The Applicability Problem for Chance Explanation

Katrina Elliott
UCLA

Abstract: I argue it is not accident that we can both predict and explain the occurrence of (at least some) chance outcomes. For good reason, influential theories about the nature of chance explanation imply that our beliefs about the chance of an event should determine our degree of confidence in that event. Such theories face a problem that has so far gone unnoticed: so long as we are not justified in believing that beliefs about chances should determine degrees of confidence, we are not justified in believing that our best probabilistic theories provide explanations of any events that occur by chance. Adapting terminology from Salmon (1967), I call this difficulty the "applicability problem". I argue that any otherwise plausible theory of chance explanation faces a dilemma: either it is not a genuine theory of chance explanation or it faces the applicability problem. I conclude by sketching a novel approach to solving the applicability problem: perhaps the relationship between chance and degree of confidence is grounded by chance's explanatory role.

0. Introduction

Every morning before work, I watch the local weather report in the hope of finding out whether it will rain. The meteorologist typically doesn't say that it will rain or that it won't. Instead, the meteorologist typically reports the *chance* of rain. Still, learning the chance of rain is useful to me. If I find out that the chance of rain is high, for example, I expect that it will rain and I bring an umbrella. So, even if the weather report doesn't settle whether it will rain, it can inform my expectations about rain.

Sometimes I find myself in a more reflective mood when I'm watching the weather report. I wonder not only if it will rain, but also why it rains if it does. The meteorologist displays a surface weather map and points out parts of the map that represent features of the environment that are relevant to whether it will rain today. It is not obvious that actual weather reports contain enough information to count as genuine scientific explanations of various meteorological phenomena, but many philosophers will agree that there could be a weather report that would explain why it rains if it does. However, on days when rain is a matter of chance, not even an ideal explanation of rain contains information that guarantees that it will rain. So, even if the weather report doesn't settle whether it will rain, it can explain why it rains if it does.

I aim to argue it is no accident that, e.g., weather reports can both explain why it rains and inform our expectations about whether it rains. For events that occur by chance to have explanations, there must be some connection between information that explains chance events and the occurrence of those chance events. I'll argue that, for good reason, many influential theories about the nature of that connection imply that our beliefs about the chance of an event should determine our expectations about whether that event occurs.

Here's why that matters. Philosophers have been searching for a metaphysical theory of chance that reveals why beliefs about chances should determine our expectations. Their hope is that, once we understand what chances are, it will become clear why beliefs about that sort of thing should determine our expectations. However, no theory of chance is widely agreed to have met this burden, the most promising attempts so far apply to only a few (and independently controversial) theories of chance (e.g., Loewer 2004, Frigg and Hoefer 2010, Schwarz 2014), and some philosophers have argued that there can be no such explanation (e.g., Strevens 1999). Adapting terminology from Wesley Salmon (1967), I call the problem of showing why it is that beliefs about chances should determine our expectations the "applicability problem". To the extent that there are explanations of chance events only if beliefs about chances should determine our expectations, theories of explanation face a problem that has so far gone unnoticed: so long as the applicability problem is unsolved, we are not justified in believing that there are explanations of events that occur by chance. In other words, theories of chance explanation inherit the applicability problem.

I'll proceed as follows. In section 1, I introduce the notion of a "chance explanation" and show why some models of explanation, such as Carl Hempel's deductive-nomological model, cannot accommodate chance explanations.[1] In section 2, I use Hempel's model of chance explanation (the "inductive-statistical" model) to illustrate one way in which a theory of chance explanation might require the correctness of some "principle of probability coordination": a principle according to which beliefs about chances should determine our confidence about the future. In section 3, I discuss why it is problematic for a theory of chance explanation to require that there is a correct principle of probability coordination (i.e., why the applicability problem is a problem). In section 4, I show that the applicability problem faces three more theories of chance explanation: Wesley Salmon's statistical-relevance model, James Woodward's manipulability model, and Peter Railton's deductive-nomological-probablistic model.

[1] I use Hempel's models of scientific explanation for illustrative purposes because they are both admirably clear and hugely influential. The actual targets of this essay are theories of scientific explanation that, unlike Hempel's, have not been unanimously rejected.

Generalizing from that discussion, I claim that any otherwise plausible theory of chance explanation faces a dilemma: either it is not a genuine theory of chance explanation or it requires that there is a correct principle of probability coordination (and so inherits the applicability problem). Though my primary goal is merely to argue that theories of chance explanation inherit the applicability problem, I conclude in section 5 by considering three possible responses to that problem. Two of these are orthodox responses (albeit implicitly, since the ubiquity of the applicability problem for theories of chance explanation has not yet been recognized): deny that there are chance explanations or ignore the applicability problem when formulating theories of chance explanation (as we often ignore other foundational problems, such as the problem of induction). After evaluating these two responses, I conclude by considering an unorthodox third response. On the view I explore, the recognition that theories of chance explanation require the correctness of some principle of probability coordination lights the way to a novel strategy for solving the applicability problem: perhaps probability coordination is not grounded merely by what chances are (i.e., by the metaphysical nature of chance), but rather by the role that chances play in chance explanations.

1. Chance Explanation and the Deductive-Nomological Model

My argument that various theories of chance explanation inherit the applicability problem requires some substantive assumptions about the nature of chance. First, I'll assume that there are such things as chances: objective features of the world that obey the standard axioms of probability, that (at least sometimes) take values between 0 and 1, and that are (arguably) referred to by mature scientific theories.[2] Philosophers disagree about which (if any) scientific theories are best interpreted as modeling chances but, partly for the sake of having a familiar toy example to work with, I'll assume that the probabilities in weather reports model chances.[3] Our intuitions about probabilistic explanations are most clear when we consider explanations provided by commonsense and non-fundamental physical theories, since it is these explanations with which we are most familiar. That said, my appeal to such non-fundamental chances is not merely an exegetical convenience; my own view is that at least some commonsense and non-fundamental theories do model chances and that these theories provide us with genuine explanations of events that occur by chance. Defending that view, however, is more than I'll do here. All that is essential to the discussion that follows is that there are *some* cases in which events occur by chance. Readers who believe these cases are best drawn from sciences other than meteorology may safely replace my cases with their own.

Second, I assume that at least some objective chances are "single-case" chances: chances of a *single event* occurring. For example, I assume that the meteorologist is not merely reporting on the chance of rain among a collection of (actual or hypothetical) days that are relevantly similar to today. Rather, the meteorologist is reporting the chance that it rains on *this particular* day. Some philosophers are skeptical of single-case chances

[2] This characterization of objective chance is insufficient to distinguish it from other arguably distinct features of the world, such as actual relative frequency, but is all that can be said about the nature of chance without courting controversy—other than that chance obeys a principle of probability coordination.

[3] The probabilities that appear in the meteorologist's report might instead be, in some sense, *subjective*. Such subjective probabilities might refer to the meteorologist's actual expectation that it will rain, or to the expectations that any agent should have given the evidence available to the meteorologist, or to some other sort of thing that is essentially relative to a particular subject or body of evidence.

(e.g., von Mises 1957, Howson and Urbach 1993, Gillies 2000) and single-case chances open the door to further difficulties such as the reference class problem (discussed in, e.g., Ayer 1963, Hajek 2007), but I'll ignore these complications in what follows.

Finally, it will be important to distinguish between an explanation of an event that occurs by chance, which I'll call a "chance explanation", and an explanation of an event's having a given chance of occurring.[4] To illustrate the difference, consider Carl Hempel's deductive-nomological (D-N) model of scientific explanation. According to the D-N model, we have an explanation if we have statements of particular facts and statements of general laws that, when conjoined, deductively entail that the event to be explained occurs. To use a standard example, the current position of the celestial bodies and the laws of planetary motion combine to explain the next lunar eclipse.

Let's try to apply the D-N model to an explanation of an event that occurs by chance. Suppose it is a general law that there is a 90% chance of evening rain on any day in which morning weather condition *C* obtains. Suppose further that condition *C* obtained this morning and that it rained this evening. Inspired by the D-N model, we might offer the following explanation of this evening's rain:

P1. On any day in which condition *C* obtains, the chance of evening rain is .9.

P2. Today condition *C* obtained.

So,

C. Today's chance of evening rain was .9.

---

[4] This choice of terminology is simply a stipulation, and may not correspond to what anyone normally means by "chance explanation".

Notice that the conclusion of this argument is not that a chance event occurs (e.g., that it rained), but is rather a specification of the chance that an event occurs (e.g., that the chance of rain was .9). If the explanation of an event must entail the occurrence of that event, then P1 and P2 do not explain why it rained. Chance explanations do not fit the D-N model.

Is there some reason to think that P1 and P2 nevertheless explain why it rained? According to Hempel, an argument that satisfies the D-N model is explanatory because, "…the argument shows that, given the particular circumstances and the laws in question, the occurrence of the phenomenon *was to be expected;* and it is in this sense that the explanation enables us to *understand why* the phenomenon occurred." (1965, pg. 337) Valid arguments are not the only kind of argument that shows that a given phenomenon was to be expected—inductive arguments do that too.

2. Probability Coordination and the Inductive-Statistical Model

Unlike a deductively valid argument, an inductive argument can be better or worse depending on its "strength". Measures of inductive strength, like chances, take real values along the open unit interval (where deductive validity is a limiting case). The strength of an inductive argument with complex premises can be hard to intuit, but we have no problem judging the relative strength of simple inductive arguments. For example, the argument that I will get hungry tomorrow since I have gotten hungry every day so far is stronger than is the argument that I will have sushi for lunch tomorrow since I sometimes eat sushi.

There are few propositions about which we are certain, and our uncertainty comes in degrees. For example, I'm certain that 2+2=4, I'm confident (but less than certain) that smoking is unhealthy, and I'm skeptical (but nevertheless open to the possibility) that keeping a pet is morally wrong. I take on the Bayesian tradition of modeling these degrees of expectation, or "degrees of confidence", as probabilities and I assume that an agent's degrees of confidence are rational only if they satisfy the standard Kolmogorov axioms of probability (Kolmogorov, 1933). Hempel (see, for example, Hempel's 1965, pg. 397-399), following Carnap (1950), saw an important connection between the inductive strength of an argument and our degree of confidence in its conclusion.[5] If I am certain of the premises of an inductive argument and I have no other information that would change the argument's inductive strength were it included as a premise, then my degree of confidence in the conclusion should equal the inductive strength of the argument.[6] Since a strong inductive argument can give us reason to expect that its conclusion is true, Hempel introduces the Inductive-Statistical (I-S) model to allow that some strong inductive arguments are explanations of their conclusions.

Consider the following inductive argument that it will rain this evening:

P1. On any day in which condition *C* obtains, the chance of evening rain is .9.

P2. Today condition *C* obtained.

So,

C*. It rained this evening.

---

[5] Neither author uses the phrase "degree of confidence". Instead, Hempel and Carnap talk of "applying" inductive logic to a given "knowledge situation".

[6] The antecedent of this conditional is an informal version of the requirement of total evidence.

What is the inductive strength of this argument? Given the connection between inductive strength and reasonable degree of confidence, the correct answer depends on how confident an agent should be that C* is true given that she is certain of P1 and P2 and has no other information that is relevant to whether it will rain.

Following Michael Strevens (1999), I use the term "principle of probability coordination" to refer to a rule that specifies how confident an agent should be in an outcome given her opinions about that outcome's chance. There is disagreement over what exactly is the most intuitive and philosophically useful principle of probability coordination (e.g. Lewis 1994, Hall 1994, Nelson 2009, Meachum 2010), but most everyone shares the strong intuition that if an agent is certain that some outcome has a particular chance of occurring, then (ignoring extraordinary cases involving, e.g., reliable crystal balls) her degree of confidence in that outcome should equal the chance of that outcome (on pain of irrationality). The information about chances in a weather report justifiably informs my expectations only if there is some correct principle of probability coordination. Without such a principle, knowledge of chances is not *applicable* to our degrees of confidence about future chance occurrences. [7]

P1 and P2 imply that the chance of rain this evening is .9. The intuitively correct principle of probability coordination implies that if an agent is certain that the chance of rain this evening is .9, then her degree of confidence that it will rain this evening is also

[7] Here is another way to see the point. The Kolmogorov axioms combined with Bayesian updating find no conflict between, say, being certain that the chance of rain is high and being skeptical that it will rain. Chances get no special purchase on degrees of confidence in a Bayesian framework unless we count as rational only those initial credence functions that obey a principle of probability coordination.

.9. So, thanks (in part) to the principle of probability coordination, the inductive strength of our argument from P1 and P2 to C* is .9.

If strong inductive arguments like the one above (in which an argument's premises specify the chance value of an event) are explanations of chance events, then there are such chance explanations only if there is a correct principle of probability coordination. For an argument with premises that specify an event's chance to count as strong, there must be some bridge between chance and degree of inductive support. If there is a bridge between chance and degree of inductive support, then there is a bridge between chance and rational degree of confidence. So, if there is no correct principle of probability coordination (or if we are not justified in believing that any such principle is correct) and so no bridge between chance and rational degree of confidence, then inductive arguments like the one above are not strong (because they have no strength) and so are not chance explanations (or we are not justified in believing that they are chance explanations).

3. The Applicability Problem

Traditionally, the problem of justifying some principle of coordination is posed not to theories of explanation, but rather to interpretations of the probability calculus. Salmon, for example, argues that any adequate explication of the concept of probability must meet the "criterion of applicability" by providing an explication of the concept of probability such that we can see why knowledge of that sort of thing should determine our degrees of confidence (1967, pp.64-65). David Lewis echoes this constraint as applied to objective interpretations of probability when he writes, "Don't call any alleged feature of reality 'chance' unless you've already shown that you have something, knowledge of which could constrain rational [degrees of confidence]." (1994, pg. 484) And Alan Hájek makes a similar point when he criticizes "no-theory" theories of chance, which "leave quite obscure why probability should function as a guide to life." (2007, pg. 563)

Unfortunately, satisfying the applicability criterion has proved to be a problem (hence "the applicability problem"). It is beyond the scope of this paper to discuss why each well-known theory of chance founders on the applicability problem, but it is worth taking a moment to illustrate one sort of difficulty these theories face. Consider the view that the chance of a given type of outcome is simply the actual frequency of that type of outcome among the total number of actual outcomes. On this view, for example, if a coin actually lands heads half of the times it is flipped (throughout its entire existence), then the coin is fair. Let's call this theory "ARF" for "actual relative frequency".

Overall, ARF is a wildly implausible theory of chance (since it implies that e.g., a fair coin cannot, on pain of contradiction, be tossed an odd number of times), but it should at least show promise of solving the applicability problem. According to ARF, if we know the chance of an outcome then we know how frequently that outcome actually occurs. How hard could it be to argue that an agent's degree of confidence that a given coin lands heads should be equal to the frequency with which it actually lands heads?

As it turns out, pretty hard. Suppose an agent is certain that her coin is fair. The intuitive principle of probability coordination implies that her degree of confidence that the coin will land heads should be .5. How might we argue that our agent should do what the intuitive principle requires? We could try arguing that she will "do better" if she obeys the principle than if she does not, in the sense that she will score higher according to some scoring rule that rewards agents for having high expectations in outcomes that occur and low expectations in outcomes that do not occur.[8] That seems like a promising start, especially because ARF rules out deviant series of outcomes, such as the coin landing heads on every toss (since ARF implies that half the flips of a fair coin land heads). Still, we cannot argue that an agent who follows the principle of probability coordination is *guaranteed* to do better than one who does not. For example, our agent might only be around to experience the half of the total coin flips that land heads. We might stipulate that our agent experiences *all* of the coin flips, but even then she might still do better by disobeying the principle. Suppose every odd numbered toss lands heads and every even numbered toss lands tails (so that the first toss is heads, the second tails, and so on…). In that case, the *best* strategy (i.e., the strategy on which the agent receives the highest possible score) would have been to disobey the intuitive principle of probability coordination by alternating between certainty that the coin lands heads and certainty that the coin lands tails. So, even if ARF is true, there can be no guarantee that an agent who follows the intuitive principle of probability coordination will do better than one who does not.

That said, while it is *possible* for an agent to do better by violating the principle, an agent who sets her degrees of confidence in accord with actual relative frequencies over a sufficiently long series of tosses is *likely* to do better than one who does not. In

[8] Never mind wondering what justifies this scoring rule over any another—the argument fails even if we grant the scoring rule. See Winkler (1996) for a discussion of scoring rules.

other words, we can show that if chances are actual relative frequencies then our agent's best chance of doing well is to follow the intuitive principle of probability coordination. But why should our agent adopt a strategy that gives her the best *chance* of doing well? This question is simply a higher-order version of the question with which we began. We set out to justify an inference from a chance to degree of confidence but we have justified only an inference from one chance to another, and so we are no closer to having justified some principle of probability coordination.

Things only get worse when we acknowledge that, *pace* the implausible theory ARF, any (non-extremal) chance value is consistent with any actual relative frequency whatsoever. For example, a fair coin might only ever land heads, or only ever land tails, or any ratio (of heads to tails) in between.[9] While most everyone acknowledges that chance values are not guaranteed to match *actual* relative frequencies, some philosophers have argued that the chance of an outcome corresponds to the *hypothetical* frequency of that outcome among a potentially infinite series of trials. I'm not sure; if the chance of an outcome is less than 1, I'm not sure that the outcome *would* occur even once given a hypothetical infinite series of trials. (If a fair coin might land tails on the first hypothetical flip, for example, why does it have to land heads on any of the other hypothetical flips?) In fact, I am not sure there are *any* true would-counterfactuals about the occurrence of chance outcomes. Even the most likely outcome might not occur, and

[9] According to a Lewisian about chance, the chance value of, e.g., rain *can* put some limits on the range of possible actual frequencies of e.g., rain. But even the Lewisian tolerates sometimes wide divergence between chance values and actual relative frequencies. (Lewis, 1994)

so there seems to be no hypothetical supposition under which it would occur. [10] At any rate, it's hard to see how an appeal to hypothetical frequencies can be of any help in justifying some principle of probability coordination, given that actual relative frequencies were not able to do the job. Hypothetical frequencies only take us further away from how actual agents fare by following some principle of probability coordination.

Of course, there are other strategies for justifying probability coordination and there are more sophisticated theories of chance. But similar problems seem to face all attempts to solve the applicability problem.[11] The general difficulty is that there seems to be no relevant non-probabilistic relation between the chance of an outcome and the occurrence of that outcome (other than probability coordination itself) in which to ground probability coordination.

That the applicability problem is a *problem* is well known to philosophers who work on theories of chance. However, that the applicability problem threatens theories of chance explanation has not yet been appreciated. If the applicability problem is unsolved, we do not merely lack an adequate metaphysics of objective chance. Instead, we face a more central problem to our understanding of scientific inquiry: we are not justified in believing that our best probabilistic scientific theories (of which there are many) provide explanations of events that occur by chance.

---

[10] These issues are closely related to well-known objections to theories of chance that attempt to reduce chances to hypothetical frequencies. The problem is that the chance of an event seems to guarantee very little about the actual or hypothetical frequency of that event. For an excellent and thorough discussion in which this point is convincingly argued for, see Hajek's (1996) and (2009).

[11] For a detailed discussion of the ways that attempts to solve the applicability problem founder, see Strevens 1999.

4. The Applicability Problem for Chance Explanation

Let's take stock. We have seen that, because Hempel's theory of explanation models chance explanations as inductive arguments, it requires that there is a correct principle of probability coordination. That is a problematic consequence, I claimed, because no principle of probability coordination has yet been justified. So much the worse, one might respond, for Hempel's theory of explanation. There are many reasons to deny that explanations are strong inductive arguments (e.g., Scriven 1959, Jeffrey 1969), so why should we care that the applicability problem gives us one more? We shouldn't, but the applicability problem is not limited to views on which chance explanations are arguments.

The reason that many theories of chance face the applicability problem is closely related to the problem faced by attempts to justify probability coordination. Just as it is hard to find suitable relations between chances and outcomes— other than probability coordination— in which to ground probability coordination, so too it is hard to find explanatory relations between chances and outcomes other than probability coordination. To demonstrate the ubiquity of the applicability problem, I show that it applies to each of a diverse sampling of theories of chance explanation.[12] First, I consider Salmon's model of chance explanation, on which explanations give us the best possible grounds for our degrees of confidence. It is easy to see that Salmon's model requires that there is a

[12] Each of the following three theories (rightly, on my view) allows that both likely and unlikely events have explanations. Some people find it more intuitive that likely events are explicable than that unlikely events are explicable, but each of these theories would still face the applicability problem if they were modified so as to apply only to likely events. That said, the strategy for solving the applicability problem that I explore in section five does require that unlikely events have explanations.

correct principle of probability coordination once we understand the role that chances play in that model. Next, I consider James Woodward's model of chance explanation, which does not explicitly link explanation with confidence but does frame explanation in terms that evoke subjects. Explanations, on Woodward's view, give us information that is relevant to controlling, changing, and manipulating our environment. I'll argue that it is difficult to see how explanations of events that occur by chance could ever provide such information, unless—thanks to a principle of probability coordination—chance explanations are relevant to the practical utility of controlling, changing, and manipulating our environment. Finally, I'll consider Peter Railton's model of chance explanation, which aims to make "explanation" a purely ontic notion: the ideal explanation of an event contains all information (notably including information about underlying physical mechanisms) that is causally and nomologically relevant to that event. I argue that Railton's model fails as a purely ontic theory of chance explanation: either the model is not a theory of chance explanation or it requires the truth of a principle of probability coordination.

4.1 The Statistical Relevance Model

In explaining his statistical relevance (SR) model of explanation, Salmon writes:

> To explain an event is to provide the best possible grounds we could have had for making predictions concerning it. An explanation does not show that the event was to be expected; it shows what sorts of expectations would have been reasonable and under what circumstances it was to be expected. To explain an event is to show to what degree it was to be expected… (1971, pg.79)

On Salmon's view, chance explanations are neither inductively strong arguments nor arguments of any kind. Nevertheless, as the quoted passage makes clear, Salmon believes that it is crucial to their explanatory power that chance explanations should determine our degrees of confidence. On the SR model, explanatorily relevant factors are *statistically relevant* factors: factors that make a difference to the probability of the event to be explained. An explanation consists of three ingredients: the prior probability of the event to be explained (i.e., the probability of the event *prior* to taking into account statistically relevant factors), the posterior probabilities of the event to be explained (i.e., roughly, the probabilities of the event given each possible combination of statistically relevant factors), and a description of which statistically relevant factors are present in the case at hand. For example, making the simplifying assumption that only cold fronts are statistically relevant to rain, an explanation of today's rain consists in the prior probability of rain, the probability of rain conditional on there being a cold front, the probability of rain conditional on there not being a cold front, and a statement that today we are (or we are not) experiencing a cold front.[13]

So, on Salmon's view, the explanation of an event should determine our degree of confidence that the event occurs. And, on Salmon's view, an explanation of an event contains every posterior probability for that event as well as a statement of which posterior probability applies to the actual situation. These two elements combine to imply that there is a correct principle of probability coordination. In any case in which the event to be explained occurs by chance, the posterior probability that applies to the

[13] For all I've said so far, the SR model obviously suffers from counterexamples. Barometer readings, for example, are statistically relevant to whether it rains, but the statistical relationship between barometer readings and rain has no place in an explanation of why it rains. Further details of the SR model aim to defang this objection, but discussing these details takes us too far afield from our main topic. Salmon (1997) ultimately concedes that the SR model should be supplemented so that it is sensitive to causal relations.

actual situation corresponds to the chance of that event. For example, if there is a cold front today then the probability that it will rain conditional on there being a cold front is the chance that it will rain today (assuming, once again, that only cold fronts are relevant to rain). The prior probability and the remaining posterior probabilities of the event clearly should not determine our degree of confidence in the event, since each of them fails to take into account at least one statistically relevant factor that actually obtains. So, if an SR explanation of a chance event should determine our degree of confidence in that event, then the posterior probability that corresponds to the chance of that event should determine our degree of confidence in that event. According to the SR model, then, there are chance explanations only if there is a correct principle of probability coordination. Without such a principle, chance explanations would not provide the "best possible grounds we could have had" for our degrees of confidence.

4.2 The Manipulability Model

Unlike Hempel and Salmon, James Woodward offers a model of explanation that, at first glance, seems to have nothing at all to do with our particular degrees of confidence in chance outcomes. On Woodward's view, causal claims give us information that is relevant to manipulating, controlling, and changing our environment and, according to Woodward, "we are in a position to explain when we have information that is relevant to manipulating, controlling, or changing nature." (2003, pg. 11) For Woodward, a sentence such as "cold fronts cause rain and we are experiencing a cold front" counts as a minimally adequate chance explanation of today's rain. Plainly, that explanation of today's rain no more determines what should be our (precise) degree of confidence that it rains than does the information that it might rain today.

Nevertheless, I will argue that there are chance explanations on Woodward's view only if there is a correct principle of probability coordination. The manipulability model is built from Woodward's theory of causation. Woodward employs formalisms involving directed graphs and structural equations to make this theory precise, but the basic idea can be formulated using only his notion of a "variable". In the simplest case, a variable represents a property, and the value of the variable (either 1 or 0), represents whether or not that property is instantiated.[14] On Woodward's view, "the claim that X causes Y means that for at least some [objects], there is a possible manipulation of some value of [the variable] X that they possess, which, given other appropriate conditions (perhaps including manipulations that fix other variables distinct from X at certain values), will change [a] the value of [the variable] Y or [b] *the probability distribution of Y* for those [objects]." (2003, pg.40, with my labels "a" and "b".)

To illustrate, let's first apply Woodward's view to a deterministic case. In deterministic contexts, only condition [a] of Woodward's theory is relevant to understanding the meaning of causal claims. Suppose that overeating is a deterministic cause of weight gain. We schematize "overeating causes weight gain" by introducing the variable X to represent the property of overeating and the variable Y to represent the property of having gained weight. Consider a group of individuals who are not overeating and who are not gaining weight. For each of them, the value of X and Y is 0. Now, suppose we were to intervene by changing the value of X from 0 to 1 (i.e., by

[14] It is important for Woodward's overall project that variables can correspond to magnitudes and can take many values, but that detail will not be relevant to what follows.

causing each individual to overeat). If X causes Y in a deterministic case, then (assuming some additional conditions are met that need not concern us here) by condition [a], the value of Y for (at least some of) these individuals changes to 1. Translating out of the formalism, the claim that overeating causes weight gain means (roughly) that, for some individuals, were there an appropriate intervention that changes whether or not they overeat then there would be a corresponding change in whether or not they gain weight.

So much for the deterministic case— since we are interested in chance explanations, we are interested in how Woodward's theory of causation works in indeterministic cases. Woodward's treatment of causation in indeterministic contexts relies on condition [b], and so is importantly different from his treatment of causation in deterministic contexts. If causal claims were semantically linked only to counterfactuals about changes that would occur under certain interventions (i.e., if satisfying [a] were necessary for the truth of a causal claim), causal claims would (arguably) always imply deterministic contexts. After all, if the relationship between X and Y is indeterministic, then it is (arguably) never the case that the value of Y *would* change under a suitable intervention on the value of X. In indeterministic contexts, the only relevant would-counterfactauls seem to be about what *the probability distribution* on Y would be given an intervention on X. But, instead of denying that there are true causal claims in indeterministic contexts or arguing that there are true would-counterfactuals in indeterminsitic contexts, Woodward includes condition [b] to allow that if the probability distribution on Y were to change under a suitable intervention on X, then X causes Y.[15]

[15] According to some philosophers, claims of the form "if P had been the case, then Q would have been the case" are true if Q is true. In contrast, Woodward (rightly) doubts

The upshot is that, for Woodward, the fact that X causes Y does not guarantee that intervening on X would change the value of Y, since [b] allows that X causes Y when changing X might leave Y unchanged (so long as the probability distribution on Y is changed).

If whether it rains is a matter of chance and cold fronts cause rain, it follows from Woodward's theory of causation that an intervention that changes whether or not there is a cold front changes the *chance* of rain. It is reasonably intuitive that the fact that cold fronts cause rain is relevant to manipulating, controlling, and changing the chance of rain.[16] So, granting Woodward's theory of explanation, it is reasonably intuitive that the fact that cold fronts cause rain explains the chance of rain. But, why is manipulating, controlling, and changing the chance of rain relevant to manipulating, controlling, and changing whether it rains?

One might be tempted to answer that control over the chance of rain allows us to control whether it rains by allowing us to control the actual frequency with which it rains. But, as we saw in section 3, that is not true. What is true is that the more likely rain is the higher the actual frequency of rain is likely to be. But again it is not at all obvious that controlling the *chance* of the frequency of rain allows us to control the frequency of rain. One might instead answer that control over the chance of rain provides us with control over whether it rains by allowing us to control what the frequency of rain would be

---

that all counterfactuals with true consequents are true in every indeterministic context. (2003, pg. 214)

[16] Since we do not have control over cold fronts, talk of controlling rain via controlling cold fronts might be confusing. But were we to have control over cold fronts, then we would have control over the chance of rain. For Woodward, our interest in causal claims is a generalization of our interest in situations over which we actually have control. (2003, chapters 1 and 3)

among a hypothetically infinite series of days. Even if control over the chance of rain does imply control over the hypothetical frequency of rain (which I doubt: see section 3), control over the hypothetical frequency of rain does not imply control over whether it actually rains.

If it is a matter of chance whether an event occurs, it seems that we can only manipulate, control or change the chance of that event and then hope for the best. The relation "is information that is relevant to manipulation, control, and change" may be an otherwise plausible candidate for being a relation that explanations stands in to that which they explain, but it is not at all clear that it is a relation that chance events can stand in to chance explanations.[17] So, if the manipulability model is to count as a model of chance explanation, Woodward owes us an account of why information that is relevant to manipulating, controlling, and changing *the chance* of an outcome is relevant to

[17] Woodward aims to address this sort of worry when he argues that type-causal generalizations, such as "latent syphilis causes paresis", are explanatory because they imply that e.g., "in the right circumstances, intervening to cause someone to have latent syphilis (i.e., manipulating the situation from one in which no latent syphilis is present to one in which latent syphilis is present will *sometimes* change whether paresis occurs; in particular, it will change the situation from one in which the probability of paresis is 0 (because there is no syphilis) to one in which that probability is greater than 0." (2003, pg. 214) I agree with Woodward that, on his view, the causal generalization "latent syphilis causes paresis" implies that changes in whether a patient has syphilis also change the patient's chance of developing paresis, but it does not follow that there are (actual or possible) circumstances in which changing the chance of syphilis from 0 to any value less than 1 will even sometimes change whether paresis occurs. Since having latent syphilis merely increases the chance that one will have paresis, it is possible (though overwhelmingly unlikely) for latent syphilis to *never* develop into paresis. Woodward's retreat in this passage from a claim about what will occur to a claim about what is likely to occur illustrates the difficulty one faces when trying to develop a theory of chance explanation that makes no appeal to probability coordination. There are (plausibly) explanatory relations that antecedent conditions stand in to chances (such as Woodward's brand of counterfactual dependence), but it is hard to see why any of those relations imply that there is an explanatory connection between those antecedent conditions and the occurrence of chance events.

manipulating, controlling, and changing *that outcome*. As it turns out, Woodward has the resources for just such an account, but only if there is a correct principle of probability coordination.

One of Woodward's primary motivations for his theory of causation is that it "makes understandable how knowledge of causal relationships has any practical utility at all." (2003, pp. 30-31) Perhaps, then, the reason that information about manipulating, controlling, and changing the chance of an outcome counts as information that is relevant to manipulating, controlling, and changing whether that outcome occurs is that the former has practical import for those who care about the latter. The most obvious way for information about chances to have practical import is via an expected utility function. Changes in the actual and counterfactual chance of rain, for example, result in changes in the expected utility of causing cold fronts, which in turn (arguably) determines whether one should (on pain of practical irrationality) cause cold fronts. However, the move from actual and counterfactual chances to expected utility functions requires an inference from chance values to degrees of confidence.[18] For example, the reason I shouldn't ride my bike to work when I know that the chance of rain is high (given that I hate getting rained on and I don't mind driving) is that I should expect that it will rain and so should expect to have a worse time if I ride my bike than if I drive. But, as is by now familiar, this

[18] Perhaps this point is more obvious when we have in mind evidential decision theory, rather than causal decision theory—especially since Lewis's influential version of causal decision theory is formulated in terms of chances rather than degrees of confidence. Nevertheless, the justification of this formulation requires an application of Lewis's version of probability coordination (i.e., the Principal Principle). (See, e.g., Lewis 1981, pg. 27 fn. 24) I have no decisive argument to give that there is *no* other way for knowledge of chances to have practical import for those who care about chance outcomes, but the most familiar transitions from chance values to claims about practical rationality go by way of probability coordination.

inference from chances to rational degrees of confidence requires that there is a correct principle of probability coordination.

4.3 The Deductive-Nomological-Probabilistic Model

The deductive-nomological-probabilistic (D-N-P) model of explanation is Peter Railton's attempt to improve on Hempel's theory of explanation in two ways: by allowing that there are explanations of even unlikely chance outcomes and by requiring that chance explanations illuminate the mechanisms that are responsible for chance processes. Recall our earlier deductively valid argument that the chance of rain is .9:

P1. On any day in which condition *C* obtains, the chance of evening rain is .9.

P2. Today condition *C* obtained.

So,

C. Today's chance of evening rain was .9.

On Railtons' view, P1 and P2 are insufficient to explain C because P1 and P2 do not illuminate the mechanisms responsible for C. He writes,

> The goal of understanding the world is a theoretical goal, and if the world is a machine—a vast arrangement of nomic connections—then our theory ought to give us some insight into the structure and workings of the mechanism…Knowing enough to subsume an event under the right kind of laws is not, therefore, tantamount to knowing the *how* or the *why* of it." (1978, pg. 208)

If we are to have a genuine explanation of the chance of rain, we must, on Railton's view, supplement our derivation of the chance of rain with a further derivation of P1 from an underlying fundamental physical theory. Of course, that derivation will be incredibly complex and is not something we have the ability to produce. But, for Railton, the ideal of explanation is to provide derivations from fundamental theories because these derivations illuminate the mechanisms responsible for the higher-order laws that subsume the events to be explained.

So, an explanation of rain includes a derivation of C from P1 and P2, as well as a derivation of P1 from fundamental physics. Can these two derivations explain why it rains? Yes, claims Railton, so long as we add a parenthetical addendum stating whether or not it rained. Railton rightly notes that, without such an addendum, these two derivations would merely be an explanation of the chance of rain. He writes, "Dropping off the addendum leaves an explanation, but it is a D-N explanation of the occurrence of a particular probability, not a probabilistic explanation of the occurrence of a [chance outcome]." (1978, pg. 217) However, how can adding Railton's addendum transform an explanation of the chance of rain into an explanation of rain? The addendum neither plays a role in the derivation nor illuminates any underlying mechanisms; it simply states that an event occurs (or fails to occur). It is hard to see how the statement that an event occurs (or fails to occur) can play *any* role in explaining that event's occurrence (or failure to occur).

Why, then, does Railton consider the D-N-P model to be a model of chance explanation? Addressing this very question, Railton writes:

> Still, does a [D-N-P explanation of an unlikely event] explain why the [event] took place? It does not explain why the [event] *had to* take place, nor does it explain why the [event] *could be expected to* take place. And a good thing, too: there is no *had to* or *could be expected to* about the [event] to explain—it is not only a chance event, but a very improbable one. [The D-N-P

> explanation] does explain why the [event] *improbably* took place, which is how it did. (1978, pg. 216)

So, the D-N-P explains why an unlikely event occurred by explaining why it "improbably took place". If "improbably" refers to the chance of the event's occurrence, it seems as if Railton is simply asserting that an explanation of an event's chance is an explanation of that event's occurrence. That may be true, but a theory of explanation must do more than stipulate that one thing explains another. If "improbably" refers to the degree of confidence an agent should have that the event occurs, then Railton's view is that the explanation of an event must also explain why our degree of confidence in that outcome should be one value rather than another. Fair enough, but if a derivation of the chance of an event explains (or determines) what should be our degree of confidence in that event, then there is a correct principle of probability coordination. I conclude that either the D-N-P model is not a model of chance explanation or it requires that there is some correct principle of probability coordination.

4.4 Generalizing the Argument

I've argued that, on three prominent theories of chance explanation (or four, counting Hempel's), there are chance explanations only if there is a correct principle of probability coordination. For each theory, my argument is an instance of a more general argument schema. Let "O" be a given proposition that a particular chance event occurred, such as the proposition that it rained this evening. Let "E" be whatever explains O according to the theory of chance explanation in question. One task for any theory of chance explanation is to identify at least one relation, call it "R", that E and O stand in when E explains O (other than the explanation relation itself).[19] If a theory of chance explanation were to not identify any such relation, the "theory" would simply be a stipulation that E explains O. The general argument faced by any theory of chance explanation takes the form of a dilemma: either the theory's candidate R implies that there is a correct principle of probability coordination (and so the theory inherits the applicability problem) or the theory's candidate R is not the kind of relation that can hold between E and O (and so the theory rules out the possibility of chance explanation). If, for example, E explains O only if E is an inductively strong argument for O, then there is a correct principle of probability coordination. If, instead, E explains O only if E is a deductively valid argument for O, then E does not explain O since E is not a deductively valid argument for O.

The SR model succeeds in identifying a suitable relation that E and O stand in when E explains O (i.e., being the best possible grounds for degree of confidence in), but E and O stand in that relation only if there is a correct principle of probability coordination. So, the SR model faces the applicability problem. The manipulability model has it that E explains O only if E is relevant to manipulating, controlling, and changing O. But, the only way I can see for Woodward to argue that E *is* relevant to manipulating, controlling, and changing O is to argue that E has practical import for agents who care about O. That argument, however, requires some principle of probability coordination. So, either the manipulability model is not a model of chance

[19] R need not be a relation that is more fundamental than explanation. On the view I explore in the final section, for example, chance explanation is more fundamental than is probability coordination. Furthermore, strictly speaking R may be identical to the explanation relation, but a theory of explanation must be able to describe R without invoking the concept "explanation".

explanation or it faces the applicability problem. Finally, the D-N-P model either simply stipulates that E explains O (and so is not a genuine theory of chance explanation) or it posits that E explains O only if E explains why our degree of confidence in O should be one value rather than another. So, the D-N-P model is a theory of chance explanation only if there is a correct principle of probability coordination. These examples illustrate how difficult it is to find a suitable relation, other than probability coordination, that explanations of chance outcomes bear to the occurrence of those outcomes. Because of this difficulty, I suspect that no otherwise plausible theory of chance explanation can make do without some principle of probability coordination, and so conclude that all theories of chance explanation inherit the applicability problem.

5. What Now?

There are two different orientations that we might take toward the connection between chance explanation and probability coordination, and which orientation we take constrains our options for responding to the applicability problem. On the first orientation, which I take to be the (implicitly) standard orientation, there are chance explanations *in virtue of* the correct principle of probability coordination. Given the standard orientation, there are two natural ways to respond to the applicability problem for chance explanation. First, we might respond by simply denying that there are chance explanations. Perhaps there is a correct principle of probability coordination, but (according to this response) justifying some principle of probability coordination is irrelevant to our understanding of scientific explanation. This response is not obviously a disaster, since some philosophers have found it independently plausible that there are no chance explanations. However, it comes with the unappealing consequence that chance events, which occur all the time, are in principle inexplicable. Second, we might respond by spotting ourselves some principle of probability coordination when formulating theories of chance explanation. That there is a correct principle of probability coordination seems intuitively obvious, we might argue, and justifying that principle is a foundational issue for the metaphysics of chance that we need not settle in order to further our understanding of chance explanation. This response is, to my mind, better than is denying that there are chance explanations. And, it is familiar enough that we must sometimes spot ourselves claims about more fundamental domains to make progress toward understanding less fundamental domains. For example, it seems appropriate that we spot ourselves a justification of induction when formulating a theory of confirmation. Nevertheless, this response is not ideal since it increases the number of claims that our philosophy of science must, at least for now, accept on faith.

Rather than endorse either of these two responses, I suggest that we explore a very different orientation toward the relationship between chance explanation and probability coordination, by viewing chance explanation as more fundamental than probability coordination. If this were the correct orientation, then the above two responses to the applicability problem would not be viable. First, denying that there are chance explanations would amount to denying that there is a correct principle of probability coordination. This would be quite a bullet to bite—though it is hard to say why beliefs about chances should determine our degrees of confidence, it is extremely intuitive that they do. Second, if chance explanation is more fundamental than is probability coordination then we should not begin our exploration of scientific explanation by first assuming probability coordination— rather, we should be formulating our theories of explanation with an eye toward grounding the correct principle of probability coordination.

If the correct principle of probability coordination holds because of the nature of chance explanation, then it is no surprise that traditional attempts to solve the applicability problem with only metaphysical theories of chance, while ignoring chance explanation, have all failed. I suspect that grounding probability coordination in a theory of chance explanation is our best hope for solving the applicability problem. Explanation is Janus-faced: it is constrained by purely objective features of the world (such as truth) but it is also tied to subjective features of our inquiry (such as understanding). As such, explanation seems like just the right tool to bridge the gulf between objective chance and subjective degree of confidence.

But making good on this strategy for solving the applicability problem is far from trivial. The rough idea is that there are epistemic norms that connect beliefs about explanations and beliefs about the targets of those explanations (such as, but not limited to, inference to the best explanation), and that these norms, combined with the role that chances play in chance explanation (e.g., that chances are explanations of their outcomes, or that chances are essential components of explanations of their outcomes, or that chances ground explanations of their outcomes, etc.), imply the correct principle of probability coordination. Perhaps finding such a derivation from otherwise attractive claims about the nature of chance explanation will turn out to be as difficult as is deriving some principle of probability coordination from otherwise attractive metaphysical theories of chance.[20] Still, it's worth a try. The applicability problem for chance explanation arises because chance explanation and probability coordination are inexorably linked—a solution to the applicability problem would deepen our understanding of both chance and explanation.

Works Cited

Ayer, A.J. 1963. 'Two Notes on Probability', in *The Concept of a Person and other Essays*, Macmillan, 188-208.

Carnap, R. 1950. *Logical Foundations of Probability*. Chicago: University of Chicago Press

Feigl, Herbert. 1970. The 'Orthodox' View of Theories: Remarks in Defense as well as Critique. In Minnesota Studies in the Philosophy of Science, Volume IV, eds. M. Radner and S. Winokur. Minneapolis: University of Minnesota Press

Friedman, Michael. 1974. Explanation and Scientific Understanding, Journal of Philosophy, 71: 5-19. G

Frigg, R & Hoefer, C. 2010: "Determinism and Chance from a Humean Perspective". In Friedrich Stadler, Dennis Dieks, Wenceslao González, Hartmann

[20] [footnote omitted for blind review]

J., Uebel Stephan, Weber Thomas & Marcel (eds.), *The Present Situation in the Philosophy of Science.* Springer. 351-72.

Gillies, D. 2000, *Philosophical Theories of Probability*, London: Routledge.

Glymour, B. 1998. Contrastive, non-probabilistic statistical explanations. Philosophy of Science 65 (3):448-471.

Hall, N. 1994: "Correcting the Guide to Objective Chance", *Mind* 103:505–518.

Hájek, A. 1996. "Mises redux" — redux: Fifteen arguments against finite frequentism. Erkenntnis 45 (2-3):209--27.

--------- 2007. The reference class problem is your problem too. Synthese 156 (3):563—585.

--------- 2009. Fifteen arguments against hypothetical frequentism. Erkenntnis 70 (2):211 - 235.

Hempel, C. 1965: 'Aspects of Scientific Explanation,' in *Aspects of Scientific Explanation*, New York, NY: Free Press.

Howson, C. and Urbach, P., 1993, *Scientific Reasoning: The Bayesian Approach*, La

Salle, IL: Open Court, 2nd edition.

Jeffrey, R. 1969. Statistical explanation vs. statistical inference. In Nicholas Rescher (ed.), Essays in Honor of Carl G. Hempel. Reidel. 104--113.

Kitcher, P.1989. 'Explanatary unification and the causal structure of the world.' In Philip Kitcher & Wesley Salmon (eds.), Scientific Explanation. Minneapolis: University of Minnesota Press. 410-505.

Kolmogorov, A. N., 1933. *Foundations of Probability*, New York: Chelsea Publishing Company, 1950.

Kyburg, H., 1965. 'Comment.', *Philosophy of Science*, 32: 147–51.

Lewis, D. 1986 (1980): 'A Subjectivist's Guide to Objective Chance', in *Philosophical Papers*, Vol. II, New York: Oxford University Press, pp. 83-113.

------- 1981: 'Causal Decision Theory', *Australasian Journal of Philosophy*, Vol. 59, No. 1; March 1981

------- 1994: 'Humean Supervenience Debugged', *Mind* 106: 321-34.

Loewer, B. 2004: "David Lewis's Humean Theory of Objective Chance", *Philosophy of Science*, 71:1115–1125

Meacham, C. J. G. 2010: “Two Mistakes Regarding the Principal Principle”, *British Journal for the Philosophy of Science* 61 (2):407-431.

Nelson, K. 2009: “On Background: Using Two-argument Chance”, *Synthese* 166 (1):165 186.

Salmon, W. 1967. *The Foundations of Scientific Inference*. [Pittsburgh] University of Pittsburgh Press.

---------- 1971, ‘Statistical Explanation’, in *Statistical Explanation and Statistical Relevance*, W. Salmon, (ed.), 29–87, Pittsburgh: University of Pittsburgh Press.

-----------1989, *Four Decades of Scientific Explanation*, Minneapolis: University of Minnesota Press.

----------- 1997, ‘Causality and Explanation: A Reply to Two Critiques.’, *Philosophy of Science*, 64: 461–477.

Scriven, M. 1959. ‘Explanation and Prediction in Evolutionary Theory”, *Science*, 130, Number 3347.

----------. 1962. ‘Explanations, Predictions, and Laws’, in *Scientific Explanation, Space, and Time* (Minnesota Studies in the Philosophy of Science: Vol. 3), H. Feigl and

G. Maxwell (eds), 170–230. Minneapolis: University of Minnesota Press.

Strevens, M. 1999. Objective probability as a guide to the world. *Philosophical Studies* 95 (3):243-275.

Von Mises R., 1957, *Probability, Statistics and Truth*, revised English edition, New York: Macmillan.

Winkler, R. 1969: ‘Scoring Rules and the Evaluation of Probability Assessors’, *Journal of the American Statistical Association*, 64: 1073-1078.

Woodward, J. 2003. Making Things Happen: A Theory of Causal Explanation. Oxford University Press.

Van Fraassen, B., 1980, *The Scientific Image.* Oxford: Oxford University Press