Perhaps 00 lands in one twenty-fifth of the times (say, only a few hundred) that our 43-slot roulette wheel is spun, rather than something close to one forty-third. In cases such as this, how can PP be justifiable?

³¹ And as we will see in the next section, it is possible for a more generic, macro-chance to trump a more specific, micro-based chance.

Notice that in cases like this, where the numbers over all history are relatively low, we cannot mount an argument similar to our basic argument for PP above, but now in favor of a rule of setting credence equal to the actual frequency. Why not? Because there is no guarantee that there will be the sort of uniformity over subsets of n trials that we knew we could appeal to when the numbers are large or infinite. Suppose our 43-slot roulette wheel was spun a total of 800 times, and 00 came up 32 times. Considering now guesses as to the number of 00 outcomes in ‘short run’ sets of $n = 50$ consecutive spins, can we assert with confidence that in most of these runs, a guess of $2/50$ will be closer to the actual frequency more often than a guess of $1/43$? By no means. It may happen that the actual pattern of outcomes is very much as one would expect based on the objective chance of $1/43$, but in the last 90 spins 00 turns up a surprising 12 times. These things happen, and given the low overall numbers, cannot be considered undermining outcomes for the objective chance of $1/43$.³² But if this is how the 00 outcomes are distributed, then a person betting on a subset of 50 consecutive spins will do better to bet with the objective chance, *most of the time* (just not in the last 100 spins!). Neither a credence level of $1/25$ nor one of $1/43$ is guaranteed to be a winner in most of the short run guessing games that might be extracted from the total 800 spins.

It is true, of course, that if we consider all the possible ways of selecting sets of 50 outcomes for our guessing games — not just consecutive spins, but even-numbered spins, ‘randomly’ chosen sets, etc. — then the frequency of $1/25$ is going to do better, overall, in this wider set of games. (The reason is that in the vast range of games where the 50-trial subsets are chosen ‘randomly’, the frequency in such subsets will be close to $1/25$ more often than to $1/43$.) But it remains true that a consequentialist argument for setting credence equal to the actual frequency rather than the HOC is much weaker here than is the argument for using HOCs in the standard case of setups with very many outcomes. And we should remember that if the divergence is too serious, and the number of outcomes at issue large enough, the Best System will have to go with the frequency rather than the symmetry-derived chance: we have frequency-deviation tolerance, but only within limits!

The credence level of $1/25$ is guaranteed to beat that of $1/43$, of course, in the limit as the ‘short’ run over which we are guessing approaches, and finally equals, the total set of 800 actual spins. To this minor extent, and for these sorts of cases, actual frequentism can claim to offer a stronger justification of PP than HOC offers. Still, it seems to me that the

³² Undermining is explained in section 6.

importance of this lacuna should not be overstated, in light of the fact that when objective chance and actual frequency *do* diverge significantly, (a) neither one is *guaranteed* to beat the other as a credence-setting strategy for short runs, and (b) any disadvantage that accrues to setting one's credences via PP rather than according to the actual frequencies will be limited in size.

5.2 *Other accounts and PP*

Whatever the difficulties or limitations of our deduction of PP may be, it should be apparent already that the standard competing accounts of objective chance (with the exception of actual frequentism) cannot offer anything even remotely similar. Actual frequentism can of course mount a consequentialist justification of PP along the lines we took in section 5.1, as long as the frequentist's position incorporates many of the 'sophistications' found in HOC: restricting the domain so that not every frequency in every conceivable reference class counts as an objective probability; insisting on a stochastic-looking distribution of actual outcomes; and so forth. Such sophistications have tended to be built into scientific and statistical practice, if not built into the official definition of probability, which is another way of saying that HOC nicely fits scientific and statistical uses of objective probability.

The traditional *hypothetical* frequentist, however, adopts a position that utterly disconnects the objective chances from frequencies (and even from random-lookingness) in finite initial subsequences, so she cannot mount an argument for the rationality of PP based on it's guaranteeing a winning strategy most of the time (for us, in finite human history). Instead she will be forced to go second-order and say that adopting PP will yield winning strategies *with high probability*, this probability being an objective one derived from the first-order objective chances themselves. But even if we grant her this second-order objective probability, the question remains: why should I *believe* that adapting my credences to the objective chances is a likely-winning strategy, just because this proposition has a high objective chance in the hypothetical frequentist's sense? Evidently, I am being asked to apply PP to her second-order chances, in order to establish that PP is justified for her first-order chances . . . a glaring circularity. Moreover, if I reflect on the literal meaning of these second order chances, they direct me to contemplate the limiting frequency of cases (worlds?) in which applying PP to the first-order chances is a winning strategy, in a hypothetical infinite sequence of 'trial-worlds'. The metaphysics begins to look excessive, and in any case

we immediately see that the problem reappears at the second level, so that we need a third-order argument to justify PP for the second-order chances . . . And so on. Infinity turns out to be an unhappy place to mount a consequentialist argument for PP.

What about a non-consequentialist justification? Howson and Urbach (1993) offer an ingenious non-consequentialist defense of PP in the context of von Mises' version of hypothetical frequentism that is successful, *pace* worries about the coherence of counterfactuals whose antecedents posit an infinite sequence of coin flips. What the argument shows is that *if* you consider objective chances to be the frequencies in hypothetical von Mises collectives generated from the chance setup, and believe the chance of outcome *A* is *p*, then it is incoherent to set one's degree of belief concerning the next outcome's being *A* to any value other than *p*. This justification of PP is in one sense stronger than mine for HOC, since it shows setting one's credences differently to be *incoherent*, not just likely to bring about unfortunate results. In another sense, the justification is weaker, and for precisely the same reason: It shows that violating PP is incoherent, but not that it is a strategy likely to bring about unfortunate results in our world.³³

The metaphysical propensity theorist does not offer a definition of objective chances (in a reductive sense) at all, yet may still claim that PP is valid. The boldest version of this position in recent times appears in Hall (2004).

Let's recall the full force of Lewis' challenge to the advocate of metaphysical propensities:

'Be my guest — posit all the primitive unHumean whatnots you like. (I only ask that your alleged truths should supervene on being). But play fair in naming your whatnots. Don't call any alleged feature of reality 'chance' unless you've already shown that you have something, knowledge of which could constrain rational credence.' (1994, p. 484).

Answering Lewis' challenge to the advocate of metaphysical whatnots, Hall writes:

³³ Strevens (1999) criticizes the Howson and Urbach argument for relying on a version of the principle of indifference, which Strevens regards in turn as a restricted version of PP. But I believe he mistakes the nature of the principle presupposed in Howson and Urbach's argument. The Howson and Urbach argument is also amenable to application to HOC; for reasons of space I will not go into these issues further here.

‘What I ‘have’ are *objective chances*, which, I assert, are simply an extra ingredient of metaphysical reality, not to be ‘reduced’ to anything else. What’s more, it is just a basic conceptual truth about chance that it is the sort of thing, knowledge of which constrains rational credence. Of course, it is a further question — one that certainly cannot be answered definitively from the armchair — whether nature provides anything that answers to our concept of chance. But we have long since learned to live with such modest sceptical worries. And at any rate, they are irrelevant to the question at issue, which is whether my primitivist position can ‘rationalize’ the Principal Principle. The answer, manifestly, is that it can: for, again, it is part of my position that this principle states an analytic truth about chance and rational credence.’ (2004, p. 106).

The problem with this move is that assertion is not argument. If one postulates a metaphysical whatnot (calling it an ‘ingredient’ doesn’t help), insists that it is irreducible to anything else *and* indeed that the actual history of occurrent events places no constraints on what these whatnots might be (numerically), then it is bold indeed to assert that these unknowable whatnots ought to constrain my degrees of belief, else I am irrational. (Notice also that Hall is deliberately declining to ‘play fair’ in Lewis’ terms, calling his ingredients ‘chances’ without *showing* anything at all about them.)

It seems to me that a position such as this writes itself completely off the map as far as real-world endeavors are concerned. Until the primitivist can overcome the ‘modest’ skeptical worry that, perhaps, there are no objective chances in our world after all — and for Hall, as for Mellor (1995) and others, this means *proving determinism false* — scientists and statisticians can safely ignore him. They may as well help themselves to my Humean objective chances, whose existence is guaranteed and which can be learned by ordinary scientific practices.³⁴

³⁴ Hall goes on to suggest that perhaps the primitivist will be able to argue that the reasonableness of PP follows from constraints on reasonable initial credence functions that are connected to “categorical” features of the world. But although Hall speaks of these constraints as if they were ‘imposed by’ categorical facts (p. 107), in fact the constraints he has in mind are just rules about how to distribute credence in light of the presence or absence of certain categorical features of things. (E.g., Hall suggests the constraints might include ‘. . . various indifference principles, carefully qualified so as to avoid inconsistency.’) He goes on to postulate a situation in which exchangeability works as a ‘categorical constraint’, allowing a hypothetical frequentist to derive the appropriate local application of PP. But

Finally, let's consider Sober's (2004) no-theory theory of chance. It has the great virtue, compared to metaphysical propensities and hypothetical frequentism, of letting us be sure that the chances exist and are nontrivial, because the chances just *are* whatever our accepted scientific theories say they are. It also beats the other two competitors when it comes to justifying PP. Whereas they can give, in the end, no justification of the reasonableness of PP at all, Sober can perhaps appeal to our successes in using objective probabilities as inductive evidence that applying PP is a good strategy.

Or rather, Sober can help himself to such inductive support in sciences where the chances are not inferred from statistics, but rather grounded aprioristically in theory (i.e., in QM and statistical mechanics). It is widely recognized that the standard procedures of inferring objective probabilities from statistical evidence *use* PP — that is, their claim to methodological soundness rests on assuming that the objective chances being guessed at deserve to govern credence. Where PP needs to be assumed in arriving at the (estimated) chances, it would seem to be circular to claim that the success of the sciences using said chances argues for the validity of PP as applied to chances in that science. This caveat drastically reduces the range of objective probabilities for which the no-theory theory can claim an inductive grounding for PP.

So the no-theory theory is a clear improvement over the more traditional theories, but only enjoys this advantage in sciences where chances are derived from pure theory.

6. The Limitations of Chance

6.1 Undermining

Undermining is a feature of Humean chances: the objective probabilities supervening on the Humean mosaic may entail that certain sorts of 'unlikely' large-scale future events have a non-zero objective chance, even though *were said events to have occurred, then the Humean mosaic would have been so different that the objective chances themselves would have been different*. The problem gets elevated to the status of a contradiction as follows: If F is our undermining future, then PP says

$$C(F|H_{tw}T_w) = x = Pr(F), \quad x > 0, \text{ where } Pr \text{ is the objective chance function given by } T_w.$$

exchangeability is by no means imposed on us by categorical facts in the world, and Hall's application of it here is in reality just a disguised application of PP itself.

But (the story goes) F plus H_{tw} jointly entail, under the Humean analysis of chance, $\neg T_w$. So the standard axioms of probability tell you that $C(F|H_{tw}T_w) = 0$. Never mind that the difference between x and zero will be just about zero, in any realistic case; never mind that neither H_{tw} nor T_w can ever really be known by us, and so forth: a contradiction is a contradiction, and is bad news.

Many papers on Humean chance since 1994 have been largely devoted to the undermining problem: arguing that it is solvable, or unsolvable, or was never a problem to begin with. There are too many turns and twists to the story to enter into it here. My view is that the contradiction problem is real, though much harder to generate than Lewis originally thought; no application of HOC that beings like us would ever undertake can have undermining potential. And as I argued in (1997), the solution to the problem is basically to ignore it.

Recall that the deduction of PP relies on the scenario of using chances to predict outcomes over small-to-medium sized sets of events. Since PP is essential to chance, the limitations on its validity are likewise limitations on the proper scope of applicability of HOC. It is a mistake — a misunderstanding of the nature of objective chance — to contemplate using the OC's for setting credences concerning fragments of the Humean mosaic big enough to have undermining potential. Humean chance is by its very nature a notion whose range of correct application must be limited. The applications of PP that might engender a contradiction due to undermining are thus not proper applications of PP to HOC at all.

6.2 *A Chance for all Setups?*

HOC may be limited in another way as well. Let's suppose for the moment that the Best System gives chances for micro-level events — *complete* chances, in the sense that for a completely-specified initial physical state of affairs, the Best System entails the chances for all physically possible future evolutions of the physical states of affairs.³⁵ Then if we suppose also a strong enough reductive supervenience thesis (all possible macroscopic states of affairs being equivalent to disjunctions of micro-physically specified states of affairs), then a question arises concerning the Best System's chances. The System may have, we said, an objective chance of the 09:37 train arriving late on an ordinary weekday. And this chance

³⁵ Dupré (1993) calls this 'causal completeness' in the probabilistic sense. The non-probabilistic version of causal completeness is determinism.

supervenes simply on the pattern of facts about train arrival times. But our causally complete micro-chances may also entail a value for this objective chance. We might consider, for example, the micro-derived chance of the 09:37 train arriving late, given the complete physical state of the world (over an appropriately large region of space) at 06:00 a.m. of the same morning. Such micro-derived chances will presumably be time-variable, and presumably often close to the macro-derived chance; but nonetheless different. Two questions arise from this scenario. First, does this scenario make the Best System self-contradictory? And second, if we *could* somehow know both objective chances, which one would better deserve to guide our credences?

The answer to the first question is negative. All OC's are referred to a chance setup, and the micro-derived, time-variable chances of lateness for the 9:37 have a different chance setup every morning, while the macro-derived chance's setup is relatively permanent. So the two competing chances do not overtly contradict each other, any more than the two breast cancer chances discussed earlier.³⁶

But they are still *competitors*. Suppose God whispers in one ear the macro-level chance, based on the entire history of 09:37 trains in my town, while a Laplacean demon who calculates the micro-derived chance whispers it in your other ear. Which should you use? Common wisdom among philosophers of science suggests that it must be the micro-derived chance. It is, we might think, the chance that is grounded on more complete knowledge, hence favored under the Principal of Total Evidence (see Sober (2004)). But on the contrary, I want to suggest that it could be the macro-derived chance that better deserves to guide credence. How could this be?

First, we'll suppose that the micro-derived chance, though time-variable, is still fairly stable from morning to morning. (We'll suppose we are given the 06:00 chance of lateness each morning, and that it generally hovers around 0.3 except when the network has a serious breakdown.) But the macro-derived chance — in a case like this, essentially the frequency — might be significantly different (say, 0.15). How could this be? Well, nothing rules it out. The micro-level chances are what they are because they best systematize the patterns of outcomes of micro-level chance setups, such as quantum state transitions. And they may do this excellently well, at all times and places in world history. But that entails nothing about

³⁶ It may be that our two posited chances are such that admissibility considerations rule out the use of one, if the other is known, as we saw in the breast cancer case. But it is not clear to me that this must happen in general.

what will happen for train arrivals. The micro-theory of chances in the Best System gets the frequencies right for micro-level events (and reasonable-sized conjunctions and disjunctions of them), to a good approximation, over the entire mosaic. This simply does not entail that the micro-theory must get the frequencies right for sets of distinct one-off setups, each being a horribly complex conjunction of micro-events, from which a chance is calculated for an even more horribly complex *disjunction* of conjunctions of micro-events, all confined to a small fragment of the Humean mosaic. To suppose that the micro-derived probability of lateness is somehow ‘more right’, or ‘the true’ probability, is to indulge in a non-Humean way of thinking about the events in the mosaic: thinking of the micro-level events as ‘brought about’ by chance-laws, rather than as being *simply and compactly characterized* by those laws.

To properly characterize this scenario and demonstrate more rigorously that *nothing* guarantees that the micro-derived chances will be close to the macro-derived chances (and hence, to the frequencies) would require many pages of discussion. I hope the basic point is clear enough: Humean micro-level chances are guaranteed (with the usual probabilistic caveats) to be good guides to credence for micro-level events, but not necessarily for macro-level events. If a significant divergence were to occur, then our argument in section 5 shows that the macro-derived chances would be the ones deserving to guide our credences. Contrary to many philosophers’ intuitions, the micro-level does not automatically dominate over the macro.

Chances are nothing but aspects of the patterns in world history. There are chance-making patterns at various ontological levels. Nothing makes the patterns at one level automatically dominate over those at another; at whatever level, the chances that can best play the PP role are those that count as the ‘real’ objective probabilities. In this section we have seen that the constraint of being able to play the PP role puts limitations on the scope of proper uses of objective probabilities.

7. Conclusion

The view of chance on offer is perhaps the *only* interpretation of chance (besides actual frequentism) able to demonstrate the reasonability of PP in a consequentialist sense, and due to its Best System character it has a number of virtues not possessed by actual frequentism, as we saw in sec. 4.5. By way of closing, I want to turn briefly to the idea of chances as explainers of events. By their very nature, Humean chances are not, in the deepest sense,

explainers of events, but rather a *product* of the events themselves — which are thus logically prior.³⁷ To some extent the existence of SNM's lets us give a non-trivial explanation of events (e.g., that approximately 50% of the coins landed heads in a 1000-flip experiment); but the explanation is grounded on the random-looking distribution of facts at a more micro-level, and *that* has no explanation based on chance. This may be the great defect of the view to those attracted by propensities. If you have metaphysical yearnings that Humean chances just cannot satisfy, then you are welcome to postulate all the 'metaphysical whatnots' you like; and I wish you luck in trying to demonstrate that they exist, or that we can find out their strengths if they do exist, or that they deserve to play the PP role in guiding belief. *Humean* chances exist whether or not there are such Popperian propensities or Hallian primitive chances in the world. If such whatnots do exist, very good; but by their very nature, the Humean chances are demonstrably apt for playing the PP role, and so deserve the title of '*the objective chances*'. Capacities and propensities may explain why the chances are what they are (though I personally have difficulty seeing such explanations as anything better than 'dormitive virtue' explanations) — but that does not make them the chances. And of course, if there are no such things, the Humean chances are still as I have described them (and they have a different explanation, or none at all — a question for another day).

Since scientists and policymakers do postulate objective probabilities, and try to find out what they are — often with apparent success — philosophers of science should not be content with anything less than a theory of objective chances that entails that they exist, that we can come to know them, and that (once known) they deserve to guide action under circumstances of ignorance. Humean objective chance, as I have sketched it, is the only genuine theory of chance that can meet these goals.³⁸

³⁷ The same goes for Humean laws, of course, *pace* Lewis. The traditional empiricist response to this awkward point is to offer a revised account of explanation – one based on deduction, for example, or systematization/unification. In these weakened senses of explanation, Humean chances may also be able to lay claim to explanatory power.

³⁸ *Acknowledgments:* This paper has had an unusually long gestation, and has been steadily improved by the criticisms and suggestions of many readers and colloquium audiences. I wish to thank the following philosophers for comments and/or discussions on chance: Robert Bishop, Nick Bostrom, Craig Callender, Nancy Cartwright, Alan Hájek, Colin Howson, Genoveva Martí, Manuel Pérez Otero, John Roberts, Jonathan Schaffer, Elliott Sober, Michael Strevens, and Henrik Zinkernagel. Special thanks for more extensive help go to Jose Díez, Roman Frigg, Marc Lange, Barry Loewer and Mauricio Suárez. Research for this paper was supported by the Spanish government via the research group projects BFF2002-01552, HUM2005-07187-C03-02, and by the Catalan government via the consolidated research group GRECC, 2001 SGR 00154.

ICREA & Autonomous University of Barcelona
carl.hoefer@uab.es

CARL HOEFER

Bibliography

- Cartwright, N. 1999: *The Dappled World: A Study of the Boundaries of Science*. Cambridge, England: Cambridge U. Press.
- Diaconis, Holmes and Montgomery (forthcoming): ‘Dynamical Bias in the Coin Toss’, to appear in *SIAM Review*; currently available at <http://www-stat.stanford.edu/~cgates/PERSI/papers/headswithJ.pdf>.
- Elga, A. 2004: ‘Infinitesimal Chances and the Laws of Nature’. *Australasian Journal of Philosophy* 82, pp. 67–76.
- Gillies, D. 2000: *Philosophical Theories of Probability*. London: Routledge.
- Hájek, A. 2003: ‘Interpretations of Probability’. *Stanford Encyclopedia of Philosophy* (online), <http://plato.stanford.edu>
- Hájek, A. 2003b: ‘What Conditional Probability Could Not Be’. *Synthese* 137, pp. 273–323.
- Hájek, A. (forthcoming): ‘The Reference Class Problem is Your Problem Too’. *Synthese*, forthcoming.
- Hall, N. 1994: ‘Correcting the Guide to Objective Chance’. *Mind* 103, pp. 504–517.
- Hall, N. 2004: ‘Two Mistakes about Credence and Chance’. *Australasian Journal of Philosophy* 82, pp. 93–111.
- Hofer, C. 1997: ‘On Lewis’ Objective Chance: Humean Supervenience Debugged’. *Mind* 106, pp. 321–334.
- Hofer, C. 2005: ‘Humean Effective Strategies’. In Petr Hájek, Luis Valdés-Villanueva, Dag Westerståhl (eds.), *Logic, Methodology and Philosophy of Science Proceedings of the Twelfth International Congress*. London: KCL Publications.
- Howson, C. and Urbach, P. 1993: *Scientific Reasoning: The Bayesian Approach*. Chicago: Open Court.
- Humphreys, P. 2004: ‘Some Considerations on Conditional Chances’. *British Journal for the Philosophy of Science* 55, pp. 667–680.
- Levi, I. 1983: ‘Review of *Studies in Inductive Logic and Probability*’. *Philosophical Review* 92, pp. 120–1.
- Lewis, D. 1980: ‘A Subjectivist’s Guide to Objective Chance,’ in Richard C. Jeffrey, ed., *Studies in Inductive Logic and Probability*, vol. II. Berkeley: University of California Press. Reprinted in Lewis (1986); page numbers refer to that edition.

- Lewis, D. 1986: 'A Subjectivist's Guide to Objective Chance,' reprint with postscripts added of (1980). *Philosophical Papers*, vol. II. Oxford: Oxford University Press, pp. 83–132.
- Lewis, D. 1994: 'Humean Supervenience Debugged'. *Mind* 103, pp. 473–90.
- Loewer, B. 2001: 'Determinism and Chance'. *Studies in the History and Philosophy of Modern Physics* 32, pp. 309–320.
- Loewer, B. 2004: 'David Lewis' Humean Theory of Objective Chance'. *Philosophy of Science* 71, pp. 1115–1125.
- Mellor, H. 1995: *The Facts of Causation*. London: Routledge.
- Sober, E. 2004: 'Evolutionary Theory and the Reality of Macro Probabilities'. *PSA 2004 Presidential Address*. Also in E. Eells and J. Fetzer (eds.), *Probability in Science*. Open Court, forthcoming.
- Strevens, M. 1999: 'Objective Probability as a Guide to the World'. *Philosophical Studies* 95, pp. 243–275.
- Suppes, P. 1993: 'The transcendental character of determinism' in P.A. French, T.E. Uehling and H.K. Wettstein (eds.), *Midwest Studies in Philosophy, Vol. XVIII*. Notre Dame, IN: University of Notre Dame Press, pp. 242-257.
- Thau, M. 1994: 'Undermining and Admissibility'. *Mind* 103, pp. 491–503.
- von Mises, R. 1928: *Probability, Statistics and Truth*. Translated 2nd edition: Dover books, 1981.