

Explaining (One Aspect of) the Principal Principle without (Much) Metaphysics

According to David Lewis's (1980) Principal Principle, our beliefs about the objective chances of outcomes (typically) determine our rational credences in those outcomes. In his (1994), Lewis influentially argues that any adequate metaphysics of objective chance must explain why the Principal Principle holds. Since no theory of chance is widely agreed to have met this burden, I suggest we change tack. On the view I develop, a central aspect of the Principal Principle holds not because of what objective chances are, but rather because of the explanatory role that objective chances play for the rational agents who believe in them.

Katrina Elliott

Department of Philosophy

UCLA Box 951451 Dodd 321

Los Angeles, CA 90025

310.267.4847

kelliott@humnet.ucla.edu

My thanks to Josh Armstrong, John Carriero, Daniela Dover, Gabriel Greenberg, Matthew Kotzen, Marc Lange, John Roberts, Seana Shiffrin and four anonymous referees. Thanks also to audiences at California State University Northridge, California State University Los Angeles, and University of California Irvine.

According to David Lewis's (1980) Principal Principle, our beliefs about the objective chances of outcomes (typically) determine our rational credences in those outcomes. In his (1994), Lewis influentially argues that any adequate metaphysics of objective chance must explain why the Principal Principle holds. Since no theory of chance is widely agreed to have met this burden, I suggest we change tack. On the view I develop, a central aspect of the Principal Principle holds not because of what objective chances are, but rather because of the explanatory role that objective chances play for the rational agents who believe in them.

1. Introduction.

To paraphrase Bishop Butler (1736), chances are the very guide to life; ordinarily, an agent's beliefs about the objective chances of various outcomes determine her rational credences in those outcomes. David Lewis (1980) famously offered a codification of the relationship between chance and credence. Roughly, Lewis's principle implies that if a rational agent (i.e., an agent whose opinions can be represented as having come from a reasonable initial credence function by conditionalizing on her total evidence) believes an outcome to have a given chance of occurring, then her credence in that outcome is equal to that chance. For example, if I believe the chance of rain tomorrow is .7, then my credence in rain tomorrow (if I am rational) is .7.

The principle holds with one exception. A reliable crystal ball might predict that there will be no rain tomorrow despite there now being a high chance that it will rain tomorrow. If I believe that the crystal ball is reliable, I should put in my lot with the ball's prediction and not set my credence to what I believe the chance of rain to be. Reliable crystal balls provide a kind of information about future outcomes that, to use Lewis's metaphor, does not "go by way of" the chances of those outcomes. Every such case (where an agent ought to ignore the chances when forming her credences), argued Lewis, is one in which an agent has "inadmissible" information.

Lewis named his formulation the "Principal Principle" (henceforth "PP") since it describes perhaps *the* essential feature of objective chance, and argued that any adequate

theory of chance must explain why we ought to form our credences in accord with the PP. Though there is debate about whether the PP best captures the relationship between chance and credence (e.g., Lewis 1994; Hall 1994; Nelson 2009; Meachum 2010), many philosophers (e.g., Salmon 1967; Hájek 2007) agree with Lewis that theories of objective chance are burdened with showing how it is that chances are the very guide to life. But this has proved a difficult burden to meet. No theory of chance is widely agreed to give a satisfactory explanation of the PP (or of some appropriately similar principle), and some philosophers (e.g., Strevens 1999) doubt that such an explanation can be found.¹

In what follows, I will argue for an alternative approach to explaining the PP. To begin, we can distinguish between two distinct aspects of the PP that call for explanation. The first is the particular quantitative relationship it implies between a rational agent's credences regarding the chances of various outcomes and her credences in those outcomes. If a rational agent is certain that the chance of some outcome is n (a real number between 0 and 1 inclusive), then her credence in that outcome is equal to n (so long as she has only admissible information). But why does rationality require that the values be *equal*?

¹ Important arguments have been offered that certain theories of chance *do* explain the PP (e.g., Loewer 2004; Frigg and Hoefer 2010; Schwarz 2014). Even if one of these arguments is correct, an explanation of the PP that can be employed by additional theories of chance would be a significant discovery.

A second aspect of the PP that calls for explanation is more basic. Never mind wondering why beliefs about chances determine rational credences in the *particular way* described by the PP; why do beliefs about chances determine rational credences in *any way at all*? There is no particular credence in rain, for example, that I am rationally obliged to have given that I believe that there are storm clouds overhead. But unlike my opinions about storms clouds, my opinions about the *chance* of rain (if I have any such opinions) rationally compel me to have a very particular credence that it rains, no matter what other admissible information I have. I call this second aspect of the relationship between chance and credence “SWAMPING”:

SWAMPING. A rational agent’s credence in the chance of an outcome determines her credence in that outcome, so long as she has only admissible information.²

SWAMPING is at the heart of the relationship between chance and credence. Chance is aptly described as the very guide to life precisely because a rational agent’s credence in an

² SWAMPING and the PP apply not only to cases where an agent is *certain* of the chance of an outcome, but also to cases where she merely has various opinions about the chance of that outcome. In such cases, an agent’s rational credence in an outcome is determined by a weighting of her opinions about the chance of that outcome.

outcome's chance "swamp" all of her other (admissible) evidence about that outcome.³

Explaining SWAMPING, then, is a significant step toward explaining the PP. Unfortunately, theories of chance have had no better luck at explaining SWAMPING than they've had at explaining the PP.⁴

Perhaps we have failed to discover an explanation of SWAMPING (and ultimately, the PP) because we have been looking in the wrong place. The goal of this essay is to show how SWAMPING might be explained without appeal to any particular metaphysics of chance. On the view I develop, SWAMPING holds not because of what chances are, but rather because of the explanatory role they play for the rational agents who believe in them.

I'll proceed as follows. In Section 2, I introduce a principle that articulates one aspect of the relationship between beliefs about *scientific explanations* of outcomes and rational expectations in those outcomes. I call this principle the "Hypothetical Explanation-Credence Principle" or "HCP" for short. In Section 3, I give an argument from HCP to SWAMPING and

³ SWAMPING is to be distinguished from the Bayesian idea of "the swamping of the priors" in which agents with different priors converge on the same credences given enough shared evidence. Still, the flavor of the two phenomena is similar in that both involve some evidence rendering some credences irrelevant.

⁴ Because it is more familiar, I focus on the PP rather than on rival formulations of the connection between chance and credence, but SWAMPING (or something like it) is at the heart of *any* plausible chance-credence principle.

I identify two problems this argument faces: the admissibility problem and the generalization problem. The admissibility problem arises because HCP, like SWAMPING, applies only to cases in which rational agents have no inadmissible information. In Section 4, I sketch a novel account of admissible information and argue that, though this account is imprecise, it plausibly solves the admissibility problem. The generalization problem arises because SWAMPING is an instance of HCP only if objective chances play a certain explanatory role for rational agents. In Section 5, I articulate a novel account of this explanatory role that I call “the causal chain account”. Rather than argue directly for the causal chain account, I show that it solves the generalization problem. I conclude in Section 6 by arguing that, though my strategy for explaining SWAMPING is consistent with a wide array of metaphysical theories of chance, some well-known theories are a worse fit with my strategy than are others.

2. The Hypothetical Explanation-Credence Principle.

According to SWAMPING, beliefs about chances determine rational credences in typical cases in which a rational agent is making a prediction (i.e., cases in which a rational agent has no “inadmissible” information). If our goal is to better understand why SWAMPING is true, a good way to start is by exploring what other kinds of beliefs stand in a similar relationship to our rational credences in typical cases of prediction.

Suppose that a rational agent is predicting whether I will get a cold. If she has standard background beliefs about colds, she bases her expectations about my future health on factors such as whether I’ve recently had contact with sick colleagues, the frequency with

which I disinfect my hands, and so on. She bases her expectations on *these particular* factors (rather than on, say, whether I drive a Volkswagen) because she believes that these factors *explain why* I get a cold if I do. The moral I draw from this example and others like it is that beliefs about explanations (typically) play an important role in determining rational credences about the targets of those explanations. The Hypothetical Explanation-Credence Principle is my attempt to articulate at least one aspect of that role.

Two points of clarification will be helpful for motivating HCP. First, it is notoriously difficult to give an analysis of “scientific explanation” that is neither too restrictive nor too broad. Thankfully, we can largely do without any particular analysis for my present purposes. I am interested in the relationship between *beliefs* about explanation and rational credence: independent of the truth of such beliefs. Accordingly, I’ll assume as little as possible about the correct theory of scientific explanation in motivating HCP.⁵

Second, our rational agent need not believe that, say, my recent contact with sick colleagues *explains* my cold in order for her confidence in the former to influence her confidence in the latter. The fact that *P* explains *Q* (at least on one common understanding of “explains”) guarantees that *P* and *Q* obtain, but our agent may be skeptical that I’ll get a cold precisely because she’s skeptical that I’ve had contact with sick colleagues. For her confidence that I get a cold to be constrained by her confidence that I’ve had contact with

⁵ After motivating HCP, I will make substantive and explicit claims about the nature of scientific explanation.

sick colleagues, it is sufficient that she believe that *if* it is true that I've had contact with sick colleagues and *if* it is true that I get a cold, then my contact with my colleagues explains my cold. In other words, it is enough that she believes my contact with sick colleagues is a “hypothetical explanation” of my cold:

Hypothetical explanation. For any hypothetical states of affairs X and P , an agent believes X to be a “hypothetical explanation” of P to the degree to which she believes that X explains P if X and P obtain.⁶

Our example of the rational agent who is predicting whether I get a cold can now be made more precise. Our agent has various degrees of confidence in propositions of the form “ X is a hypothetical explanation of P ”.⁷ She may, for example, be very confident that my

⁶ I treat an agent's credence that X is a hypothetical explanation of P as a credence in an indicative conditional, rather than as a conditional credence. Though Lewis (1976) demonstrates that probabilities of conditionals are not conditional probabilities, an agent's conditional credence that X explains P given X and P is a reasonably good guide to her credence that X is a hypothetical explanation of P . Accordingly, I suspect that nothing essential to what follows hinges on my choice.

⁷ I will not be strict about treating “ X ” and “ P ” as variables that range over hypothetical states of affairs. Sometimes I write as if they range over propositions (so as to avoid writing, e.g.,

recent contact with sick colleagues is a hypothetical explanation of my cold and sure that my car's make is not a hypothetical explanation of my cold.⁸ Furthermore, she has various degrees of confidence that X 's she takes to be (or not to be) hypothetical explanations obtain. She may be unsure whether I have had contact with sick colleagues, for example, but very confident that I own a Volkswagen.

Our agent's expectations about whether I get a cold are sensitive to her opinions about which hypothetical states of affairs are hypothetical explanations of my cold in tandem with her opinions about which hypothetical states of affairs actually obtain. Because she is confident that my contact with sick colleagues is a hypothetical explanation of my cold, for example, her expectation about whether I get a cold may vary with her credence that I've had contact with sick colleagues. Similarly, if she is confident that I've had contact with sick colleagues, her expectation about whether I get a cold may vary with her credence that my contact with sick colleagues is a hypothetical explanation of my cold. Finally, because she is sure that my car's make is not a hypothetical explanation of my cold, her credence that my car is a Volkswagen makes no difference to her expectation about my cold.

“ X obtains”) and sometimes I treat them as dummy letters that stand in for particular states of affairs or propositions (so as to introduce as little notation as possible).

⁸ More realistically, she believes that contact with sick colleagues combined with further facts is a hypothetical explanation of my getting a cold. I omit these further facts for ease of discussion.

This example is one instance of a broad range of cases in which a rational agent's expectations about P are sensitive to her opinions about propositions of the form " X is a hypothetical explanation of P " and " X obtains". In such cases, her credence in P may change when she acquires new information that changes her opinions about whether X is a hypothetical explanation of P or about whether X obtains. However, I will argue that once a rational agent is certain that a particular X obtains and certain that X is a hypothetical explanation of P , new information typically makes no difference to her credence in P . The rare cases in which a rational agent does change her credence in P despite being certain that X obtains and that X is a hypothetical explanation of P are ones in which the agent has inadmissible information. This claim is the germ of a principle that I call the "Hypothetical Explanation-Credence Principle":

Hypothetical explanation-credence principle (HCP). For any hypothetical states of affairs P and X , any admissible bodies of evidence $E1$ and $E2$, and any reasonable initial credence function C , if $C(X \text{ hypothetically explains } P|E1) = C(X \text{ hypothetically explains } P|E2) = 1$, then $C(P|X \& E1) = C(P|X \& E2)$.

Roughly, HCP says that if a rational agent is certain that X is a hypothetical explanation of P and she has no inadmissible information, the value of her credence in P given that X obtains is independent of the content of her admissible body of evidence.

Is HCP true? That depends on what it is for information to be “admissible”. But if “admissible” turns out to have the same meaning when used in HCP as it does when used in SWAMPING, then HCP implies that a rational agent’s admissible evidence about P is *swamped* by her beliefs about hypothetical explanations of P in much the same way that a rational agent’s admissible evidence about P is *swamped* by her beliefs about the chance of P .

3. From HCP to SWAMPING.

Consider an instance of HCP in which X is that the chance of P is equal to x (a real number on the unit interval). Suppose some rational agent is certain that the chance of P is equal to x and certain that the chance of P (i.e., the fact that P has a given chance) is a hypothetical explanation of P .⁹ Finally, assume she has only “admissible” information (whatever that might mean). Then, by HCP, her credence in P is determined by her initial credence function C (i.e., her credence function prior to updating on X and her admissible body of evidence) and the content of X and P . Because X is the proposition that the chance of P is equal to x , X specifies the content of P . So, in this special case of HCP, C and the content of X determines her credence in P . Holding fixed whatever C with which our agent began, it follows that her credence in P is determined by her opinions about the chance of

⁹ If, for example, P is a given coin’s landing heads and x is .5, then the agent is certain that the chance of the coin’s landing heads is equal to .5, and certain that the .5 chance of the coin’s landing heads is a hypothetical explanation of the coin’s landing heads.

P .¹⁰ So, in this special case, HCP implies that an agent's opinions about chances *swamp* all of her other admissible evidence.

HCP also has implications for agents who distribute their credences over (a potentially infinite) number of (mutually exclusive and collectively exhaustive) propositions about the value of the chance of P . If a rational agent is certain that *each* such chance proposition is a hypothetical explanation of P and she has only admissible information, then HCP and the theorem of total probability imply that her credence in P is determined by C , the content of each chance proposition, and her credence in each chance proposition. In this case, HCP (like SWAMPING) implies that admissible information has no impact on her credence in P except by way of changing her credence in some chance proposition.

What would it take for this for this argument to apply to every case to which SWAMPING applies? First, "is admissible" as used in HCP must be coextensive with "is admissible" as used in SWAMPING in every case in which X is that the chance of P is equal to

¹⁰ Notice that HCP does not imply that any two rational credence functions that satisfy HCP *agree* on the credence in P . If our goal is to explain the PP, we must ultimately explain why two different rational agents' credences in P are not only each determined by their credences in X , but are also determined to be the same value.

x .¹¹ To avoid confusion, let “admissible_{PP}” express the concept that “admissible” expresses as it appears in SWAMPING (and the PP). Let “admissible_{HCP}” express the concept that “admissible” expresses as used in HCP. On one hand, HCP requires a somewhat narrow reading of “admissible_{HCP}”, since potential counterexamples to HCP grow as admissibility_{HCP} is broadened.¹² In particular, “admissible_{HCP}” must be narrow enough to exclude all information that is inadmissible_{PP} in every case in which X is that the chance of P is x . Otherwise, canonical cases of inadmissible_{PP} information (e.g., the predictions of a reliable crystal ball) are counterexamples to HCP. On the other hand, “admissible_{HCP}” must not be given *too* narrow a reading. SWAMPING is meant to apply to all typical cases of prediction, and so all information that we typically have when we make predictions is admissible_{PP}.¹³ Similarly, in cases in which X is that the chance of P is x , “admissible_{HCP}” must be broad enough to include all information that is had in typical cases of prediction. The “admissibility problem” is to provide an account of admissible_{HCP} information that is narrow

¹¹ My argument does not require that “admissible_{HCP}” is coextensive with “admissible_{PP}” in cases in which X is something *other* than the chance value of P , because SWAMPING does not apply to such cases.

¹² Analogously, SWAMPING and the PP require a somewhat narrow reading of “admissible_{PP}” to protect them from spurious counterexamples.

¹³ When I refer to cases to which HCP or SWAMPING “apply”, I refer to cases in which HCP or SWAMPING are not merely trivially true.

enough to protect HCP from spurious counterexamples but broad enough for HCP to apply to every case to which SWAMPING applies.¹⁴

Second, the argument from HCP to SWAMPING seems to generalize only to cases in which a rational agent is certain that the chance of an outcome is a hypothetical explanation of that outcome. But, as we'll see in Subsection 5.1, there are cases in which a rational agent obeys SWAMPING but does not regard the chance of an outcome to be a hypothetical explanation of that outcome. The "generalization problem" is to provide an account of the explanatory role that objective chances play that allows the argument from HCP to SWAMPING to generalize to every case to which SWAMPING applies.

At this stage, some readers will doubt that there *is* a solution to the generalization problem, on the grounds that SWAMPING cannot be accounted for by the explanatory role that

¹⁴ Perhaps, as many philosophers have argued, the true chance-credence principle makes no reference to admissible information. I suspect that understanding admissibility is nevertheless crucial to understanding *why* that principle is true, because the role that chances play in determining our rational credences when we have only admissible information is importantly different from the role they play when we have some inadmissible information. Even if the true chance-credence principle (unlike the PP) applies to both cases, it does not apply to both cases for the same reason.

chances play. I consider two objections that might motivate such skepticism.¹⁵ The first is simply that there are rational agents who do not believe that objective chances are hypothetical explanations.¹⁶ Chances seem to play *no* explanatory role for such agents and so HCP does not apply to them. One reason why a rational agent might deny that objective chances are hypothetical explanations is that she believes that, though chance outcomes have hypothetical explanations, objective chances are merely parts of those hypothetical

¹⁵ Cases of possible outcomes with a chance of 0 (e.g. a random continuous variable's taking a particular value) present a further difficulty, since it is unintuitive that an outcome's having no chance of occurring explains its occurrence. Nevertheless, I am optimistic that either a chance of 0, despite appearances, is explanatory (since the occurrence of an outcome with a chance of 0 is no more inexplicable than is the occurrence of an outcome with a very low chance) or that the problem can be avoided by the discovery that events that have a chance of 0 according to standard probability theory in fact have infinitesimal chances because a non-standard probability theory holds.

¹⁶ Rationality does not require agents to believe that there are non-extremal objective chances (i.e., chances that take values between 0 and 1), but denying that an outcome has a non-extremal chance is consistent with being certain that any non-extremal chance of that outcome is a *hypothetical* explanation of that outcome. SWAMPING (like the PP) has interesting implications only for rational agents who are considering outcomes that at least might, by their own lights, have non-extremal chances of occurring.

explanations. Among such agents, distinguish those who believe that objective chances (by themselves) are not explanatory from those who believe that objective chances (by themselves) merely fail to satisfy some ideal of explanation. For HCP to apply to the latter kind of agent, “hypothetical explanation” must be understood as referring to minimally adequate, rather than ideal, explanations. I suspect that HCP remains plausible even when so interpreted, though an argument to that effect would require a delicate discussion of minimally adequate explanations. The greater difficulty is with the former sort of agent, who believes that chances are parts of hypothetical explanations but denies that they are (by themselves) even minimally adequate hypothetical explanations. I doubt that such agents have sufficient conceptual grasp of explanation and objective chance to count as believing in objective chances, but defending that claim would require a serious discussion of scientific explanation and its conceptual ties to objective chance. I aim only to show that *if* objective chances play a certain explanatory role for the rational agents who believe in them, then SWAMPING follows from HCP. That said, if I meet this goal then that is some (admittedly defeasible) evidence that chances do indeed play that explanatory role.

Another reason why a rational agent might deny that objective chances are hypothetical explanations is a commitment to the view that chance outcomes have no explanations. I find this view implausible, but it is neither irrational nor conceptually

confused. The possibility of such agents motivates our second objection: it seems that SWAMPING is true of agents who are not certain of *any* hypothetical explanations.¹⁷

Consider a rational agent who wonders whether a coin will land heads on a particular toss. Suppose she is certain that the relative frequency with which this coin actually lands heads (over the entire span of its existence) is .5. Further, suppose she doubts that this relative frequency, or anything else, is a hypothetical explanation of the coin's landing heads on a given toss. Though it seems that HCP does not apply to such an agent (since there is nothing she believes to be a hypothetical explanation of the coin's landing heads), her credence in heads is (plausibly) determined by her certainty that the actual relative frequency of heads is .5. Beliefs about actual relative frequencies obey the analogue to SWAMPING that we get by replacing "chances" with "actual relative frequencies". It is implausible that HCP is a genuine explanation of SWAMPING if it is silent about such closely related principles.

Despite appearances to the contrary, HCP *does* apply to this case. To see why, start by imagining that we are about to draw a ball from an urn filled with an equal number of black and red balls. How confident should we be that a red ball will be drawn? Simply knowing that half the balls are black and half are red is not, I claim, enough to answer that question; instead, we need some further information about *chances*. For example, if the chance that a given red ball is picked is greater than is the chance that a given black ball is picked, our confidence that a red ball is picked should be greater than .5. If, instead, all of the balls have the same chance of being picked, our confidence that a black ball is selected

¹⁷ Thanks to an anonymous reviewer for bringing this powerful objection to my attention.

should be .5. If we have no information about the chances, we are allowed to proceed *as if* all of the balls have the same chance of being picked.¹⁸ But knowing that half the balls are black and half are red, without any further assumption, is insufficient to determine our credence that the drawn ball is red.

Observing one of a series of actual coin tosses, half of which land heads and half of which land tails, is importantly like drawing one ball from an urn that contains an equal number of red and black balls. Just as knowing that half the balls in the urn are red does not (by itself) determine our credence that the drawn ball is red, knowing that half the coin tosses land heads does not (by itself) determine our agent's credence that the next toss lands heads. Instead, if she does not know each toss's chance of being the one she observes, she proceeds *as if* each toss has the same chance of being the one she observes. If she is certain that chances are hypothetical explanations, HCP entails that her beliefs about the chances of observing each coin toss determine her credences about which coin toss she will observe. In turn, those credences combine with her certainty that half the tosses land heads to determine her credence in heads. So, HCP helps to explain why rational agents set their credences in accord with actual relative frequencies even if those agents do not believe that actual relative frequencies are hypothetical explanations.

¹⁸ The intuition that we are so allowed is similar to intuitions that motivate the principle of indifference.

What of a rational agent who believes that actual relative frequencies are not hypothetical explanations and that chances just are actual relative frequencies? HCP does not apply to such an agent, since she is not certain that chances are hypothetical explanations. Perhaps that result is not so bad. The moral of the analogy between the urn case and the coin case is that beliefs about actual relative frequencies do not, by themselves, determine our rational credences. It is not clear that agents for whom beliefs about chances are just more beliefs about actual relative frequencies (and who deny that actual relative frequencies are hypothetical explanations) obey SWAMPING.¹⁹

Furthermore, it is independently plausible to me that SWAMPING and the PP do not apply to agents who believe that chance outcomes are inexplicable. For example, we know precisely how confident I should be that a fair coin lands heads, but we have no idea what my precise confidence should be that my wallet inexplicably begins to sing. I submit that part of the reason there are definitive answers to questions about how confident I should be in the outcomes of coin tosses, but no definitive answers to questions about how confident I should be that inanimate objects begin to sing, is that I am certain of hypothetical explanations of the former but regard the latter as inexplicable. For an agent who believes that chance outcomes are as inexplicable as I believe singing wallets to be, perhaps there are no answers (beyond

¹⁹ Intuitions to the contrary might be the result of overgeneralizing from cases in which beliefs about actual relative frequencies, when not conceived of as chances, obey an analogue to SWAMPING.

mere Bayesian updating) to questions about how confident to be in the occurrence of chance outcomes.

4. The Admissibility Problem.

According to HCP, a rational agent who is certain that X and that X is a hypothetical explanation of P does not change her credence in P when she learns new $\text{admissible}_{\text{HCP}}$ information. Recall that, in order to get from HCP to SWAMPING, we need an account of $\text{admissible}_{\text{HCP}}$ information that is narrow enough to protect HCP from spurious counterexamples but broad enough for HCP to apply to all typical cases of prediction (i.e., to all cases to which SWAMPING applies).

Suppose a rational agent is certain that I have had contact with sick colleagues and that having contact with sick colleagues is a hypothetical explanation of my cold. Does her expectation about my cold change if she learns new information? Unsurprisingly, that depends on the kind of information she learns.

A rational agent's certainty that X and that X is a hypothetical explanation of P screens off a great deal of information that might otherwise make a difference to her credence in P . Learning that I attended a department meeting, for example, makes no difference to our agent's credence about whether I get a cold since it is evidence that I get a cold only in virtue of being evidence that I've had contact with sick colleagues. Similarly, learning that cold viruses are passed through touch makes no difference to our agent's credence that I get a cold, since she is already certain that contact with sick colleagues is a

hypothetical explanation of my cold. So, if a rational agent is certain that X and that X is a hypothetical explanation of P , then the following kinds of information make no difference to her credence in P :

- (1) Information that is evidence about P only in virtue of being evidence about X .
- (2) Information that is evidence about P only in virtue of being evidence about whether X is a hypothetical explanation of P .

I have no precise account of “evidence” to offer.²⁰ The rough idea is that information is evidence about some P for an agent exactly in case she should, given her other opinions, take that information to confirm or disconfirm P .²¹ The sense of “evidence” I employ, then, is always relativized to the background beliefs of a particular agent. Since this understanding of evidence will be central to my account of admissible_{HCP} information, whether information is admissible_{HCP}, on my view, is similarly relativized.

Notice that not *all* evidence about P is screened off by a rational agent’s certainty that X and that X is a hypothetical explanation of P . If, for example, our agent learns about events

²⁰ Furthermore, I rely on the reader’s intuitive understanding of what it is for information to be, e.g., evidence about P “only in virtue of” being of evidence about X .

²¹ Providing a Bayesian precisification of “evidence” that does not run afoul of the problem of old evidence (see Glymour 1980) is more than I’ll attempt here.

that occur subsequent to my catching a cold, she may change her credence that I get a cold. Furthermore, learning about events that occur prior to my catching a cold, such as the predictions of a reliable crystal ball, may change her credence that I get a cold. However, we do not have access to such information in typical cases of prediction and so it can safely be regarded as inadmissible_{HCP}.

What happens if a rational agent learns that something other than X , say Y , might be a hypothetical explanation of P ? If she is certain that there is a *unique* explanation of P , then her credence in P does not change. Because she is certain that X and that X is a hypothetical explanation of P , *and* certain that there is only *one* explanation of P , she must be either certain that Y is not a hypothetical explanation of P or certain that Y does not obtain. Either way, learning that some Y is a hypothetical explanation of P makes no difference to her credence in P .

However, rationality allows agents to believe that there are compatible yet distinct explanations.²² Cases involving such agents are potential counterexamples to HCP because, to put the matter metaphorically, compatible hypothetical explanations of P might give a rational agent conflicting advice about what should be her credence in P . Suppose our agent is certain that I have had contact with sick colleagues *and* that I have a vitamin C deficiency. Suppose she then becomes convinced that my vitamin C deficiency, like my contact with

²² For an excellent discussion of compatible yet distinct explanations of a single event, see Salmon's (1989) treatment of the issue.

sick colleagues, is a hypothetical explanation of my cold. Then, her updated credence that I get a cold—made in light of the new information that my vitamin C deficiency is a hypothetical explanation of my cold—may differ from her old credence that I get a cold. We cannot argue that information about a hypothetical explanation of my cold other than my contact with sick colleagues is somehow inadmissible_{HCP} for this agent, because information about alternative hypothetical explanations of chance outcomes is surely admissible_{PP}. So, we seem to have shown that HCP is false.

This potential counterexample relies on the assumption that *X* and *Y* might give a rational agent incompatible advice about what to expect with respect to *P* when she is certain that both *X* and *Y* are hypothetical explanations of *P*. In advocating HCP, I reject this assumption. It seems to me central to our conception of explanation that we are not certain of why an event occurs (if it does) unless we are also certain that we are in at least the best possible epistemic position (relative to that event) available to agents who have only admissible_{HCP} information (in the sense of “admissible_{HCP}” developed below). If that is right, then if a rational agent with only admissible_{HCP} information is certain that *X* obtains *and* that *X* is a hypothetical explanation of *P*, then either she is not certain that *Y* is a hypothetical explanation of *P* or she has ruled out *Y* or learning that *Y* obtains would not (by her own lights) put her in a better epistemic situation with respect to forming expectations about *P*.

Influential theories of statistical explanation (Hempel 1965; Salmon 1971; Railton 1978) imply a similarly high standard for certainty about why an event occurs, but it may strike some readers as implausibly demanding (or at least rational to deny). For example,

given that we know that we are rarely in the best possible epistemic situation that is consistent with having admissible_{HCP} information, we are rarely rationally certain of some X that we are certain is a hypothetical explanation of P . I find this implication intuitive, but admittedly it rules out any theory of explanation according to which it is easy to know why events occur. At any rate, if we assume that rational agents never make incompatible predictions about P in light of (what they believe to be) compatible hypothetical explanations of P , then learning that Y obtains or that Y is a hypothetical explanation of P makes no difference to a rational agent's credence in P once she is certain that X and that X is a hypothetical explanation of P .

There is one final wrinkle brought out by this case. In virtue of what is the information that I am vitamin C deficient admissible_{HCP} for our rational agent who is predicting whether I get a cold while a reliable crystal ball's predictions about whether I get a cold are not? The answer is that our agent believes that it is at least a *live-option* that a vitamin C deficiency is a hypothetical explanation of my cold, whereas she has *ruled out* the possibility that the predictions of a reliable crystal ball are a hypothetical explanation of my cold (so long as she has background beliefs like ours). Whatever other hypothetical explanations of my cold there may be, she is sure that the predictions of reliable crystal balls are not among them.

In light of these considerations, our list of information that is available in typical cases of prediction but that makes no difference to a rational agent's credence in P once she is certain that X and that X is a hypothetical explanation of P grows to include the following:

(3) Information that is evidence about P only in virtue of being evidence about some hypothetical state of affairs Y such that it is a live-option (for the agent) that Y is a hypothetical explanation of P .

(4) Information that is evidence about P only in virtue of being evidence about whether some hypothetical state of affairs Y is a hypothetical explanation of P .

Finally, in a typical case of prediction, the vast majority of information a rational agent has is not evidence about P . Since information that is not evidence about P makes no difference to a rational agent's credence in P , the following information makes our list:

(5) Information that is not evidence about P .

An account according to which information is admissible_{HCP} (with respect to a particular rational agent's credence function and to particular hypothetical states of affairs X and P) if and only if it satisfies at least one of (1)-(5) would be adequately broad and would protect HCP from *most* spurious counterexample—but not all. Suppose, once again, that a rational agent is predicting whether I get a cold. She is certain that I have had contact with sick colleagues and certain that my contact with sick colleagues is a hypothetical explanation of my cold. Suppose also that it is live-option for her that vitamin C deficiencies are hypothetical explanations of colds, though she is presently not very confident that I have a

vitamin C deficiency. Finally, suppose she is certain of the following disjunction: either I do not have a vitamin C deficiency or I will not get a cold. What happens to her credence that I get a cold if she learns that I hate citrus fruit? If she has normal background beliefs about the connection between citrus and vitamin C, she may increase her credence that I am vitamin C deficient. But, because she is certain that either I do not have a vitamin C deficiency or I do not get a cold, she may also *decrease* her credence that I get a cold. If all of her information in this case is $\text{admissible}_{\text{HCP}}$, HCP is false. Condition (3) implies that the information that I hate citrus fruit is $\text{admissible}_{\text{HCP}}$. Whether the disjunction (that either I do not have a vitamin C deficiency or I will not get a cold) is $\text{admissible}_{\text{HCP}}$ depends on whether it counts as evidence about my cold, and my use of “evidence” (and “in virtue of”) in (1)-(5) is too imprecise to yield a definitive answer.

Nevertheless, I am optimistic that the admissibility problem can be solved with an account of “ $\text{admissible}_{\text{HCP}}$ ” that is in the spirit of (1)-(5). Conditions (1)-(5) are not a disjointed collection of ad hoc stipulations. Information that fails to satisfy at least one of (1)-(5) is, roughly, information that is evidence about P but not only in virtue of being explanatorily relevant to P . The information that I hate citrus fruit is intuitively *not* $\text{admissible}_{\text{HCP}}$ because it provides our agent with evidence about my cold via *both* an explanatory route (because it satisfies (3)) *and* a non-explanatory route (through her confidence in the disjunction). It is hard to say what it is for information to provide evidence through an “explanatory route” rather than not, or to be “explanatorily relevant” rather than not, but the intuition behind these distinctions is reasonably clear. Every example of

inadmissible_{PP} information I know of involves evidence about a chance outcome that is not merely explanatorily relevant to that outcome, so it is plausible that this rough characterization of “admissible_{HCP}” is sufficiently narrow to protect HCP from spurious counterexamples.

Furthermore, this rough characterization of admissible_{HCP} information seems to be broad enough to allow HCP to apply in all typical cases of prediction (i.e., cases to which SWAMPING applies). The examples that motivate (1)-(5) suggest that, in typical cases of prediction, all of our information is either not evidence about P , or is explanatorily relevant to P , or is screened off by information that is explanatorily relevant to P . I conclude that it is plausible that HCP is true on a reading of “admissible_{HCP}” that is coextensive with “admissible_{PP}” in every case in which X is that the chance of P is x .²³

5. The Generalization Problem.

5.1 The Causal Chain Account.

In arguing from HCP to SWAMPING (in Section 3), I began by assuming that some rational agent is certain that the chance of P is a hypothetical explanation of P . But rational agents cannot be certain that the chance of *any* given outcome is a hypothetical explanation

²³ A thorough defense of this claim requires a discussion of how admissibility_{HCP} compares to other accounts of admissibility_{PP} (e.g., Thau 1994; Hofer 2007; Meacham 2010).

of that outcome, on pain of contradicting HCP. To see why, suppose a rational agent is considering whether she will develop lung cancer and is certain at time $t/$ that her current chance of developing lung cancer (in, say, the next twenty years) is .25. Furthermore, suppose that she has only admissible information. If she is certain that her present chance of developing lung cancer is a hypothetical explanation of lung cancer, then HCP entails that she does not change her credence that she develops lung cancer when she acquires new admissible information. Now imagine she learns the admissible information that her chance of developing lung cancer has increased to .5. By the PP (i.e., the very principle HCP is supposed to illuminate!), a rational agent's credence that she develops lung cancer does not remain .25 when she learns that her present chance of lung cancer is .5. So, if rational agents consider the chance of any given outcome to be a hypothetical explanation of that outcome, HCP is false.

The solution is to adopt a more nuanced view of the explanatory role that objective chances play. Consider the distinction between being a hypothetical explanation of an outcome and being a hypothetical explanation of the occurrence of some causal chain ending in that outcome. A hypothetical explanation of the occurrence of some causal chain ending in an outcome is not guaranteed to be a hypothetical explanation of that outcome. For example, my dog and I head toward the park at three o'clock every afternoon. Sometimes I change my mind on the way and we go for a hike instead. Other times we make it to the park and we play fetch. Occasionally, traffic is bad and we come home after sitting in the car for an hour. But no matter what we do after we head toward the park, my dog ends up

exhausted. Today, we played fetch at the park. Why did some causal chain, starting at 3 o'clock and ending now with my dog being exhausted, occur? Because my dog and I headed toward the park at 3 o'clock. Why is my dog now exhausted? If all you know about my afternoon is that my dog and I headed toward the park at 3 o'clock, then you do not know why my dog is exhausted. My dog is exhausted because we played fetch.

Of course, this example hardly settles the matter. Many readers will be reasonably resistant to the view that an explanation of the occurrence of some causal chain ending in an event is not guaranteed to be an explanation of that event, and an argument that scientific explanation is not transitive may be required to defend it. But if that view were correct, it would motivate the theory of the explanatory role of chance that I call the "causal chain account":

Causal chain account. If it is part of a rational agent's background beliefs that there are events that form a causal chain between the occurrence of a particular outcome and conditions that obtain at some earlier time t , then she is certain that the chance at t of that outcome is a hypothetical explanation of the occurrence of some causal chain beginning at t and ending in that outcome, rather than a hypothetical explanation of that outcome. If instead she believes that conditions at t are a direct cause of that

outcome, or that the outcome is uncaused, then she is certain that the chance at t of that outcome is a hypothetical explanation of that outcome.²⁴

Rather than argue directly for the causal chain account, I show that the generalization problem is solved if the causal chain account is true.

5.2 Solving the Generalization Problem.

With the causal chain account now available, return to the lung cancer case. Let “CC_{t1}” be that some causal chain occurs that begins at $t1$ and ends in our agent developing lung cancer. Let “CC_{t2}” be that some causal chain occurs that begins at $t2$ and ends in our agent developing lung cancer. If our agent has background beliefs like ours, she believes that if she develops lung cancer then there are carcinogenic events that form a causal chain between conditions at $t1$ and her developing lung cancer, as well as between conditions at $t2$ and her developing lung cancer. By the causal chain account, she is certain that her chance at $t1$ of developing lung cancer is a hypothetical explanation of CC_{t1}, rather than of her

²⁴ C is a “direct cause” of E exactly if there is no F that is both an effect of C and a cause of E. By a “causal chain”, I mean any ordered sequence of events such that every event in the sequence is a direct cause of the next event (if there is one). For example, if C is a direct cause of E and E is a direct cause of F, there is a causal chain from C to E and C to F.

developing lung cancer. Similarly, she believes that her chance at t_2 of developing lung cancer is a hypothetical explanation of CC_{t_2} , rather than of her developing lung cancer.

What does HCP say about our agent's credence at t_1 that she will develop lung cancer? Nothing directly, because HCP does not apply to our agent with respect to her credence that she will develop lung cancer; there is no X such that she is certain that X is a hypothetical explanation of lung cancer. However, HCP *does* apply to our agent with respect to her credence in CC_{t_1} . According to the causal chain account, she is certain that her chance at t_1 of developing lung cancer is a hypothetical explanation of CC_{t_1} . Therefore, HCP entails that her credence in CC_{t_1} is determined by the chance value she assigns to her developing lung cancer. She believes that she will develop lung cancer if and only if CC_{t_1} is true, and for that reason her credence in lung cancer is equal to her credence in CC_{t_1} . So, if the causal chain account is true, HCP entails SWAMPING in this case: our agent's credence that she develops lung cancer is determined by her opinions about the chance of lung cancer.

What happens at t_2 when our agent learns that her chance of developing lung cancer has increased to .5? Once again, HCP does not apply to our agent with respect to her credence that she develops lung cancer. Furthermore, HCP no longer applies to our agent with respect to her credence in CC_{t_1} . Our agent's chance at t_2 of developing lung cancer is inadmissible_{HCP} with respect to CC_{t_1} (according to our rough understanding of admissible_{HCP} information) because it is evidence about CC_{t_1} but not via an explanatory route. After all, the chance of lung cancer at t_2 clearly does not explain the occurrence of a causal chain that *begins at t_1* and ends in lung cancer. However, HCP does apply to our agent with respect to

her credence in CC_{t2} . According to the causal chain account, our agent is certain that her chance at $t2$ of developing lung cancer is a hypothetical explanation of CC_{t2} . By HCP, her credence in CC_{t2} is determined by her opinions about her chance of developing lung cancer. She believes that she develops lung cancer if and only if CC_{t2} is true and so her credence in lung cancer is equal to her credence in CC_{t2} . Once again, then, HCP implies that SWAMPING is true of this case.

Is the generalization problem thus solved? Nearly. The only remaining difficulty arises if our rational agent believes that conditions at $t1$ and at $t2$ are both direct causes of her developing lung cancer, or if her lung cancer is uncaused. There could be such an agent, but she would *not* believe that her chance of developing lung cancer changes from $t1$ to $t2$.

6. Conclusion: Implications for a Theory of Chance.

Lewis suspected that the PP is explained only if chances somehow reduce to (or, at least, supervene on) the Humean mosaic. He writes, “I think I see, dimly but well enough, how knowledge of frequencies and symmetries and best systems could constrain rational credence. I don't begin to see, for instance, how knowledge that two universals stand in a certain special relation N^* could constrain rational credence about the future instantiation of those universals,” (Lewis 1994, 484).

Lewis's thought, roughly speaking, is that if chances supervene on what actually happens (as far as facts not involving chances are concerned), then it is not entirely mysterious why one's opinions about chances should constrain one's expectations about

what actually happens—including the outcomes of chance processes. But, *contra* Lewis, it is far from obvious that theories according to which chances supervene on the Humean mosaic can provide non-circular explanations of the PP.²⁵

Rather than follow Lewis in arguing that the PP is explained by the metaphysics of objective chance, I have argued that one aspect of the PP (i.e., SWAMPING) follows from a theory of the particular explanatory role that chances play (i.e., the causal chain theory) and a principle governing the relationship between beliefs about scientific explanations and rational expectations (i.e., HCP). One *prima facie* advantage of this approach is that, strictly speaking, it is neutral with respect to what (if anything) chances reduce to or supervene on. That said, my explanation of SWAMPING might nevertheless have some implications for a theory of chance.

The argument from HCP to SWAMPING requires either the truth of the causal chain theory or of a relevantly similar theory according to which chances are hypothetical explanations for the rational agents who believe in them. My strategy for explaining SWAMPING, then, requires that objective chances are the kind of thing that agents can rationally believe to be hypothetical explanations. But theories of chance according to which chances supervene on the Humean mosaic have a *prima facie* difficulty allowing that chances are explanations of (any part of) the Humean mosaic. The worry is that if the Humean mosaic helps to *constitute* the chances, then the chances are ill placed to *explain* the Humean

²⁵ For a detailed discussion of why, see Strevens (1999).

mosaic.²⁶ In contrast, the very theories that Lewis uses the PP to argue against (i.e., theories that deny that chance facts supervene on the Humean mosaic) have no similar difficulty; they allow that chances are responsible for features of the Humean mosaic, but not vice versa. Of course, these are only *prima facie* considerations. But they suggest that the challenge of explaining SWAMPING can be met more easily by theories that do not purport to reduce chances to features of the Humean mosaic.

I have argued that HCP entails SWAMPING given the causal chain theory. That result, of course, does not establish that HCP and the truth of the causal chain theory are what *explain* SWAMPING. SWAMPING might instead be part of an explanation of HCP or of the truth of the causal chain theory. But the analysis of $\text{admissible}_{\text{HCP}}$ information I sketched in Section 4 makes essential reference to hypothetical explanation. This suggests that our rational opinions about hypothetical explanations are at the heart of SWAMPING (and ultimately the PP). Furthermore, it strikes me as less mysterious for beliefs about hypothetical explanations to constrain rational credences than for beliefs about chances to constrain rational credences. For these reasons, I am optimistic that the direction of explanation goes from HCP and the causal chain theory to SWAMPING.

²⁶ There have been recent contributions on both sides of the debate over whether Humean theories do justice to the explanatory power of various nomic features of the world by Cohen and Callender (2009), Loewer (2012), and Lange (2013).

If my explanation of SWAMPING is correct, then we have an explanation of a notoriously puzzling aspect of the PP that is prima facie consistent with a wide variety of theories of objective chance. Theories that deny that chances supervene on the Humean mosaic have no principled difficulty employing my explanation. Far from being the hopeless non-starters that Lewis envisioned, such theories may be able to use my explanation of SWAMPING—and thus to ground Butler’s maxim that chances are the very guide to life.

References

- Butler, Joseph. 1736. *The Analogy of Religion*. 2nd ed. London: Knapton.
- Cohen, Jonathan, and Craig Callender. 2009. "A Better Best System Account of Lawhood." *Philosophical Studies* 145 (1): 1–34.
- Frigg, Roman, and Carl Hoefer. 2010. "Determinism and Chance From a Humean Perspective." In *The Present Situation in the Philosophy of Science*, edited by Friedrich Stadler, Dennis Dieks, Wenceslao González, Hartmann J, Uebel Stephan, Weber Thomas, and Marcel, 351–72. Springer.
- Glymour, Clark N. 1980. *Theory and Evidence*. Princeton, N.J: Princeton University Press.
- Hájek, Alan. 2007. "The Reference Class Problem Is Your Problem Too." *Synthese* 156 (3): 563–85.
- Hall, Ned. 1994. "Correcting the Guide to Objective Chance." *Mind* 103 (412): 505–18.
- Hempel, Carl. 1965. "Aspects of Scientific Explanation." In *Aspects of Scientific Explanation: And Other Essays in the Philosophy of Science*, 331–496. New York: Free Press.
- Hoefer, Carl. 2007. "The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance." *Mind* 116 (463): 549–96.
- Lange, Marc. 2013. "Grounding, Scientific Explanation, and Humean Laws." *Philosophical Studies* 164 (1): 255–61.
- Lewis, David. 1976. "Probabilities of Conditionals and Conditional Probabilities." *The Philosophical Review* 85 (3): 297–315.

- . 1980. “A Subjectivist’s Guide to Objective Chance.” In *Studies in Inductive Logic and Probability*, edited by Richard C. Jeffrey, 2:83–132. University of California Press.
- . 1994. “Humean Supervenience Debugged.” *Mind* 103 (412): 473–90.
- Loewer, Barry. 2004. “David Lewis’s Humean Theory of Objective Chance.” *Philosophy of Science* 71 (5): 1115–25.
- . 2012. “Two Accounts of Laws and Time.” *Philosophical Studies* 160 (1): 115–37.
- Meacham, Christopher J. G. 2010. “Two Mistakes Regarding the Principal Principle.” *The British Journal for the Philosophy of Science* 61 (2): 407–31.
- Nelson, Kevin. 2009. “On Background: Using Two-Argument Chance.” *Synthese* 166 (1): 165–86.
- Railton, Peter. 1978. “A Deductive-Nomological Model of Probabilistic Explanation.” *Philosophy of Science* 45 (2): 206–26.
- Salmon, Wesley C. 1967. *The Foundations Of Scientific Inference*. Pittsburgh: University of Pittsburgh Press.
- . 1971. *Statistical Explanation & Statistical Relevance*. Pittsburgh: University of Pittsburgh Press.
- . 1989. *Four Decades of Scientific Explanation*. Minneapolis: University of Minnesota Press.
- Schwarz, Wolfgang. 2014. “Proving the Principal Principle.” In *Chance and Temporal Asymmetry*, edited by Alastair Wilson, 81–99. Oxford University Press.

Strevens, Michael. 1999. "Objective Probability as a Guide to the World." *Philosophical Studies* 95 (3): 243–75.

Thau, Michael. 1994. "Undermining and Admissibility." *Mind* 103 (412): 491–504.