

The Role of Chance in Explanation*

Bradford Skow

Abstract

‘Those ice cubes melted because by melting total entropy increased and entropy increase has a very high objective chance’. What role does the chance in this explanation play? I argue that it contributes to the explanation by entailing that the melting was *almost necessary*, and defend the claim that the fact that some event was almost necessary can, in the right circumstances, constitute a causal explanation of that event.

1 Introduction

When a chance event is explained what role does its chance play in the explanation? To answer this question we must, first, say something about what a body of fact must do to achieve the status of an explanation, and second, say how facts about E’s chance contribute to the achievement of that status.

At least some explanations of events are causal explanations. And there is a close connection between *being a cause of E* and *raising the chance of E*. So it is easy to identify a role that E’s chance may play in an explanation of E. It may help to determine what E’s causes are.

But this cannot be the whole story. There are chancy explanations in which chances do not play this role. So what role do chances play there? Here is my view in brief. There is a style of explanation that I call necessity explanation. A necessity explanation of E shows that E *had* to occur. (This needs to be refined before we have a definition of necessity explanation; I get to that in section 5.) There are also

*Published in *Australasian Journal of Philosophy* 92 (2014): 103-123.

‘almost necessity’ explanations, which show that the event being explained almost had to occur. Another role that E’s chance can play in an explanation of E is to ground the fact that the occurrence of E was almost necessary.

2 Chance, Causal Explanations, and Two Examples

Let me begin by describing in more detail the role chance plays in paradigmatically causal explanations. For now let us say that a body of fact constitutes a causal explanation of E iff it is a body of fact about what did or did not cause E. (I will later say that this is too restrictive.) Then the mere fact that E had a certain chance of occurring (high, low, or middling) cannot by itself causally explain why E happened. For the fact that E had, say, a 30% chance of occurring is compatible with *any* story about what caused E. It is even compatible with the possibility that E is uncaused.

(Really? Doesn’t the fact that E had a 30% chance of occurring at least imply something quite general and abstract about E’s causal history, namely, that its causal history is not deterministic? I do not think that even this is true. Even if determinism is true events can have non-trivial objective chances. This is a controversial claim, but I agree with the many others who take the nature of the probabilities in statistical mechanics to establish its truth. For an argument against this claim see [Schaffer 2007]; for an argument for it see [Loewer 2001].)

If the fact that E had a certain chance cannot causally explain E, what about facts about how E’s chance *changed*? Facts like that cannot causally explain E either. The fact that E’s chance went up from 30% to 90% just before it occurred does not even entail that E *was* caused. Maybe the change in chance was spontaneous.

Still, facts about E’s chance are not *completely irrelevant* when it comes to causally explaining E. What facts about E’s chance can do is ‘ground’ facts about what E’s causes are, and thereby ground facts that do causally explain E. Let me say something brief about (i) grounding and (ii) how chances ground causes.

Roughly speaking, fact F grounds fact G when G obtains ‘in virtue of’ the fact that F obtains. Exactly what this amounts to is a tricky question in metaphysics that we can safely ignore. For our purposes it is enough to point to paradigm cases. Here is one: the fact that the gas in this box is hot is grounded in facts about how

its constituent molecules are moving.¹

When E is a chancy event then the fact that some other event C is a cause of E is grounded in facts about how the chance of E depends on the occurrence of C. In straightforward cases causes raise the chances of their effects.² So when a claim about chance appears in a paradigmatically causal explanation of some event E (one that fits the restrictive conception of causal explanation I am working with), the role of that claim is not really to *help explain why E occurred*. Instead its role is to *ground a fact that helps (causally) explain why E occurred*.³

These ideas about the role of chance in causal explanation resemble some things Michael Scriven said a long time ago about the role of laws in explanation. Carl Hempel said that explanations must cite laws of nature. Scriven objected: no, explanations just need to cite causes. But, he continued, of course there is a connection between laws and causes: facts about the laws help ground facts about what causes what. But (again) G may ground a fact that explains E without itself explaining E (see the discussion in [Hempel 1965: 359]).

This story about the role of chance in explanation fits lots of examples. One famous example it fits well is due to Michael Scriven [Scriven 1959]:

Paresis. Jones develops paresis. Why? Because he earlier contracted syphilis and left it untreated. (Only people who have untreated syphilis

¹One way to explain a fact F is to show how F is grounded in deeper facts. This is a ‘grounding’ explanation. None of my examples are grounding explanations and I will ignore grounding explanations in this paper.

²Not all cases are straightforward. If Billy hangs back because he sees Suzy throw her rock then it is Suzy’s throw that causes the window to break. This is so even if Suzy has worse aim and her throwing actually lowered the chance of the window’s breaking. So just how do facts about chance ground facts about causes? To answer this question is to give an analysis of causation partly in terms of chances. The guiding idea behind this kind of analysis is the idea that typically, causes increase the chance of their effects. But this is not the place to get in to the nitty-gritty details of what these analyses look like, for which see [Hitchcock 2010].

³This view resembles the Statistical Relevance model of explanation [Salmon 1989: 62-67], though that model incorporates a specific analysis of causation in terms of probability-raising. It is also similar to Humphreys’ theory of chancy explanation [Humphreys 1989].

develop paresis, but not all of them do: there is a 30% chance that those with untreated syphilis develop paresis.)

Here what explains the fact that Jones developed paresis is the fact that he contracted syphilis and did not have it treated. And the syphilis explains the paresis because it is a cause of the paresis. The claim about the chance of getting paresis that appears above is not really part of the explanation of the paresis; it is just there to underwrite the fact that the syphilis is a cause.

However, this story about the role of chance in explanation does not fit all examples. Here are two:

Radioactive Decay. A radioactive atom is created in the lab. It decays within 10 seconds. Why? Well, radioactive decay is a genuine indeterministic process.⁴ The atom was an atom of an element with a very short half-life: atoms of this element have a physically basic 95% chance of decaying in 10 seconds.

The Ice Cubes. I have a water glass on my desk. I put ice in the water and it melts. Why, once it was in there, did the ice melt, rather than increase in size, or remain frozen? If the ice cubes had increased in size or remained frozen then the combined entropy of the ice and water would have decreased or remained the same. But the second law of thermodynamics says that if a system is adiabatically isolated and not in thermal equilibrium then the chance that its entropy will go up is extremely high. And the glass of ice water is adiabatically isolated (for all intents and purposes, since the rate at which the water and the air exchange heat is much slower than the rate at which the water and the ice exchange heat).

Both of these are examples of explanations. But the role the chances cited in these explanations play is not that of grounding facts about a cause of the event being

⁴This is a controversial claim about the interpretation of quantum mechanics. But even if it is false, what matters is that this example would be an example of an explanation if the claims it makes were true. Throughout the paper I will assume that those claims are true.

explained. In **Radioactive Decay** we are told that the atom's chance of decay is 95%, but the explanation does not identify any prior factors that raised or lowered the atom's chance to 95%. And facts about an event's chance can ground a fact about what caused it only if they are facts about what prior events changed its chance. That is what happens in **Paresis**: untreated syphilis raises the chance of paresis from 0 to 30%. Nothing like that happens in **Radioactive Decay**.

One might reply that the creation of the atom is a cause of the decay, and that the explanation works by citing this cause. But that is not so. For we may suppose that the explanatory information in **Radioactive Decay** was delivered to people who, like we did, already knew that the atom had been created (and then decayed). What they wanted, and what they got in **Radioactive Decay**, was explanatory information that they did not already have. The fact that the atom had a 95% chance of decaying is, for them, additional explanatory information. But if the only role that the 95% chance plays is that of grounding the fact that the atom's creation is a cause of its decay then the citation of that chance is not additional explanatory information. So the chance must play some other explanatory role.

Might there be other causes of the decay that I have not considered, that have their status as causes grounded in the 95% chance of decay? Maybe, for example, the atom's being one isotope rather than another caused the decay by raising its chance. My reply is the same. Even for someone who knows about these other causes (in this case, someone who knows what isotope the atom is, and knows that its being that isotope raised its chance of decay), the fact that the atom had a 95% chance of decay is additional explanatory information.

Similar thoughts apply to **The Ice Cubes**. It is true that the fact that the ice was placed in the water when it could have continued to languish in the freezer instead raised the ice's chance of melting. But the explanation above is not the causal explanation that cites this cause. For we may assume the audience for this explanation already knows about this cause. What the audience wants to know is why, *given* that they are in room-temperature water, the ice cubes melt. And the explanation does not cite any of the further causes of the melting (the most salient of which is the transfer of energy from the water to the ice). It just gives the chance of melting (and sketches a derivation of that chance from the condition of the ice

and the second law).

If the chances in **Radioactive Decay** and **The Ice Cubes** do not aid the explanatory project by grounding facts about causes then what role do they play?

3 Necessity Explanations

I am not at all sure that the chances in these two explanations play the same role. But I do think I know what role the chance in **The Ice Cubes** plays: the role it plays is that of showing the event being explained to have been *almost necessary*. What I will do is argue that showing an event to have been almost necessary (I will explain what I mean by this) is a way to explain it, and that, in **The Ice Cubes** at least, the fact that the melting had a very high chance of occurring entails that it was almost necessary. After that I will say a few things about the radioactive decay example.

I will begin my argument by discussing a kind of explanation that I call *necessity explanation*. I am going to introduce this kind of explanation by example. Necessity explanations are explanations relevantly like these two:

The Supernova. A star explodes and all the bits of matter that fly away from it move at less than the speed of light. Why? Because they have to. The laws of special relativity forbid massive bodies from accelerating from below to above the speed of light.

The Hiking Habit. Jones enjoys hiking along a particular trail up and down Mt. Washington. He has a habit of checking his watch while he hikes and he soon notices that every time he does the hike there is some positive real number r such that he is at the same place on the mountain after he has been hiking up for r hours as he is when he has been hiking down for r hours. Why? Because it is impossible for there to fail to be such an r . Since Jones moves through space continuously we can represent any of his possible hikes by a pair of functions f and g , where $f(s)$ is the distance he has hiked up after s hours of ascent and $g(s)$ is the distance he has hiked down after s hours of descent. Then the function $h(s) = g^{-1}(f(s))$ is a continuous function from the closed interval $[0, 4]$ to itself (the hike never takes longer than 4 hours).

It is a theorem of analysis that every continuous function from a closed interval to itself has a fixed point. And if h has a fixed point x , so that $h(x) = x$, then $g(x) = f(x)$. But this just means that Jones is at the same place x hours up as he is x hours down.

These are examples of necessity explanations. But what *is* necessity explanation? I am not going to propose a definition yet. Right now all I want to say is this: a key feature of these two examples is that they show that the target events *had* to happen, that their non-occurrence was impossible. The fact that the explanations show this has a lot to do with their being explanations.

I hasten to add that I do not think that showing an event to be necessary is sufficient for a body of fact to constitute a necessity explanation. That thesis gives the wrong verdicts about examples that illustrate the asymmetry of explanation. It is physically necessary that the flagpole have a certain height, given the laws of geometric optics and the length of the flagpole's shadow. But showing that the flagpole's height is necessary on this basis does not explain it.

Is necessity explanation a kind of explanation that is entirely distinct from causal explanation? If so there are (or can be) necessity explanations that are not causal.⁵ Or, alternatively, is necessity explanation just a species of causal explanation? If so then by showing an event to have been necessary one has (perhaps indirectly) provided information that counts as *causal*-explanatory information. I think the second view is correct. But for now I want to set the question whether necessity explanations are causal aside. The only claim I am relying on now is that there is such a thing as necessity explanation.

Although I have not offered a definition of necessity explanation I want to point out a few things about it before going on. First, necessity explanations do not all need to invoke the same variety of necessity. **The Supernova** shows the event being explained to have been physically necessary while **The Hiking Habit** shows the event being explained to have been necessary in a stronger sense (meta-

⁵A grounding explanation of F shows F to be necessary in light of some deeper facts. And grounding explanations are not causal. But, again, I am ignoring grounding explanations, and stipulate that 'necessity explanation' as I use it does not apply to grounding explanations.

physically necessary). Second, neither explanation shows an event to have been unconditionally necessary. Each explanation works by showing the target event to have been necessary *given* certain other facts: in **The Supernova** these facts include the fact that the supernova occurred at all; in **The Hiking Habit** these facts include the fact that Jones moves continuously. So conditional necessity is all that is required for a necessity explanation.

4 ‘Almost Necessity’ Explanations

So there is such a thing as necessity explanation. The next thing I want to do is argue that there is such a thing as ‘almost’ necessity explanation. (I will eventually say that chances can play a role in almost necessity explanations; but first things first.)

One argument starts from the fact that ‘explanatoriness’ comes in degrees. A theory of explanation is too crude if the only distinction it draws is between bodies of fact that are explanatory and those that are not. Some bodies of fact are more explanatory than others. Similarly, a theory of modality is too crude if the only distinction it draws between truths is between truths that are necessary and those that are contingent. Some contingent truths are ‘closer’ to being necessary than others. The larger the size of the set of possible worlds in which a proposition is true, the closer it is to being necessary. Now, importantly, it will not do to say that a proposition is closer to being necessary the larger the *number* of possible worlds in which it is true; too many propositions are true in infinitely many worlds. Making sense of how close a proposition is to being necessary requires something like a measure on the set of possible worlds. As I will discuss below, objective chances are one place to look for such a measure.

So we have a way of understanding how close a proposition is to being necessary. This immediately gives us a way to understand what it is for a proposition to be almost necessary. Almost necessary propositions are just those that are very close to being necessary. (There are more sophisticated notions of closeness to necessity. For example, we might want to say that P is closer to being necessary than Q even if P and Q are true in the same proportion of possible worlds, provided that

the nearest Q world is farther away than the nearest P world. Or we might want to say that P is closer to being necessary than Q even if they are true in the same proportion of possible worlds, provided that the *average distance* to a Q world is greater than the average distance to a P world. I am not going to make use of any of these notions in this paper, even though I think they are interesting and important.)

I have here been talking about how close a proposition is to being unconditionally necessary, even though the main notion at work in necessity explanations is conditional necessity. But closeness to conditional necessity can be understood along similar lines. A proposition is closer to being necessary given F when it is true in a larger proportion of the F-worlds.

So explanatoriness and necessity both come in degrees. Now a body of fact that shows an event to have been necessary can (if it meets other conditions) constitute a necessity explanation; and necessity explanations are typically very good explanations. I say that a body of fact that shows an event to have been *almost* necessary can also constitute an explanation, one that, if it is less good than a necessity explanation, is only slightly less good.

The alternative view, the one I am arguing against, is that showing some event to be almost necessary cannot constitute an explanation at all. But why should a small change in the degree to which a body of fact shows some event to be necessary correspond to a huge change in how explanatory that body of fact is? It seems to me that how explanatory some body of fact is should vary continuously with how necessary it shows the target event to have been (provided we hold fixed that it continues to do other things that a body of fact must do to be an explanation). If some body of fact F1 shows the occurrence of E1 to have been (conditionally) necessary; and some other body of fact F2 does something *extremely similar* for some other event E2, namely showing it to have been almost (conditionally) necessary; then surely the second body of fact is not doing a whole lot worse with respect to meeting the aim of explanation than the first one is. And the first body of fact meets that aim extremely well.

5 Necessity Explanations are Causal

That was a pretty abstract argument. Now for a less abstract one. I want to argue that it is evident, in light of the definition of necessity explanation, that showing a fact to be almost necessary can (under the right conditions) constitute an explanation. Now I have not yet given a definition of necessity explanation. So far I have just stipulated that necessity explanations are explanations relevantly similar to **The Supernova** and **The Hiking Habit**. But, I claim, **The Supernova** and **The Hiking Habit** are special kinds of causal explanations. The most interesting respect in which an explanation might resemble these two paradigms is to also be a causal explanation of that special kind. That is, I identify necessity explanation with that kind of causal explanation.

If it seems that my two paradigm necessity explanations cannot be causal that is only because we have been working with too crude a conception of what causal explanation is. I said earlier that a body of fact constitutes a causal explanation of some event E iff it is about what did or did not cause E. But causal explanation encompasses a broader class of bodies of fact than this.

One way to see why is to start with James Woodward's theory of causal explanation (with which I am in broad agreement). He says that a body of fact is causal-explanatory iff it can be used to answer a 'what-if-things-had-been-different' question. Woodward puts it like this: 'the explanation must enable us to see what sort of difference it would have made for the explanandum if the factors cited in the explanans had been different in various possible ways' [2003: 11]. This characterization entails that citing a cause can explain. For learning that C is a cause of E lets us see what difference it would have made for E if C had not occurred (and certain other features of the situation had been as they actually were).⁶ But Woodward's characterization counts other kinds of facts as causal-explanatory as well. The fact that if some specific alternative to C had occurred in C's place then some specific alternative to E would have occurred in E's place constitutes a causal

⁶Much of Woodward's book is about what features are to be 'held fixed' in these counterfactuals. The theory needs a way to solve the problem of preemption: cases where if the cause had not occurred, a backup cause would still have produced the effect. This is an important problem but throughout this paper I will ignore it.

explanation, because it also answers a what-if-things-had-been-different question. And the explanation it constitutes is not identical to the one constituted by the fact that C was a cause of E. For the facts provide answers to different what-if-things-had-been-different questions.⁷

Now I should say that I doubt that the condition Woodward states is necessary as well as sufficient. For it seems to me that the fact that some event E was uncaused constitutes a causal explanation of it.⁸ But learning that E was uncaused does not enable us to see what sort of difference it would have made for E if the factors cited in ‘E was uncaused’ had been different; for no factors are cited. Still, Woodward’s condition does seem to me sufficient for a body of fact to constitute a causal explanation, and that is what matters.

Woodward’s condition seems to me sufficient; but I think that he puts the focus in the wrong place. A shift in focus provides a better sufficient condition. Woodward’s focus is on ‘factors’ (paradigmatically, events) cited in the explanation. The way he puts it, an explanation qualifies as causal if it names some factors and says something about how wiggling one or another of those factors changes how or whether E occurs. But I think the focus should be on *the event being explained*. An explanation qualifies as causal, I say, if it says something about what it would have taken for E not to occur, or (more specifically) what it would have taken for some alternative or range of alternatives to E to have occurred instead, or what it

⁷Lurking in the background here is a counterfactual analysis of event causation: C is a cause of E iff C and E are distinct and had C not occurred, and certain facts been just as they actually are, then E would not have occurred. But Woodward treats causation as a relation between variables rather than events. (A variable is a kind of ‘generalized event’. While events must either occur or fail to occur, a variable may in principle have more than two values. The current atmospheric pressure is an example of a variable: when measured in atmospheres it can take any smallish real number as a value.) Taking causation to be a relation between variables makes the distance between ‘C is a cause of E’ and ‘if C* had occurred instead of C then E* would have occurred instead of E’ smaller. It is also smaller on David Lewis’s analysis of causation as influence [Lewis 2000], and the view that causation is contrastive [Schaffer 2005].

⁸Lewis [1986] uses the idea that an event can be causally explained by saying it was uncaused to respond to some alleged examples of non-causal explanations.

would have taken for E to have occurred in a different way. If a body of fact can do this without naming or citing any particular factors, it should still count as a causal explanation.

Even after this change of focus the fact that C is a cause of E still qualifies as a causal-explanatory fact. But by seeing that the class of causal explanations is broader than explanations that just describe causes we can see that there is a subclass of causal explanations that deserve to be called necessity explanations.

The limiting case of an explanation that says what it would have taken for E not to have occurred is an explanation that shows that *nothing* could have prevented E. Now this ‘nothing’ need not be completely unqualified. Causal explanations are usually offered against a background of assumptions about the existence of some causes of the target event. So an explanation that shows that nothing could have prevented E may just show that nothing *that is possible in some relevant sense* could have prevented E, *given* the occurrence of certain of E’s causes. (The relevant sense of possibility might be physical possibility or it might be metaphysical possibility or it might be some other variety of possibility.) An explanation that does this is, I say, a necessity explanation, and the assumed causes are the facts that the necessity is conditional on.

Woodward thinks that a causal explanation must cite some factor on which E depends. (He denies that **The Supernova** is a causal explanation, and this is why [Woodward 2003: 193].) But if there are no (relevant) factors on which E depends, if E is (conditionally) necessary and so *independent* of all (relevant) prior events, then the fact that it is independent of prior events, that fact that there are no factors to cite, should constitute a causal explanation of E.

So: paradigmatic causal explanations answer ‘what would it have taken for E not to have occurred?’ by citing specific factors that, had they been different, would have prevented E. What necessity explanations do is give the null answer: there are no such factors. But they are still answering the same question, and for that reason count as causal explanations.⁹

⁹In a recent paper Marc Lange argues that ‘distinctively mathematical explanations’ of physical phenomena are non-causal [Lange forthcoming]. His description of these explanations makes them look a lot like necessity explanations. (He would

6 Causal Explanation and Almost Necessity Explanation

The fact that an event was *almost* necessary can also constitute a causal explanation of that event. For learning that E was almost necessary can give us a large part of the answer to ‘What would it have taken for E not to have occurred?’ By learning that E was almost necessary we may learn that among all the (relevant) changes to the (relevant) factors, only a *very small* proportion of them would have prevented E. When we do the facts we have learned constitute an almost necessity explanation.

Compare an example of a paradigmatic causal explanation with an example where some event is shown to have been almost necessary. Suppose we see a window cracked in a beautiful star pattern and we want to know why it cracked in that pattern. A materials scientist tells us just how the cracking pattern depends on facts about the impact, such as facts about how fast the rock was going, the angle at which it hit the glass, and so on. She shows us exactly how varying those factors varies the cracking pattern. So she situates the actual cracking pattern in a range of alternative patterns, and situates the rock’s actual impact parameters in a range of alternative parameters, and shows us just which parameters produce just which patterns. Now we have an excellent causal explanation of the cracking pattern. (My language here again draws on Woodward’s way of thinking about causal explanation; he writes in one place that causal explanations ‘locate their explananda within a space of alternative possibilities and show us how which of these alternatives is realized systematically depends on the conditions cited’ [Woodward 2003: 191].)

Here is a more abstract way to put what we acquire when we acquire this explanation. Use numbers to parameterize the space of possible cracking patterns (since this is a toy example it does not matter how this is done), and let the variable **C** take these numbers as possible values. Similarly let **P** take as possible values regard **The Hiking Habit** as distinctively mathematical.) If he is right then I am wrong to say that necessity explanations are causal. Now I am not sure that he and I are using the same criteria to classify explanations as causal. Setting that aside, the most his argument shows is that **The Supernova** and **The Hiking Habit** should not be classified together. But what matters for my argument is that there are some explanations that fit my definition of necessity explanations as a kind of causal explanation. And there are. **The Supernova** is one. It does not matter if there are other explanations that do not.

the numbers that parameterize the possible impact states of the rock. Then what we have learned when we learn why the window cracked in a star pattern, what we learn when we learn why $\mathbf{C} = s$, where s represents the star pattern, is which function maps an impact parameter to the cracking pattern it produces (in the circumstances of Suzy's throw).

Now let's think through an example in which an event is shown to be almost necessary. We see a different window, made out of a new and experimental kind of glass, cracked in a strange circular pattern. Our materials scientist's eyes light up when we ask her about this window. She says she is very excited about this new kind of glass. It was designed to crack in that circular pattern for almost any impact. Billy, whose rock hit the window, would have to have either been extremely lucky or had extremely good aim (and extremely good knowledge of the window's structure) for his rock to have cracked the window in any other pattern. What the scientist has told us is that given Billy's throw the window's cracking in that pattern was almost necessary.

The scientist has not done what she did in the first example and given us a relatively complete paradigmatically causal explanation. She has not told us just which function g maps impact parameters to the cracking patterns they produce. Nor has she given us a necessity explanation of the cracking pattern. She has not told us that the cracking pattern was independent of the details of the throw, that g is a constant function. But she has still given us *a lot* of information about g . She has told us that for almost any impact parameter p , $g(p) = c$ (c represents the circular pattern). And to give a lot of information about this function is to give a causal explanation.

Certainly it makes no sense to require a causal explanation to give complete information about this function. For simple causal statements give less than complete information. On a theory of causation like Woodward's, the statement that \mathbf{P} is a cause of \mathbf{C} gives the information that *some* (maybe just one!) non-actual value of \mathbf{P} gets mapped by the function to some non-actual value of \mathbf{C} . The fact that $\mathbf{C} = c$ was almost necessary gives way more information about the function g than the simple causal statement that \mathbf{P} is a cause of \mathbf{C} . And if it does more of what simple causal explanations do then it is also a causal explanation.

7 Chance and Almost Necessity

Are high chance events almost necessary? If they are then citing the high chance of an event can be way to explain why that event occurred, and we have found another role for chances to play in explanations: that of establishing that the event being explained was almost necessary.

Let's look back at the examples. In **The Ice Cube** the high chance that the ice cubes melt comes from statistical mechanics. When you chase down where statistical mechanical chances come from you find this. The ice and water system is in a certain macrostate, which may be characterized by quantities like the temperature of the water and the volume of the ice cubes. Any given macrostate can be realized by (infinitely many) microstates, where a microstate specifies the exact position and momentum of each of the system's molecular constituents. The laws of thermodynamics only concern the system's macrostate, while the underlying fundamental dynamical laws concern the system's microstate. Statistical mechanics comes with a measure that assigns sizes to sets of microstates. And it identifies the chance that one macrostate M will evolve into another N with the (relative) size of the set of microstates that start in M and evolve under the micro-dynamics into microstates in N .¹⁰ The size of the set of a system's microstates that lead to behaviour that violates the second law of thermodynamics is extremely small. Since the theory identifies chances with relative sizes, this means that the statistical mechanical chance that a system will behave anti-thermodynamically is small.¹¹ So here, at least, the claim that thermodynamic behaviour has a very high chance entails that thermodynamic

¹⁰I am here assuming the Boltzmannian way of thinking about the relationship between the second law of thermodynamics and statistical mechanics, as defended in, for example, [Albert 2000].

¹¹Another theory might not identify chances with relative sizes. But since probability functions are special cases of measure functions, a theory like that in effect has *two* ways of measuring sizes of sets of states or possible worlds (only one of which is identified with chances). That means that a proposition could be almost necessary relative to one of the measures but not relative to the other. Which kind of almost necessary is explanatorily relevant? I say that only the chance measure is; the other is not. I do not have space here to give an extended defence of this idea. But, briefly, I think it is only the chance measure that has real physical meaning.

behaviour is almost necessary. (In Woodwardian terms, almost any intervention that changed the system's microstate would not affect whether the ice cubes melt.)

So the role of the chance in **The Ice Cubes** is to show the melting to have been almost necessary. What about the role of the chance in **Radioactive Decay**? That chance has a very different metaphysical underpinning (on interpretations of quantum mechanics that make it an indeterministic theory). It is not true that a radioactive atom with a very short half-life might be in any one of a bunch of microstates, a large proportion of which lead deterministically to decay in a short time. Instead there is one and only one microstate — the one the atom is in — and it is undetermined by the laws whether an atom in that microstate will decay. The fact that it has a 95% chance of decay in any given 10 second stretch physically basic; the physical theory does not ground it in anything else.

So the fact that the atom had a high chance of decay does not entail that there was any prior factor such that the decay is almost independent of how that factor is manipulated. Does this mean that **Radioactive Decay** is not an almost necessity explanation? Maybe it does. Maybe the chance in that explanation plays a different role from the chance in **The Ice Cubes**.

I admit to feeling the pull of the idea that these two explanations work in the same way. I do not give that pull much weight. It happens all the time that phenomena that seem similar turn out, on deeper analysis, to be very different. Still, let me briefly speculate about how the two explanations could be similar. What does the fact that the radioactive atom had a high chance of decay tell us, if it does not tell us that a high proportion of the ways of realizing the atom's macrostate lead to decay? In a more metaphysical mood one might think that the fact about chance entails this: if we were to repeatedly rewind history to the time at which the atom was created in the laboratory, then (it is extremely likely that) the atom would decay within 10 seconds almost every time. This claim has the same sort of flavour as the claim that almost every way of intervening to change the ice cubes' microstate still leads to melting. The difference is that the intervention we are speaking of when we speak of rewinding history is a much weirder and remote possibility (if it is possible at all) than is changing the ice cubes' microstate. I am left unsure what to say. Maybe this similarity is enough for **Radioactive Decay** to be an almost

necessity explanation. Or maybe answering questions about what would happen if history were rewound does not have anything to do with explaining (causally or otherwise) why some event occurred.

8 Further Reflections on Chance and Explanation

I said that necessity explanations are often conditional: they may show that an event had to happen, given some other facts, including facts about some of its causes. One might suspect that even some paradigmatically causal explanations turn out to be necessity explanations on my definition. The derivation of the length of the flagpole's shadow from its height, the position of the sun, and the laws of optics both identifies several causes of the shadow's length and also shows that its having that length was necessary given those causes. Isn't this a necessity explanation? I agree that it is but I do not think its being so is any objection to what I have said. *In this case* one learns a body of fact that simultaneously constitutes both a paradigmatically causal and a necessity explanation. That does not mean that all necessity explanations are like this. And in fact they are not: **The Supernova**, for example, is not.

Here is another way to think about the difference between necessity explanations and paradigmatically causal explanations. Paradigmatically causal explanations aim to identify causes, or at least to narrow down the list of possible causes. Necessity explanations aim not to identify causes, but to give information about the *nature of the connection* between the event being explained and the causes already identified. The fact that E has certain events among its causes does not by itself entail that E is *necessary* given those events. This is the additional explanatory information that necessary explanations can provide.

Similarly, almost necessity explanations can be thought of as giving information about the connection between the event being explained and its causes. In paradigmatically causal explanations we care about how E's actual chance compares to the chance it would have had if some event C had not occurred; but we do not care about the absolute value of its chance. In almost necessity explanations (that invoke chance) it is the other way around. We care nothing for how its actual

chance compares to its hypothetical chance, but we do care about the absolute value of its actual chance (the chance conferred on it by its causes, or some relevant subset of its causes). For that value tells us something about how easily E could have failed to obtain.

9 Other Views

There are other views out there about the role the chances in **The Ice Cubes** and **Radioactive Decay** play. It is too big a project to say where I disagree with each philosopher who has a view on this topic, so I will limit my attention to three: two who have recently defended views about the role of chance in explanation — Robert Batterman and Michael Strevens — and one whose work forms an important backdrop to theirs — Peter Railton.

Railton (primarily in [Railton 1978] and [Railton 1981]) says that the chances in **Radioactive Decay** and **The Ice Cubes** play different roles. **Radioactive Decay** is his paradigm case of a chancy explanation. Railton says that this explanation should be regimented as (i) a derivation of the fact that the atom had a 95% chance of decay from the laws of physics and facts about the circumstance in which it was created, followed by (ii) a comment to the effect that the atom actually did decay. If that were Railton's whole theory it would not be very illuminating (which is not to say that it would be false). It would amount to saying: citing physically basic chances just is explanatory; there is nothing interesting to say about *why* it is explanatory.

But his claim about the form of chancy explanations is not the whole of Railton's theory. He has more general views about what a body of fact must do to constitute an explanation. How do the chances in explanations like **Radioactive Decay** help those explanations do those things?

Railton says that to explain is to give an 'account of' or to 'elucidate' the mechanism at work in the production of the event being explained:

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. [Railton 1978: 208]; see also [Railton 1981: 242]

So does citing the fact that the atom had a 95% chance of decaying elucidate the mechanism behind the decay of the atom?

When Railton introduces the idea that explaining is about elucidating mechanisms he does so with an example. Suppose it starts to rain and I say that this happened because the atmospheric pressure was falling. Railton complains that ‘No connection between cause and effect, no mechanism by which falling atmospheric pressure produces a change for the worse in the weather, has been revealed’, and that is why ‘we remain in the dark as to *why* the weather will turn bad’ [1978: 208].¹² In this case it is clear how to respond to a request for further elucidation of the mechanism responsible for the decline in the quality of the weather. Just describe the intermediate events in the causal chain that leads from declining pressure to rain, and also perhaps cite the laws that help ground the fact that those events are related as cause to effect.

But now suppose the atom in **Radioactive Decay** has decayed and I say that the decay was an indeterministic event, that prior circumstances did not physically necessitate its occurrence. This explanation can certainly be improved by adding that the decay had a (physically basic) chance of 95% of occurring. But how does adding this fact about the chance constitute an elucidation of the mechanism responsible for decay? It does not say anything about events intermediate between some earlier cited cause and the decay. It does not fill in the kind of gaps that Railton wanted filled in in the bad weather example.

Railton’s talk of ‘elucidating the mechanism’ suggests that there is some simple or complex physical system that one could tinker with that plays a role in making the atom decay, and that learning about the atom’s chance of decaying is learning something about it. But there is no such thing. Saying that the atom had a 95% chance of decaying is not an elucidation of the mechanism responsible for decay in the most natural sense of that phrase. It does not ground any (relevant) facts about

¹²For the record, I do not agree (and doubt many others would) that we remain *completely* in the dark. We have learned something about the causes of the weather. But maybe Railton’s point is just that there is still a lot of explanatory information we do not have.

the decay's causes.

One way to interpret Railton has him using 'elucidate the mechanism' in something other than its most natural sense. On this interpretation one elucidates the mechanism by describing the laws under which an event falls, even when this implies nothing about what caused that event or how its causes are connected to it. (Railton does say in one place that citing laws counts as part of describing the mechanism [1981: 242].)

I am not sure how well this interpretation fits with other things Railton says. He seems elsewhere to be offering a causal theory of explanation. If so then this interpretation, on which some bits of explanatory information (including some parts of chancy explanations) may be non-causal, cannot be right. But set that issue aside. We were interested in the idea that to explain is to elucidate the mechanism because we wanted to see why facts about chances in explanations like **Radioactive Decay** were explanatory facts. It looked like Railton had an answer: because those facts help elucidate the mechanism. But this only helps if the following two claims are pre-theoretically intuitive: (i) any information that elucidates the mechanism is explanatory; and (ii) facts about basic physical chances help elucidate the mechanism. Now (i) is plausible on the most natural sense of 'elucidate the mechanism'. But I do not think (ii) is. On the other hand, (ii) is true more or less by stipulation on the extended sense of 'elucidate the mechanism' that Railton seems to be using. But on that extended sense (i) is not pre-theoretically intuitive. As far as I can tell Railton offers no defence of it; and I see no reason to believe it.

There is another reason to doubt that Railton has identified the role that the chance in **Radioactive Decay** plays in that explanation. His theory entails *size egalitarianism*: explanations that cite high chances are just as good as explanations that cite low chances. But size egalitarianism is false. Michael Strevens has made a strong case against size egalitarianism (see [Strevens 2000]). Instead of rehearsing his argument here is another one. Suppose we do not know the law governing the radioactive decay of this atom. We have two theories. Theory 1 assigns the decay a high chance per unit time; theory 2 assigns it a low chance per unit time. Which should we believe? Certainly theory 1 (especially if we perform the experiment several more times and the atom decays quickly each time). We should believe

theory 1 because it is a better explanation of the quick decay than theory 2. But then egalitarianism is false.

I have been arguing that Railton has not identified the role that the chance in **Radioactive Decay** plays in that explanation. I am not certain that I have identified that role either. But it is worth explaining why I think Railton is wrong even if I do not know the right answer.

At one point in his discussion Railton wrote that ‘Probabilistic explanation...is not a second-string substitute for deterministic explanation, showing that the explanandum *almost* had to occur’ [1981: 237]. But I think that this is almost right. Chancy explanation, at least sometimes, does work by showing the target event to have almost had to occur. (I would not call this kind of explanation second-string or describe the alternative style as ‘deterministic explanation’.)

Railton also says something about the role of the chance in **The Ice Cubes**. Railton holds that the ‘ideal’, complete explanation of the melting of the ice cubes does not cite any chances. Instead it describes, in complete microphysical detail, the causal process that constituted the melting. No human being ever grasps or delivers to anyone else the ideal complete explanation of any event. What we do grasp and deliver to others are partial explanations. A body of fact constitutes a partial explanation of some event iff it is information about the ideal explanation. Officially Railton focuses on a set of sentences that express the ideal explanation; he calls it the ‘ideal explanatory text’ and says that a body of fact is explanatory iff it is information about the ideal explanatory text. In a well-known passage he says that the role of the chance in **The Ice Cubes** is to convey information about the ideal explanatory text: the explanation citing it

illuminates a *modal* feature of the causal processes involved and therefore a modal feature of the relevant ideal explanatory texts: this sort of causal process is such that its macroscopic outcomes are remarkably insensitive (in the limit) to wide variations in initial microstates. The stability of an outcome of a causal process in spite of significant variation in initial conditions can be informative about an ideal causal explanatory text in the same way that it is informative to learn, regarding a given causal explanation of the First World War, that a world war

would have come about (according to this explanation) even if no bomb had exploded in Sarajevo. [1981: 251]

Railton is certainly on to something here. He is right that the chance in **The Ice Cubes** contributes to the explanatoriness of that body of fact by conveying modal information. But this claim is not consistent with Railton's background theory of explanation. He makes them appear consistent by trading on two interpretations of 'X is information about the ideal explanatory text'.

Railton says that someone who had complete explanatory information about some event E would know exactly what the ideal explanatory text for E says. And he says that partial explanatory information is information 'about that explanatory text'. But he does not really mean any information that is about the text *without qualification* is explanatory information. He writes that explanatory information is information that 'eliminate[s] some degree of uncertainty about what is contained' in the ideal text [1981: 240]. But information like that is information about *what the text says* (and possibly how the text says it). Information that is just plain about the ideal text need not also be information about what the text says.

The problem is that the modal information Railton points to in the quotation, while certainly information about the ideal text, is not information about what the text says or how it says it. The fact that the ideal explanatory text for the melting of the ice cubes would have looked much the way it actually does if the ice had been in a different initial microstate does nothing to reduce *my* ignorance about what that text says. In this regard it is a lot like the fact that there is a guy in Australia who is quite fond of the ideal text for the melting. That fact is also about the ideal text. But it is not explanatory.

(Railton does say that 'the relevant ideal text contains more than 10^2 words in English' is the right kind of information about the relevant text [1981: 246]. He just says that it is a *very small amount* of information, and so is only a *little* explanatory. Could he say something similar about the modal information the chances in **The Ice Cubes** convey? I doubt it. Facts about how many English words a text contains certainly constrains what that text says and how it says it. But why do the modal facts constrain it? In any regard, if the strongest thesis available to Railton is that **The Ice Cubes** is a tiny bit explanatory then his theory is still false. That expla-

nation, and the modal information it provides, is *a lot* of explanatory information, even though it conveys only a little information about the ideal text's contents. So its 'explanatoriness' does not consist in what it says about that text's contents.)

Batterman and Strevens raise objections to Railton's theory of chancy explanation allied to the ones I have raised. A central claim for them is that Railton's ideal texts are too overloaded with detail. They argue that some of those details are explanatorily irrelevant and that bodies of fact that contain less detail will be more explanatory. Each goes on to defend a theory opposed to Railton's.

Batterman [1992] focuses on explanations like **The Ice Cubes** that do not use physically basic chances.¹³ He complains that on Railton's view the chances in these explanations are not closely enough connected to physical laws. On his view chances that explain must be derived from some physical theory (not necessarily a fundamental physical theory) [1992: 343].

But this part of Batterman's view does not tell us how the chances these explanations invoke contribute to their 'explanatoriness'. It is the second part of his view that says this. He says that good chancy explanations must cite a chance very close to one. And, he continues, they must do that because explanations must cite facts that would have constituted very good evidence for the occurrence of the event being explained had we had it before the event occurred [Batterman 1992: 344]. In this respect his theory draws on Hempel's claim that explanatory facts are facts that show that the thing being explained was to have been expected [Hempel 1965].

But then Batterman's theory is false for the same reason Hempel's is: explanatory facts do *not* need to do this. **Paresis** establishes this. No one should have expected Jones to develop paresis knowing he had untreated syphilis; that development was not particularly likely. Still the fact that he had syphilis explains his developing paresis.¹⁴

¹³Actually Batterman would say that he focuses more on explanations of regularities than explanations of single events. He would say that the kind of example he is interested in is an explanation of the fact that *every time* anyone has put ice in a glass in a warm room the ice has melted. But (as Strevens notes in a different context) this explanation of a regularity can still be regarded as an explanation of a single event: it is just a large and spatiotemporally scattered event.

¹⁴Hempel attempted to reconcile his view with **Paresis**. An excellent discussion

Batterman is right to reject size egalitarianism. Explanations (like those in statistical mechanics) that cite chances without saying anything about changes in chance need to cite high chances. But his story about why those chances need to be high to explain is wrong. I maintain that they need to be high because only then do they convey a relatively large amount of causal-explanatory information.

Strevens, like Batterman, rejects size egalitarianism. Unlike Batterman he does not say that high chances are explanatory because of any epistemic properties they have. Saying why, on his view, high chances are explanatory will take some setting-up.

Strevens focuses, for simplicity, on explaining the behaviour of the wheel of fortune.¹⁵ He intends what he says to also apply to examples like **The Ice Cubes**. I will focus on the role he says chance plays in these examples, since he says relatively little about explanations like **Radioactive Decay**.

The wheel of fortune consists in a wheel with pegs around its circumference that may spin about its centre, and a pointer that sticks between the pegs. The pegs divide the wheel into sectors that alternate between red and black.

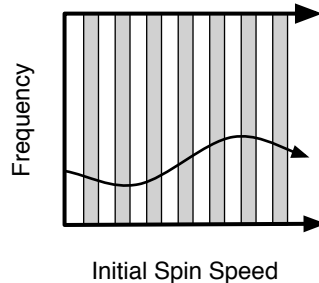
Sam spins the wheel one hundred times, and lands on red close to half the time. Why? Strevens compares a deterministic explanation and a chancy explanation. The deterministic explanation (the one he thinks is the best deterministic explanation) looks like this. The ‘evolution function’ of the wheel gives the colour it will land on as a function of the initial spin speed. We may picture it by representing initial spin speeds on a horizontal axis and colouring white and grey vertical strips over this axis. Initial spin speeds in a grey strip result in the pointer pointing at black; initial speeds in white result in red (see figure 1). Notice that two adjacent strips are close to the same width.

If for each vertical strip in the diagram we plot the frequency with which Sam spun the wheel with that speed, and then we draw lines connecting those dots, and smooth out those lines a little bit, we will get a curve that looks something like the

of what is wrong with that attempt may be found in chapter 4 of Woodward’s book [Woodward 2003].

¹⁵What follows summarizes material from in and around chapter 10 of [Strevens 2008].

Figure 1: Evolution function with the frequencies of Sam's initial spin speeds



one in figure 1.

Importantly, this curve's altitude changes very little between one strip and the next. Strevens calls sets of initial spin speeds that have this property 'macroperiodically distributed'. It is not hard to suspect that if the wheel is spun many times and the initial speeds of those spins are macroperiodically distributed then the wheel will land on red roughly half the time. Look at any one white strip; the frequency with which the wheel is spun with those initial speeds (and so lands on red) is very close to the frequency with which the wheel is spun with a speed in the grey strip to the right (and so lands on black). And those strips are almost the same width. This is true for every strip, so the wheel lands on red roughly half the time.¹⁶

So the best deterministic explanation of Sam's having spun red close to half the time cites the fact that Sam's initial spins were macroperiodically distributed and the features of the wheel and the laws governing it that entail that its evolution function is the one pictured in figure 1.

Strevens also believes that there is a chancy explanation of the fact that the wheel landed on red close to half the time. The chancy explanation cites, not the fact that the initial spin speeds were macroperiodically distributed, but the fact that *the probability that the initial spin speeds were macroperiodically distributed was very high*. Strevens says this about the relationship between the deterministic and the chancy explanation:

¹⁶For a mathematically rigorous statement and proof of this claim see [Strevens 1998] and [Strevens 2003].

The two models are rather similar. They share the same causal mechanism, which is to say that they represent in exactly the same way the causal process by which the frequency of red is produced. The micro-constant [i.e., chancy] explanation differs from the statistical deterministic explanation in that it adds to this mechanism a basing generalization that subsumes the macroperiodicity of the actual initial conditions under a broad pattern of entanglement, namely, the entanglement between the property of being a set of initial spin speeds on a wheel of fortune and the property of being macroperiodically distributed. [Strevens 2008: 382]

Strevens here says that the two explanations provide all the same causal information about the outcomes of the spins. Where they differ is that the chancy explanation provides an extra bit of information, information that is not causal. And it is the claim about the high probability that provides this information.

When Strevens says just what this non-causal information is he uses the terms ‘basing generalization’ and ‘entanglement’, which are technical terms in Strevens’ theory of explanation. I cannot hope to explain the rich meanings he gives these terms here. Fortunately, Strevens gives a more succinct (if less exact) characterization of this non-causal information is later: the explanatory role of the probability claim ‘is to imply the existence of the right sort of pattern, namely, a prevalence of macroperiodicity in the relevant initial conditions’ [2008: 394]. By ‘prevalence’ he means prevalence in the actual and nearby possible worlds. So the deterministic explanation says that the speeds of Sam’s spins were macroperiodically distributed; the claim about chance in the chancy explanation adds to this a fact about *other, similar spins in the actual world* and a fact about *other, similar spins in nearby counterfactual worlds*. It says that almost every other time this and similar wheels were spun, and will be spun, by people like Sam, the initial spin speeds were, and will be, macroperiodically distributed. It says that almost all nearby worlds where Sam spins the wheel differently, his initial spin speeds are macroperiodically distributed. And it says that other people’s spins also have this modal property, not just Sam’s.

Strevens’ view about the role of the chance in **The Ice Cubes** is that it is the

same as the role of the chance in the wheel of fortune explanation. And clearly the role he describes is broadly similar to the one I have described. Like mine, his view can be summed up by saying that what a claim about high probability contributes to an explanation is the fact that the event being explained is almost necessary. But, importantly, what ‘almost necessary’ means here is different from what I mean by it. When I say that some event E was almost necessary I mean that almost every relevant way of manipulating its relevant causes would still have produced E. In a characterization of Strevens’ view it means something else (or perhaps just something stronger). For there it entails (among other things) that most of the *other* people who will later spin the wheel of fortune spin with speeds that, like Sam’s, are macroperiodically distributed.

Here is my view about the wheel of fortune example: the fact that it was highly likely that Sam’s initial spin speeds were macroperiodically distributed, to the extent that it is explanatory, is explanatory because it entails that, in almost all of the nearby worlds in which Sam spins differently, his spin speeds are (still) macroperiodically distributed. (Part of what makes this true is that there are limits to Sam’s ability to control exactly how his hands and arms move) How other people spin in the actual or nearby counterfactual worlds is explanatorily irrelevant.

I will say why I prefer my theory, but arguing in any detail for this preference (and discussing Strevens’ defence of his view) is too big a project for the end of this paper. I just doubt that facts about the speeds at which Joe (for example) will later spin the wheel of fortune are relevant to why Sam landed on red roughly half the time. Part of my doubt stems from the fact that Joe’s spins are causally independent of Sam’s spins. I think that if we had a complete causal explanation of the fact that Sam got red roughly half the time, our understanding would not be improved by learning anything about how Joe spun or what results he got.

Of course if we do not start with a complete causal explanation then learning about Joe might improve our understanding; but that might just be because it tells us something about Sam’s spinning abilities. The fact that no one in history has managed to spin the wheel a bunch of times and get a frequency of red appreciably different from one half might tell us that in all or almost all nearby counterfactual worlds *Sam’s* spins are still macroperiodically distributed; and that kind of infor-

mation, I allow, is causal-explanatory information.

10 Conclusion

Wesley Salmon distinguished the ontic conception of explanation, by which he more or less meant causal explanations, from the modal conception [Salmon 1985: 293-96]. Explanations that fit the modal conception look like necessity explanations: they show that the event being explained had to happen. Salmon thought that these two conceptions were opposed, and rejected the modal conception because he did not think it covered statistical explanation. I think Salmon was wrong twice: first by thinking that modal explanations are not ontic, and second by thinking that explanations that use chances are never modal.

My aim in this paper has been to describe a role that chances can play other than the role they play in paradigmatically causal explanations. That role is to show the event being explained to have been almost necessary. Along the way I have discussed necessity explanations and argued that both they and almost necessity explanations are causal. I am confident that the role I have described is the role that the chance in **The Ice Cubes** plays. I am less confident that the chance in **Radioactive Decay** plays that role. There may yet be other roles that chances can play in explanations.

References

- Albert, David Z. (2000). *Time and Chance*. Cambridge, MA: Harvard University Press.
- Batterman, Robert W. (1992). "Explanatory Instability." *Nous* 26: 325–348.
- Hempel, Carl G. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: The Free Press.
- Hitchcock, Christopher (2010). "Probabilistic Causation." In *The Stanford Encyclopedia of Philosophy*.

- Humphreys, Paul (1989). *The Chances of Explanation*. Princeton: Princeton University Press.
- Lange, Marc (forthcoming). “What Makes a Scientific Explanation Distinctively Mathematical?” *British Journal for the Philosophy of Science* .
- Lewis, David (1986). “Causal Explanation.” In *Philosophical Papers*, volume 2, 214–240. New York: Oxford University Press.
- (2000). “Causation as Influence.” *The Journal of Philosophy* 97: 182–197.
- Loewer, Barry (2001). “Determinism and Chance.” *Studies in the History and Philosophy of Modern Physics* 32: 609–620.
- Railton, Peter (1978). “A Deductive-Nomological Model of Probabilistic Explanation.” *Philosophy of Science* 45: 206–226.
- (1981). “Probability, Explanation, and Information.” *Synthese* 48: 233–256.
- Salmon, Wesley C. (1985). “Scientific Explanation: Three Basic Concepts.” In P. D. Asquith and P. Kitcher (eds.), *PSA 1984*, volume 2, 293–305. East Lansing, MI: Philosophy of Science Association.
- (1989). *Four Decades of Scientific Explanation*. Pittsburgh: University of Pittsburgh Press.
- Schaffer, Jonathan (2005). “Contrastive Causation.” *The Philosophical Review* 114: 297–328.
- (2007). “Deterministic Chance?” *The British Journal for the Philosophy of Science* 58: 113–140.
- Scriven, Michael (1959). “Explanation and Prediction in Evolutionary Theory.” *Science* 130: 477–482.
- Strevens, Michael (1998). “Inferring Probabilities from Symmetries.” *Nous* 32: 231–246.

— (2000). “Do Large Probabilities Explain Better?” *Philosophy of Science* 67: 366–390.

— (2003). *Bigger than Chaos: Understanding Complexity Through Probability*. Cambridge, MA: Harvard University Press.

— (2008). *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.

Woodward, James (2003). *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.