Philosophical Studies (2026) 183:481–510
https://doi.org/10.1007/s11098-025-02454-1

There is no such thing as a statistical explanation

James Woodward[1]

Received: 9 July 2024 / Accepted: 12 November 2025 / Published online: 27 November 2025

Abstract
This paper argues, in contrast to what is usually assumed in the philosophical literature, that there is no such thing as a statistical explanation (hereafter SE). By SE I mean a putative explanation in which the occurrence of an individual outcome e or a collection of these is claimed to be explained by the ascription of a probability p to the outcome or collection, where p may be less than 1. Different accounts of explanation have different implications for the existence of SEs, with the no SE position supported by a dependency or what-if-things-had-been-different account of explanation and other accounts of explanation being friendlier to SEs. Recent attempts to arguing for the existence of SEs via an appeal to inference to the best explanation (IBE) are unsuccessful. Moreover, contrary to what a number of writers claim, IBE is ill-suited to serve as a basis for probability ascriptions.

Keywords Statistical explanation · Dependency account of explanation · Inference to the best explanation · Relation between explanation and evidence

1 Introduction

This paper explores a set of issues having to do with "statistical explanation" (hereafter SE). By this I mean a putative explanation in which the occurrence of an individual outcome *e* or a fact about a collection of these is claimed to be explained by the ascription of a probability *p* to the outcome or collection, where *p* may be less than 1. It is typically assumed in the philosophical literature that *p* is ascribed on the basis of some generalization concerning the probability of events of some kind K to which *e* belongs, although one might also imagine a version of SE which does not require this, and instead employs a singular claim specifying the probability *p* attaching, so to

✉ James Woodward
jfw@pitt.edu

[1] University of Pittsburgh, Pittsburgh, USA

speak, directly to *e*. (The probability of heads on a coin toss is 0.6.) Thus the standard form of the proposed explanatory schema is something like.

(S): (1.1) Events of kind K have probability *p* of occurring

(1.2) *e* is an event of kind K

(1.3) *e* occurs

where the dotted line indicates that (1.1) and (1.2) if true, explain (1.3).

Alternatively, the proposed explanatory schema may be something like:

(S*) *e* has probability *p* of occurring

----

e occurs

A standard example of schema S which I will label R for future reference is:

(R) All radium atoms have a probability *p* of decay within time interval dt.

a is a radium atom

----

a decayed within time interval dt

Discussions of SE divide on the question of whether, for successful explanation, the probability *p* must be "high" or whether low values for *p* also explain. Following Strevens (2000) I will call the former view "elitist" and the latter "egalitarian". There are other possible positions—perhaps SEs that assign high probabilities provide better explanations than those that assign lower probabilities, but the latter also explain, albeit less well. (This is Strevens' own view which he calls "moderate elitism".) Perhaps an SE explains an outcome if it renders it more probable than some or all alternative outcomes that might have occurred. Perhaps an outcome may be explained by citing a factor that increases its probability substantially with respect to some baseline (that is, it is the *change* in probability that matters). And so on. In what follows I will mainly focus on the contrast between elitist and egalitarian SE; explicit treatment of the other possibilities will not be necessary since my arguments imply that there is no such thing as SE of *any* form.

A very large majority of those who have addressed the issue have agreed that there is such a thing as SE, although there is considerable disagreement about the criteria SEs must meet. Defenders the claim that there are SEs include Hempel, Salmon, and more recently Strevens and Clatterbuck, among others. Very few (e.g., Kitcher, 1989; Watkins, 1984; Woodward, 1989) explicitly reject the possibility of SE.

I have described SEs as "putative" explanations because one of my principal claims is that we have no good reason to regard them as genuine explanations. More bluntly: there is no such thing as statistical explanation of individual events of the form described by S or S*. In claiming this I do *not* mean that quantum mechanics and other theories that assign non-trivial probabilities to individual outcomes are unexplanatory—on the contrary. Rather I claim that what such theories explain (when the occurrence of a particular outcome is not entailed by the theory and assumptions about initial and boundary conditions) are the *probabilities* of outcomes (or related features such as expectation values, transmission coefficients, variances and so on that are characterized with reference to a probability distribution over outcomes), rather the occurrences of individual outcomes themselves. For example, quantum mechanics explains the probabilities of alpha decay and similar phenomena by assigning these probabilities, but it does not explain individual decay events. Information about the composition of a coin and the circumstances of its tossing can explain why it has probability *p* of landing heads (see Keller, 1986, also Engel, 1992) but this fact about probability does not explain why the coin lands heads on a particular toss, and this is so whether *p* is high or low. Moreover, probabilistic theories like quantum mechanics also do not explain relative frequencies of outcomes such as why most radium atoms within a certain time interval have decayed even if the probability of this happening is high. Nor do they explain why these frequencies fall within some interval. I will say more about this in Sect. 8, but basically this is because the strongest conclusions about frequencies that follow from claims about the probability of decay are (following the various laws of large numbers) claims about the *probability* of various frequencies occurring or the probability of these falling within some interval where this probability is less than 1 for any finite frequency. Thus, in claiming that facts about relative frequencies are the explananda of statistical explanations, one is still relying on the idea that such outcomes are explained by assigning them a probability less than one. This is still a version of statistical explanation in the sense under discussion—some fact about relative frequencies is claimed to be explained via appeal to information about the probability of that fact occurring. And, as argued below, once this claim is made, it is hard to see how to avoid the conclusion that individual events can also be explained by being probabilified to the same degree.

2 Motivation: connections with other issues concerning explanation and evidence

The issue of whether there are statistical explanations of individual outcome is interesting in its own right—particularly given the consensus (which I claim is mistaken) that there are such explanations. But it is also entangled in interesting ways with other issues in the theory of explanation as well as issues having to do with how explanation connects to evidential support, including the status of so-called Inference to the Best Explanation—as I illustrate, this last is particularly prominent in recent discussion. In this section I allude briefly to some of these connections, as a way of providing an overall guide to what follows.

(**2.1**) As has long been recognized, issues about the status of SE are bound up with more general issues about the criteria for successful explanation. If one thinks that explanation is fundamentally a matter of providing grounds (perhaps nomically based) for expecting that an explanandum will obtain, then, as Hempel (1965) argued, an elitist version of SE will seem plausible. If instead, one thinks of explanation as having to do with showing the extent to which to which an explanandum is expectable (where this is successfully achieved as long as the expectation is accurate, regardless of the probability value involved) this is congenial to an egalitarian version of SE such as Salmon's SR model (Salmon, 1971). Similarly, if one thinks of explanation as just a matter of subsumption under a pattern of some appropriate kind and takes the relation between a particular outcome and a generalization specifying the probability of this outcome to be one of subsumption,[1] then one will also be sympathetic to some form of SE. On the other hand, as I will argue, the ideas just described contrast fundamentally with an alternative way of thinking about explanation—that this has to do with the exhibition of dependency relations between explanans and explanandum and the successful answering of what-if- things- had-been-different questions (-w questions) in the sense described in Woodward (2003). This is a view according to which explanations work by citing difference-makers. As I will argue, the most natural way of understanding such dependency accounts lead to rejection of the claim that there is such a thing as SE. Moreover, there are good reasons to prefer dependency accounts to the alternatives described above. At the same time there are also independent reasons—that is, independent of the general considerations having to do with explanation just described—for rejecting SEs (see 2.3). This in turn provides grounds for skepticism about treatments of explanation that vindicate SEs.

Of course, some who are convinced that there are SEs will take the tension between dependency accounts of explanation and the existence of SEs to provide grounds for rejecting dependency accounts. But even for such readers, recognition of this tension should be noteworthy since it has not, to my knowledge, been made explicit hitherto.

The different views about explanation sketched above also have implications for a more specific issue, discussed by Salmon (1984), among others: can the same explanatory factor *E* can explain both explanandum *M* and an alternative explanandum inconsistent with *M*? Those who hold that there are low probability SEs of particular outcomes answer "yes". I will argue that reasonable versions of dependency accounts imply that the answer is "no".

(**2.2.**) Although, as noted above, different ways of thinking about explanation have different implications concerning the existence of SEs, these are not the only considerations relevant to this topic. We can also ask whether there is anything in scientific practice that corresponds to SE and whether we *need* a notion of statistical explanation to evaluate statistical theories with respect to their explanatory power—that is, whether we need the idea that statistical theories explain individual outcomes to capture claims that some such theories do better on explanatory grounds than others.

[1] By this I mean roughly the assumption that (1.1) in S subsumes the conjunction of (1.2) and (1.3), with the latter being regarded as an "instance" of a generalization of form (1.1).

I will argue for a negative answer to these questions. Despite frequent claims to the contrary, to understand and evaluate scientific theories invoking probabilities we do not need to interpret them as committed to the claim that there are SEs. Relatedly, there is an obvious explanation for the absence of SE-like notions in scientific practice: as I argue below, we would reach the same conclusions regarding how well theories that make probabilistic predictions explain if we took those theories to explain the probabilities of outcomes but not individual outcomes. Moreover, contrary to what is sometimes claimed, we don't need the idea that there are SEs to explain the warrant for probability assignments—we can understand this in terms of standard procedures for statistical testing. Thus, when it comes to the description of scientific practice or theory evaluation or probability assignments, the idea that there are SEs is an idle wheel. I emphasize that this is an additional argument against the existence of SEs that is independent of the argument described in **2.1**.

(**2.3**) As noted above, recent discussion of statistical explanation is bound up with issues having to do with inference to the best explanation (IBE) and with other related claims concerning the connection between explanation and evidential support. In particular several writers (e.g., Clatterbuck, 2020; Emery, 2015, 2017; Strevens, 2000, 2008) have appealed to supposed connections between explanation and evidence to support the claim that there are SEs. Consider the following three claims which I distinguish for purposes of later discussion:

> 2.3.1. (IBE) Suppose *e* is our candidate evidence and that (i) if *h* were true, it would provide the best explanation (from among some set of alternative potential explanations all of which are consistent with the available evidence) of *e* and that *h* is a "good enough" explanation of *e*.
>
> Then fact (i) (the potential explanatoriness of *h* with respect to *e*) provides a defeasible consideration in favor of *h*, a reason for belief in *h* etc.[2]
>
> 2.3.2. (EV/EX) *e* is evidence for *h* only if *h* explains *e* or, more strongly, is the best explanation of e. (evidenceexplanation link)
>
> 2.3.2 (EX/EV) *e* is evidence for *h* if *h* explains *e* or, more strongly, is the best explanation for *e*. (explanationevidence link)

Obviously, these are distinct claims. IBE claims that potential explanatoriness is a reason for belief in *h* over and above what would ordinarily be considered evidence for *h*. (EV/EX) links evidence to actual (not just potential) explanation and (EX/EV) links actual explanation to evidence. Despite this distinctness, there is a common thread linking reasons for belief and explanatory considerations and it is not surpris-

[2] Formulations may vary in the strength of the support that the potential explanatoriness of *h* is claimed to provide for *h*. However, for our purposes such differences won't matter.

ing to find writers moving back and forth among the different claims 1–3, as my discussion below illustrates.[3]

All three claims can be (and, I claim, have been) used to support the existence of SEs.

Consider the probabilistic hypothesis *h* that a coin has a fixed bias of 0.7 toward heads. The coin is tossed (independent trials) 1000 times and the very surprising result *e* is exactly 700 heads. On any reasonable theory of statistical testing, *e* is evidence for *h* or at least grounds for not rejecting *h* (if one is doing significance testing) and, moreover, grounds for rejecting some alternatives to *h*. If one assumes (EV/EX) it will follow, given this evidential role for *e*, that *h* must explain *e*—indeed this will follow even though (as in this case) the probability of *e*, given *h*, is quite small.[4] Thus (one might think) it follows that there are low probability SEs of facts about relative frequencies. And once this is concluded there seems no reason to resist the further claim that there are also low probability SEs of individual outcomes. If a fact about relative frequencies can be explained by an *h* by showing that this fact has low probability given *h*, then we seem committed to the more general claim that *h* can explain outcomes by correctly assigning them a low probability. If so, it also follows that *h* also explain individual outcomes to which it assigns low probabilities.

The argument just sketched may seem to support egalitarianism but there is a nearby argument that may seem to support elitism. Suppose that in the above example, the coin is tossed just once and the outcome e is (i) heads or (ii), alternatively, tails. It may seem that (i) is evidence that *h* (Pr (heads) = 0.7) is correct and that (ii) is evidence against *h*.[5] Again, assume a link between evidential support and explanation along the lines of (EV/EX). It follows that *h* must explain *e*—that is, that some version of SE is correct. Furthermore suppose, as a reductio, that egalitarian SE is right. Then it follows that *h* explains (ii) (tails) were it to occur just as well as (i). But then, according to IBE, the fact that tails occurs is also a reason for belief in *h*—presumably as good a reason as the occurrence of heads. Alternatively, it follows from the fact that *h* explains the occurrence of tails and (EX/EV) that (ii) is evidence for *h*. Since these conclusions seem obviously misguided, we have an argument for rejecting SE egalitarianism in favor of SE elitism. I suggest below that something like this argument can be found in Strevens (2000) and that a version of the previous argument for egalitarianism can be found in Clatterbuck (2020).

Sometimes when a set of premises apparently can be used to support inconsistent conclusions, a reasonable response is to reject at least one of the premises. I think this is the case here. I will not discuss IBE, EV/EX and EX/EV in their fully generality[6] but I will claim (as my remarks above suggest) that they are highly problematic when "explanation" is understood to encompass SEs or when we are talking about evidence

[3] For example, the classic discussion in Harman (1965) seems to claim that *all* judgments of inductive support involve IBE, which gets us something in the neighborhood of (EV/EX).

[4] The probability of getting exactly 700 heads in 1000 tosses, assuming that the coin has bias 0.7, is 0.0275.

[5] If this claim does not seem plausible to you, I agree. Genuine evidence regarding the bias of the coin requires tossing the coin many times. But many philosophers including advocates of SE do think that the outcome of a single coin toss can serve as evidence in the manner suggested. Here I'm just describing an argument based on this assumption.

[6] For a more general critical discussion of IBE, see Woodward (2025).

for statistical hypotheses. First, it is mistaken to suppose, as implied by (EV/EX), that evidential support for statistical hypotheses needs to be understood in terms of how well those hypotheses explain such evidence. Instead, information about frequencies can be evidence for a statistical hypothesis *h* or at least provide ground for accepting or rejecting that hypothesis (if one is a classical statistician) without it being true that *h* explains that evidence. This is reflected in the absence of any discussion of criteria for explanation in standard statistics texts devoted to hypothesis-testing, parameter estimation and so on. Relatedly, we don't need to assume that there are SEs or settle whether egalitarianism or elitism is correct to make sense of how statistical evidence can support or refute claims about probabilities. More generally, whether the probabilities ascribed by a statistical theory or hypothesis are high or low does not matter in itself for the acceptability of that theory or the goodness or badness of the explanations it provides. Instead, insofar as claims about probabilities enter into the picture, all that matters is whether those probabilities are objectively correct (empirically accurate, etc.) as shown by statistical tests and, to repeat, such assessments do not require a detour through explanatory considerations.

2.4. Although we don't need to invoke claims about SE to make sense of how evidence bears on statistical hypotheses, it is nonetheless true, as suggested in **2.1** that certain views about explanation fit better with assumptions like IBE, EV/EX and EX/EV. Suppose you hold that *h* is well supported by evidence *e* if and only if *Pr*(*e*/*h*) is high. Suppose you also hold, with Hempel (1965), that *h* explains *e* iff *h* shows *e* to be highly probable and *h* is a statistical law, on the general grounds that explanation is a matter of showing that there are nomic grounds for expecting an outcome. Then you naturally will think that there is a close connection between evidential support and explanation of the sort captured by EV/EX, EX/EV and IBE. Showing that *e* provides evidential support for *h* will be pretty much the same thing as showing that *h* explains *e,* assuming that *h* is a law, and a best explanation of *e* will appeal to a lawful *h* for which *Pr*(*h*/*e*) is high and thus will be a case in which *e* provides strong support for *h.*

In this "Hempelian" example, the connection between evidential support and explanation is "internal", following from the way in which these are conceptualized. By contrast, there is no such close or internal connection for an account that takes explanation to have to do with exhibition of dependency relations between explanans and explanandum along the lines described above. On a dependency account, General Relativity explains (*E*) why most objects released near the surface of the earth fall to the ground since it correctly tells us what such behavior depends on but it is not part of scientific practice to regard (*E*) as evidence for GR. One reason for this (built into the most plausible accounts of evidential support, including Bayesianism) is that evidence for a theory or hypothesis *h* needs to distinguish between it and alternatives to *h*. All gravitational theories that merit serious consideration, including Newton's, also imply *E* (and if true would explain *E*), so that *E* does not differentially support GR over these alternatives. By contrast, the classic tests of GR such as the deflection of starlight by the sun, do differentially support GR over various alternative gravitational theories and thus are appropriately regarded as evidence for GR. Roughly speaking, on a dependency conception of explanation and a conception of evidential support that focuses on the exclusion of alternatives, whether a theory (hypothesis etc.) *T* explains *M* has to do with the relationship between *T* and *M* and perhaps the relationship between T and other explananda M^*. By contrast whether E is evidence for T depends in part on the relationship between E and alternatives to T.

The example above suggests that explanatoriness does not always entail evidential support, contrary to EX/EV. It is also arguable that evidential support can sometimes be present without explanation, contrary to what EV/EX claims. Measurements of two masses and the distance and force between them, together with the Newtonian gravitational force law can provide good evidence for the value of the gravitational constant G but the above information does not explain why G has the value that it does. G, the inverse square law, the value of one of the masses m_1 and the distance to the second mass m_2 provides strong evidence for the value of the second mass, but this information does not explain why it has that value. (In both cases explanation requires citing factors on which the target explananda, G and m_2 depend.) As another illustration, if (e) measurements from my thermometer agree with the measurements of many other thermometers that are known to be well calibrated, this is evidence that (h) my thermometer is reliable, but this fact (h) does not explain why the other thermometers are well calibrated and agree with mine (e).

As I said above, *I* do not claim that these remarks amount to a full discussion of the status of EV/EX, EX/EV and IBE. They are intended to suggest that, depending on how one thinks of explanation and evidential support, the links between these need not be straightforward.

Moreover, it is not obviously anomalous to hold that facts about individual outcomes or frequencies can provide evidence for probability assignments without those assignments explaining outcomes or frequencies and that we need not think of such assignments as based on IBE.

Once we drop the assumption that assumptions like EV/EX and EX/EV are the only possible way of thinking about the relation between explanation and evidential support (including evidential support for ascription of probabilities), I claim (Sect. 8) that it follows that either (i) the assumption that there are SEs is unnecessary for the evaluation of statistical theories (this is the case for egalitarian SE) or, alternatively, (ii) that this assumption leads to leads to unacceptable results, as is the case for the simplest version of elitist SE. Moreover, if elitist SE is modified so that avoids these unacceptable results, it ends up, like egalitarian SE, being unnecessary for the evaluation of statistical theories.

3 Some additional background assumptions

Before turning to this argument, however, let me make some additional background assumptions explicit: First, I will assume, in accord with most current discussion, that if there is such a thing as SE, the probabilities that figure in such explanations must be "objective", "physical" probabilities in the sense that these have to do with features of the world that pertain to the behavior of the systems we are trying to explain rather than, e.g., people's degrees of belief about the behavior of those systems (or the degrees of belief they would have if rational). Relatedly what SEs claim to explain are also facts about the world—e.g., that an atom has decayed—rather than claims about why it is rational to expect it will decay. Probabilities in the sense of rational degrees of belief might figure in an explanation of (or, better, a justification for) why it is rational to expect certain outcomes, but this is different from explaining the occurrence of those outcomes, which is the goal of SE.

Second, I will assume that the objective probabilities that figure in statistical explanations obey the usual axioms of probability theory—probabilities take real values in the interval [0, 1], are countably additive, are measures defined on a sigma-field and so on. One consequence is that probabilities cannot be identified with either actual relative frequencies, or with the "hypothetical" relative frequencies that (allegedly) would obtain under an infinite number of trials.[7] This follows from the fact that relative frequencies—either actual or hypothetical—lack the mathematical structure possessed by probabilities but, more relevantly for our purposes, it also reflects the fact, noted above, that the strongest connections between probabilities and relative frequencies are the probabilistic relations given by the various laws of large numbers. Although probabilities cannot be *identified* with frequencies, I assume, in accordance with my remarks above, that information about frequencies can be *evidence* for claims about probabilities or can at least provide results that motivate the rejection or acceptance of claims about probabilities, with statistical methodologies, either classical or Bayesian, providing accounts of how this works.

Next, although I assume that the probabilities figuring in SEs must be objective, physical probabilities, I will not ascribe to them some of the features ascribed to "chances" in the recent philosophical literature. In particular, several recent accounts (e.g., Elliott, 2021; Emery, 2015, 2017) build into the notion of chance the claim that these have the role of "explaining" outcomes or frequencies of outcomes. Since I deny that there are SEs, I also deny that probabilities play this role: To assume a notion of probability that has this explanatory role is to beg the question of whether there are SEs. Here I simply follow the standard mathematical treatments which do not build into the notion of probability any connection with explanation or causation.[8]

Third, a point of some delicacy that is rarely addressed in the literature on SE: in the standard axiomatic treatment of probability, the fact that an outcome has probability zero does *not* mean that it is impossible for it to occur—it just means that the outcome belongs to a set having measure zero. For similar reasons that an outcome has probability one is consistent with its failing to occur. Given a well-behaved probability density function for the random variable X, the probability of X taking exactly some real value x will be zero, despite the fact that on any given occasion X must take some real value. It is a nice question whether advocates of egalitarian SEs think that

[7] See van Fraassen (1980) for additional discussion.

[8] In the literature on causal modeling and discovery it is standard to assume various connections between causal claims and probabilistic claims—the Causal Markov condition is a well-known example. (See, e.g., Pearl, 2000.) But this does not mean that probabilities themselves are treated as causes or quasi-causes or explainers—instead, a central theme in this literature is that it is important to keep probabilistic and causal claims distinct. See, e.g., Pearl (2000, p. 38ff.). For Pearl (and for me) causal notions should be understood in terms of responses to interventions and this is not something that can be defined in terms of purely statistical relationships. So-called probabilistic theories of causation of the sort that flourished in philosophy in the 1970s and 1980s conflated probabilistic and causal notions, roughly by construing $Pr(E/C)>Pr(E)$ as the claim that C *causes* an increase in the probability of E.

outcomes are explained when they have probability zero but that such events occur is inescapable consequence of the mathematics of probability.

Fourth, there is considerable controversy in the philosophical literature concerning whether "objective" probabilities can be ascribed to systems that (presumably) are deterministic at some appropriate level of analysis—coin tosses, roulette wheels and so on. However, discussions of SE regularly make use of examples involving such systems and I will follow this practice. Readers who are unhappy with this may substitute quantum mechanical examples for the deterministic cases I discuss.

Next, there is an ambiguity or unclarity concerning the role of probability in some (perhaps many) accounts of SE that is worth highlighting. One possibility is that the role of probability is to describe the *relationship* between the factors cited in an explanans and an explanandum. For example, in an SE of form (S) above, one might think of the explanans as including the information that *a* is a radium atom and the explanandum as the fact the decay occurs, with the the information about the probability of decay describing the relationship between this explanans factor and the explanandum, thus (supposedly) showing that this relationship is an explanatory one. This contrasts with another possible view according to which the fact about probability itself is an explanatory factor (part of the explanans), rather than something that describes the relationship between the explanans and explanandum. To use somewhat prejudicial language, the idea is that the probability contributes to making the explanandum-phenomenon happen—the probability itself has "umph" or "biff" or is "quasi-causal" in the sense that behaves in a cause-like way, whether or not the explanation that results is causal. Philosophers who think of probabilities in the context of SE as "powers" or "propensities" or as disposition-like (in analogy with the way in which the fragility of a glass allegedly explains its breaking) often seem to have this second picture in mind. In discussions of SE, it is often unclear which of these possibilities (or perhaps some combination of them) is in play. However, I think it plausible that earlier writers like Hempel and Salmon tend toward the first, relational picture—e.g., for Hempel this relation tells us whether the explanandum is expectable. By contrast some more recent writers (e.g., Emery, 2015) may have something more like the quasi-causal role for probabilities in mind. Arguably this second view fits better with the idea that probabilities are postulated to "explain" (via some form of IBE) individual outcomes or facts about relative frequencies.

4 Some examples and their consequences

In philosophy of science, it is always desirable to have real examples before us. I accordingly begin with a sketch of quantum mechanical example in which what is explained is a claim about the probability with which a kind of system will exhibit certain behavior. One of my goals is to contrast this explanation with the discussions in the philosophical literature of SE.

Consider (EX 4.1), in which the explanandum is the probability that a particle of mass m with kinetic energy E will penetrate a square potential barrier of width $2a$. The potential is $V(x)$ within the barrier and 0 outside of it. The explanans includes the Schrodinger equation (as a law) and information about the potential barrier and the kinetic energy of particle as initial and boundary conditions. Solving the Schrodinger equation for this system, leads, after some calculation, to an explicit expression for the approximate probability of transmission through the barrier:

$$|T|^2 = e^{-2\int_{-a}^{a}\sqrt{2m}/h\times[V(x)-E]}$$

(here h* is h-bar and the entire expression following $2m$ is under the square root).

In particular, in contrast to the classical case, there is a positive probability of transmission even if $E < V$ (non-classical barrier penetration). In this example the probability of barrier penetration is *derived* (modulo various approximations) from the Schrodinger equation, assumptions about the initial and boundary conditions characterizing the system of interest and the Born rule for obtaining probabilities from the square integral of the state vector. However, there is another feature of this analysis to which I wish to draw attention and which I claim contributes crucially to its explanatory power. This has to do with the fact that the explanation satisfies the what-if- things-had-been-different criterion (w-condition) briefly described in Sect. **2.1**. That is, the explanans identifies conditions such that variations or changes in those conditions would have led to a change in the explanandum. For example, the derivation enables us to see how the probability of barrier penetration would have been different had the potential been different or had the the kinetic energy of the particle been different. In this way it shows us how the probability of barrier penetration depends on these factors—that is, the explanation proceeds via the exhibition of dependency relations. Moreover, as described in standard textbooks, the Schrodinger equation, when combined with other assumptions about the Hamiltonians characterizing other sorts of systems, can be used to answer a range of additional questions about how those systems will behave under different conditions. For example, solving the equation for the behavior of a particle in an infinite two-dimensional potential well shows how changes in the depth of the well and the particle mass affect its behavior and allowable energy levels. One can also derive similar results for other model systems—for example, one can show how the probabilities of various behaviors of a quantum harmonic oscillator depend on (change in response to changes in) such factors as the particle mass and its angular frequency. In all these cases the common feature of using the Schrodinger equation in combination with assumptions about initial and boundary conditions to answer a range of w-questions is present.

We can describe this exhibition of dependency relations or satisfaction of the-what-if- things-had-been-different condition in a more general way. Suppose that we think of the explanandum in the above explanations as a claim that some variable E takes a particular value $= e$ (for example, a particular value p for the probability of barrier penetration). Then satisfaction of the w-condition criterion requires that there be a set of true counterfactuals connecting variations in the value of E with variations in the conditions or variables cited in the explanans X taking the following form:

(**W**) If X had been different in such and such a way, (e.g., $X = x_1$ rather than x_2) the value of E would have been different (where in this case E is a probability value or a particular probability distribution or some other quantity defined by reference to a probability distribution such as an expected value). When this condition is met, X is a difference-maker for E.

As argued in Woodward (2003), these counterfactuals should have an interventionist or non-backtracking interpretation, these being the sorts of counterfactuals that are suitable for capturing explanatory (including causal) relationships.[9]

As will become clear below (cf. also Woodward, 2003), **W** does *not* require that the factors cited in the explanans are nomically sufficient for the explanandum- phenomenon or that a description of the latter be deducible from the former. A singular causal claim of the form "*c* causes *e*" can sometimes satisfy the w-condition criterion even if the occurrence of *e* is not derivable from the occurrence of *c*. (This is the case, for example, when the counterfactual "if *c* had not occurred, *e* would not have occurred" holds.)

In contrast to examples like EX 4.1, SEs do not satisfy the w-condition or cite dependency relations, when these are understood in terms of **W** above. Returning to R, (Sect. 1) that some radium atom *a* has a probability *p* of decaying in some time interval or that all radium atoms have this feature is not information about what the occurrence of the decay depends on or what made a difference to the decay, at least under the most obvious interpretations of these requirements. R does convey information about a pattern in the behavior of radium atoms and perhaps invites us to see the decay of atom *a* as an "instance" of that pattern, but this is different from identifying a factor which made a difference for whether the decay occurred or on which the decay depends, which is what the w-criterion captures. In particular, there are no true counterfactuals of either of the following forms:

4.1) If the probability of decay had been different from *p*, decay would not have occurred.

4.2) If the probability of decay had been different from *p*, decay would have occurred.

One reason[10] for this is that even a very high value for *p* does not ensure that decay will occur and even a very low value does not ensure that decay will not occur. As noted above, this is true even if the probabilities in question are 1 or 0 given the usual

[9] It is an interesting question whether we should think of EX4.1 as a *causal* explanation, but whatever one's view about this, EX4.1 does identify factors which are difference-makers for its explanandum and on which its explanandum depends. The account of causal explanation in Woodward (2003) counts thus this as a case of causal explanation in a broad sense. As explained in Woodward (2018), there is a more general notion of explanation as involving difference-making and the exhibition of dependency relationships can be extended to non-causal relationships.

[10] An additional consideration is this: within an interventionist framework it does not seem to make intuitive sense to think of a probability value as the sort of thing which can be an object of intervention, a cause or a difference-maker. One can intervene on what a probability depends on—e.g., the mass distribution of a coin—but not, independently, on the probability itself. Probability is a modal notion and like other modal notions (e.g., possibility) it does not seem to be the sort of thing that can make something happen. (The fact that *e* is possible does not explain why *e* occurs.) However, I lack both the space to further explore this idea and the ability to make it precise.

measure theoretic understanding of probability. The falsity of these counterfactuals is reflected of our judgment that whether decay occurs does not depend on the value of *p* and that the value of *p* is not a difference-maker for whether decay occurs.

One might respond to this by searching for some alternative w-like condition that SEs satisfy. I think, however, that it is far more plausible to conclude that there is a real difficulty with fitting explanations of form (S) into an overall framework in which explanations work by citing dependency relations or difference-making information. In other words, if you think that there are SEs, you are likely thinking of them as explaining in some other way besides conveying dependency information—e.g., explaining in virtue of subsumption under a pattern. The question then becomes whether one of these alternative views of explanation is acceptable or whether instead (as I will argue in Sect. 6) there are reasons to privilege dependency accounts of explanation.

5 Singular causal explanations

To further explore these issues let us compare SEs of form S with another example which has figured prominently (but, I will claim, misleadingly) in discussions of SE.[11] Suppose, following Scriven (1959), that Jones has paresis *e* and that this condition is caused by untreated syphilis *s*. Assume, as is standard, that the probability of *e*, given *s*, is low,—e.g. 0.25 (following Salmon's exposition of this example, 1989, p. 49). Assume also, as Scriven and Salmon do, that only those with *s* develop *e*. It seems very natural to claim that.

(EX5.1) Jones untreated syphilis caused his paresis

is an explanation of why he developed paresis. Since Jones's paresis is an individual event and since that event has a probability less than one in the presence of *s*, it has seemed to many that this and other similar cases show that there must be such a thing as statistical explanation of individual outcomes. (This example is treated by Salmon, 1971; Strevens (2000), and Clatterbuck (2020), among others, as an example of an SE). In other words, the assumption of Salmon and others is that (EX5.1) works in the following way: (EX5.1) explains *e* which is an individual event and it does so by conveying the information that the probability of *e* is 0.25, given the presence of syphilis, so that what does the explaining is this probability assignment. That is, (EX5.1) is treated as conforming to schema S in the following manner.

(S1) Events of the kind K= paresis have probability 0.25 of occurring among those with syphilis

(S2) Jones's paresis is an event of kind K and Jones has syphilis

[11] Although I will not undertake a systematic survey, I think that a number of other examples that are cited in the philosophical literature as examples of low probability SEs are in fact cases of singular causal explanation that function like the paresis example.

(S3) Jones has paresis

Thus, Jones's paresis (S3) is explained by (S1) and (S2) because these premises are true and the correct probability (0.25) for S3 follows from them. Moreover, if (EX5.1) is accepted as a legitimate example of SE, this seems to show that SEs need not conform to a high probability requirement.

In my view, these last two inferences are mistaken. Although (EX5.1) is a genuine explanation on a dependency or w-condition account, it wrong to treat it as an SE. Thus accepting (EX5.1) as a genuine explanation does not show that there are SEs. The w-condition requirement, understood as described above, is the key to recognizing the differences between explanations like (EX5. 1) and SEs. (EX5.1) conforms to that requirement but as argued above, SEs do not.

We noted above that it is a background assumption to the example that the only cause of paresis is untreated syphilis. Given this assumption, the following counterfactual is true:

(5.2) If Jones had not suffered from untreated syphilis, he would not have developed paresis.

(5.2) *does* convey information about the conditions under which the explanandum phenomenon (the occurrence of paresis) would have been different—it says there would have been no paresis is the absence of untreated syphilis.[12] Thus, the w-condition requirement is satisfied by (5.2).

On this analysis, the truth of (5.2) is central to why EX 5.1 is explanatory. For purposes of comparison, suppose that even in the absence of s, Jones has probability $p=0.1$ of developing paresis if some alternative cause c_1 is present. Now suppose that s is present along with c_1 and that the presence of s increases the probability of paresis to, say, 0.35. My view is that in such circumstances and in the absence of further information the counterfactual (5.2) is false since, given what we know, we cannot exclude the possibility that paresis might have occurred even in the absence of s. Accordingly, in these circumstances we cannot conclude that Jones' paresis is caused or explained by his untreated syphilis. Jones' paresis *might* be caused or explained by his untreated syphilis, but it might not be, his paresis instead being due to c_1.[13]

We noted earlier that causal explanations work at least in part by citing difference-makers for their effects. EX5.1 is an example of this. Even though Jones' untreated latent syphilis is not sufficient for his paresis, it is what made a difference to his developing paresis, as reflected in the truth of the counterfactual (5.2). Again, this contrasts

[12] Compare (5.2) with

(5.3) If Jones had untreated syphilis, then Jones would have developed paresis.

The Lewis-Stalnaker account of counterfactuals treats (5.3) as true but my view is that when interpreted as an interventionist counterfactual ("If an intervention were to cause Jones to have untreated syphilis, he would develop paresis") (5.3) is false, even if Jones in fact develops paresis. Nonetheless the the truth of (5.2) and other associated counterfactuals is enough to ensure that (EX5.1) satisfies the w-criterion.

[13] You may be tempted to claim that in this case both s and c_1 cause (and explain) e. However, as explained in Woodward (2021) there are compelling reasons to reject this suggestion, since it is inconsistent with plausible (invariance-based) assumptions about how causes that operate probabilistically behave.

with explanations of form S in which the cited explanatory factor—the probability *p*—is not in fact a difference-maker. EX 5.1 is a form of *causal* explanation in which the difference-maker is a a cause (syphilis) that operates probabilistically. It is not an explanation in which the explanatory factor is a probability.

In addition, there are several other features of a singular causal explanation like EX5.1 that distinguish it from SEs and that will be relevant to our subsequent discussion. First, singular causal claims exhibit a kind of asymmetry with respect to what they explain that is not present in SEs, whether or not these are understood as obeying a high probability requirement. Compare Jones, whose syphilis explains his paresis, with Smith who also has untreated syphilis but does not develop paresis, an outcome which has probability 0.75 of occurring. On an egalitarian version of SE, both Jones's paresis and Smith's failure to develop paresis are explained by this information about their untreated syphilis and the cited probabilities.[14] Indeed, both outcomes are explained equally well. On an elitist version of SE, Smith's failure to develop paresis is explained (assuming 0.75 counts as a high enough probability) but not Jones' development of paresis.

Both versions of SE thus imply assessments that we ordinarily think of as mistaken: we think that the Jones' syphilis explains his paresis, but Smith's syphilis does not explain his non-paresis. This reflects the fact that *s* has an asymmetric explanatory relation with respect to paresis: *s* explains the occurrence of paresis but not its non-occurrence. Moreover, this asymmetry does not seem to track the probability values in the examples—we think that *s* explains paresis whether or not the probability of paresis, if *s* were present, is high, but we also think that *s* does not explain the non-occurrence of paresis even if that probability given *s* is high. Indeed, it seems that as long as it is true that *s* causes paresis and *s* is the only cause of paresis, and Jones has *s* and paresis, we can appeal to these facts to explain why Jones has paresis even if we don't know anything about the probability of paresis in the presence of *s*[15] and even if that probability is not cited in the explanation we provide. In other words, the value of the probability of paresis in the presence of *s* does not seem to contribute to the explanatory import of EX5.1 at all. By contrast the various philosophical accounts of SEs take a specification of that probability to be a crucial part of the explanation.[16] Moreover, accounts of SE either do not impose an explanatory asymmetry at all or impose a very different asymmetry from the one associated above with singular causal claims. Egalitarian accounts mistakenly treat both the occurrence or non-occurrence of the outcomes in the paresis case symmetrically with respect to whether they are explained. Elitist accounts treat the occurrence or non-occurrence of the outcome asymmetrically but are mistaken in automatically taking the outcome with a high probability (if any) to be the one that is explained. Again, that these consequences are mistaken is a reflection of the fact that, correctly understood, what

[14] This contravenes Salmon's Principle 1, discussed below,

[15] Recall that it is possible for paresis to occur even if its probability is zero. Even if one rejects this claim, all that is required is that the probability is non-zero.

[16] Even an egalitarian account, like Salmon's SR model, takes the specification of correct probability values to be a crucial part of the explanation.

does the explaining in EX5.2 is the fact that Jones' syphilis *causes* paresis and not the value of the probability *Pr* (paresis/syphilis).

One concern readers may have is whether the w-condition is unnecessarily strict and whether some loosening of this condition might still capture the claim that explanations must capture dependency relations, while permitting SEs. One possibility that has been suggested to me appeals to a "might" counterfactual: when *C* and *E* obtain, *E* depends on (and is explained by) *C* as long as it is true that (i) if *C* had not obtained, *E might* not have obtained. Then (the argument goes) this condition permits statistical explanation: take *E* to the decay of an atom (it in fact decays) and *C* to be the probability *p* of decay. The if *C* had not obtained (if probability *p* had been different or at least different and not equal to zero), *E* might not have decayed, and so in this sense the decay "depends" on *C*. However, this seems clearly inadequate as an account of "depends" that is relevant to explanation. In the example under discussion the following "might" counterfactuals are true: if *C* had obtained, *E* might have obtained, if *C* had obtained, *E* might not have obtained, if *C* had not obtained, *E* might have obtained, and if *C* had not obtained, *E* might not have obtained. This describes a situation in whether or not C obtains is *irrelevant* to whether or not *E* obtains, rather than a situation in which *E* depends on *C* or in which *E* is explained by *C*.

As another illustration, suppose that *E* is the (indeterministic) decay of an atom in Los Angeles within time interval *dt*. Let *C* obtain if it is sunny in Shanghai on the same day. Suppose both *C* and *E* hold. Then if *C* had not obtained, *E* might not have obtained, but of course it does not follow that *C* explains *E* (or that *E* depends on *C*), even if this might counterfactual is interpreted along interventionist or non-back tracking lines.[17]

Although I cannot survey all possible weakenings of the condition **W** that would permit SEs, I see no reason to believe that there is any such weakening that can be formulated *only* in terms of "might" counterfactuals and that successfully captures an explanation-relevant notion of "dependence". Instead, capturing this notion requires the holding of at least some, determinate, non-might counterfactuals, as in **W**.

Several other consequences follow from these observations. First, it is a mistake to think (as is sometimes claimed) that rejection of the possibility of SE commits one to a rejecting the possibility of explanation of individual outcomes under indeterminism. Given our analysis of how (EX5.1) works, it would still count as an explanation even if (contrary to what I assume is the case) the relationship between untreated syphilis and paresis is irreducibly indeterministic. Second, it should be clear from our treatment of (EX5.1) that rejection of the claim that there are SEs of form S does not commit us to the "deductivist" claim that in all genuine explanations the occurrence of the explanandum must be deducible from the conditions cited in the explanans. (What Kitcher, 1989 calls "deductive chauvinism"). In (EX5.1) the explanandum—

[17] Another proposal from the same commentator: *C* explains *E* (and *E* depends on *C*) if had *C* not obtained (or had *C* been different in a particular way), then *E*'s *chance* would have been different. One obvious response is that this is a condition that captures when the chance of *E* depends on *C*, rather than capturing when *E* depends on *C* and that it is the former but not the latter that is explained by *C*. Quite apart from this, consider the following example: *E* says that a (fair) coin is heads up on a table and *C* says the coin has not been tossed, where both *C* and *E* are true. If the coin had been tossed, the chance of the coins being heads up would have been different. But it does not follow that *E* depends on *C* or that *C* explains *E*.

occurrence of paresis—is not deducible from the factors cited in the explanans but (EX5.1) is still a bona-fide explanation. So arguments that if we reject the claim that there are SEs, we must be assuming that all explanations are deductive are also misguided.

Let me conclude this section by commenting on an influential claim of Salmon's (cf. 1984). He describes the following principle, which he calls Principle 1:

It is impossible that, on one occasion, circumstances of type *C* adequately explain an outcome of type *E* and, on another occasion, adequately explain an outcome of type *E′* that is incompatible with *E*.

Salmon rejects this principle, as of course he must, since he advocates an egalitarian version of SE: he holds that in, e.g., the binary case, one can explain both *E* and not *E*, when each occurs, by citing the same explanans *C*, and their probabilities of occurrence. By contrast, in my view any plausible version of a dependency theory of explanation must accept Principle 1. As noted above, it is hard to see what could meant by "dependency relation" if that relation can hold both between *C* and *E* and between *C* and not *E*—instead the obtaining of such a relation means that *E* does not "depend" on *C*, in any sense relevant to the presence of an explanatory connection. I thus conclude that acceptance of a dependency conception of explanation requires acceptance of Principle 1. Note also that EX5.1 (as well as EX4.1) respects Principle 1. Untreated syphilis can explain the development of paresis, but it does not also explain failure to develop paresis. Even though Salmon rejects Principle 1, he acknowledges that it is quite intuitive. It is a point in favor of a dependency conception that it respects this principle.[18]

[18] An anonymous referee describes the following as a possible case in which the same explanans *S* explains both *E* and some incompatible *E*', thus showing that there are cases in which Salmon is correct in rejecting Principle 1:

Suppose that if you've never contracted disease *X*, it's impossible for you to contract disease *Y*. Once you've contracted disease *X*, however, you are susceptible to randomly catching disease Y from your environment. That being said, disease *X* sometimes has the effect of directly causing you to contract disease *Y*, but on other occasions can give you immunity from catching disease *Y* in the future.

It thus seems that this is a case in which contracting disease X can both explain contracting disease Y and (via the indirect route) and explain not contracting disease *Y*. I see this example as having a more complicated structure than the examples involving alleged SEs that motivated Salmon's rejection of Principle 1. In particular, in contrast to Salmon's example, the present case is one in which in which there are two different pathways or routes from *X* to *Y*—a direct route by which infection with *X* sometimes causes infection with *Y* and an indirect route by which infection with *X* sometimes causes immunity *I* which in turn is (negatively) causally relevant to infection with *Y*. If we switch the representation a bit and let *X*, *Y* and *I* be variables (so that *X* can take two values, disease present or disease absent and similarly for *Y* and *I*, the example is a "triangular" structure in which there is an arrow from *X* to *Y*, an arrow from *X* to *I* and an arrow from *I* to *Y*. To bring out the relevance of the W condition to examples of this sort we need to follow a procedure that is widely recognized in the causal modeling literature, which involves intervening to "wiggle" some variable while holding others fixed so as to see the effect of the former variable on some target. For example, if we intervene to fix *I* at some value (e.g., no immunity) and wiggle whether disease *X* occurs, this will make a difference for *Y*, and the W-condition is satisfied. Salmon's Principle 1 is also respected since it is not true that the same value of *X* such as disease *X* occurs explains the occurrence and non-occurrence of disease *Y* once *I* is fixed in this way. Similarly, the non-occurrence of disease *X* does not explain both the occurrence and non-occurrence of disease *Y*. The relationships between *X* and *I* and between *I* and *Y* follow a similar pattern of respecting the W-condition and Salmon's Principle 1.

6 Are non-dependency accounts of explanation defensible?

So far one of my arguments has been that if there are SEs, they must be understood in terms of a conception according to which explanation does not to have to do with the tracing of dependency relations but instead has to do with providing other kinds of information. This should not be surprising since, as noted above, the philosophers who first introduced issues about the structures of SE relied on just such non-dependency conceptions.

This raises the general question of whether these non-dependency-based views of explanation are defensible. I take it that one of the lessons of recent discussion of explanation is that they are not. Counterexamples to the claim that explanation has to do just with providing grounds or for expecting or patterns of subsumption are legion. On the assumption that untreated syphilis *s* is the only cause of paresis, Jones' having paresis provides very strong grounds for expecting that he has *s*—$Pr(s/paresis) = 1$—but his paresis does not explain *s*. Similarly, that Jones, a male, is taking birth control pills provides strong grounds for believing that he will not get pregnant but his taking the pills does not explain this outcome.

This raises the obvious question of why we should accept SEs as genuine explanations when they apparently don't provide such dependency information and instead seem to rest on conceptions of explanation that are apparently subject to many counterexamples.

One possible response is to invoke some variety of pluralism about explanation: even if there are problems with non-dependency views (at least in connection with some examples), such views do seem to fit other paradigmatic examples of explanation and to that extent are acceptable—they capture some if not all features that we associate with explanation. Hence, if SEs don't fit well with a dependency framework, why not understand them in terms of one of these alternative frameworks? Going further one might wonder whether there is a need to connect SEs with *any* more general framework for thinking about explanation. Perhaps SEs are a *sui generis* form of explanation or involve some novel way of thinking about explanation, the details of which have not been worked out but which we should acknowledge because SEs are, on intuitive grounds, clear cases of explanation. (cf. Emery, 2017).

This line of argument is unconvincing for several reasons. First, it seems ad hoc. Second, I take it to be a common assumption among defenders of SE that the discovery of explanations is an important goal in science and elsewhere. (I share this assumption—more on this below.) However, in order for this goal to provide useful guidance there must be non-trivial constraints on what counts as an explanation and these constraints must lead to consistent judgments. If there are a number of acceptable theories of explanation with different and (apparently) inconsistent implications for the evaluation of various examples (as we have seen is the case for purported SEs), this threatens to undermine any principled basis for the assessment of the explanatory credentials of different hypotheses. I agree of course that the fact it would be desirable to have such a principled basis does not mean that one exists, but if there is no such basis, one wonders what the argument over, say, whether there are SEs or if there are, whether egalitarian or elitist SE is the correct account of them is about or how it might be settled in a non- arbitrary way. (Why not postulate two different forms of SE, corresponding to two different types of explanation, one elitist and the other egalitarian, each the correct reconstruction of different examples?) I suspect that the conclusion we should draw from strong forms of pluralism or laissez-faire about what counts as an explanation is that whether or not they provide "explanations" is not a very useful standard for evaluating scientific theories.

7 The assumption that there are SEs is not necessary for the evaluation of statistical theories or the correctness of probability ascriptions

I begin this section with the assumption that we have some way of determining (independently of IBE-based considerations and on the basis of standard statistical inference procedures, whether classical or Bayesian) whether probability ascriptions are correct or accurate. As noted above, one obvious motivation for this assumption is that standard statistical inference procedures do not seem to be based on explanatory considerations. For example, when the hypothesis that a coin is fair is rejected (or not) on the basis of a significance test, the rationale for this is based on the error characteristics of the test, and not on how well the observed evidence is explained by the bias of the coin. Although I regard this assumption as extremely plausible, I recognize that it may seem question-begging to some defenders of SE and so will relax it later in my discussion.

Consider first an egalitarian version of SE according to which (it is claimed) individual outcomes are explained by correctly specifying their probabilities of occurrence but according to which outcomes with low probabilities are just as well explained as outcomes with high probabilities. Now compare this with an alternative account (NO SE) according to which there are no SEs—individual outcomes are not explained by assigning them probabilities or subsuming them under statistical generalizations, although claims about the probabilities of such outcomes are explained if we have a theory implying such probabilities for which conditions like the w-criterion are satisfied, as in the quantum mechanical example considered earlier.

Let us compare egalitarian SE and NO SE for a theory *T* that makes accurate predictions about the probabilities of individual outcomes. Strictly speaking, SE and NO SE say nothing about how to assess theories that purport to explain probabilities (as opposed to individual outcomes) but presumably both can be consistently combined with whatever standards of assessment for such theories we think appropriate—that is, both can be combined with the claim that *T* explains the probabilities that it entails to the extent that it satisfies the w-condition and whatever other criteria are thought to be appropriate for assessing such explanations. In this respect there will be no disagreement between SE and NO SE about the explanatory merits of *T*. Of course, egalitarian SE holds that the probability ascriptions of *T* will, in addition, explain individual outcomes and NO SE denies this, but this difference does not carry any additional implications either for the assessment of *T* or for the probability ascriptions it implies.

We can see that there are no such additional implications simply by observing that because egalitarian SE holds that all SEs of particular outcomes are equally good, regardless of the probability values they ascribe as long as those values are accurate, it provides no basis for discriminating either among different candidate theories on the basis of the probability ascriptions they imply—we cannot say that, e.g., T_1 explains better than T_2 on the grounds that T_1 assigns high probabilities to some outcomes (and high probability outcomes are better explained) while T_2 does not.[19] To provide a basis for such discrimination we need a version of SE that discriminates among theories on the basis of the probability ascriptions they imply—for example, a version according to which high probability assignments are preferable. However, this is a version of elitist SE, rather than egalitarian SE. Thus, given the background assumption that probability ascriptions must be accurate, the claim of egalitarian SE that individual outcomes are explained does no additional work in evaluating either theories that imply probabilities (like quantum mechanics) or individual probability ascriptions themselves beyond what is provided by NO SE—both SE and NO SE reach exactly the same conclusions about which theories that imply probabilities are explanatory (and how well they explain), again assuming that such implications are accurate. Of course, egalitarian SE differs from NO SE in adding, as a kind of addendum, that individual outcomes are also explained, but this seems, so to speak, an idle wheel—it adds nothing to whatever standards are employed in assessing either theories or individual probability assignments.

Suppose instead we adopt an elitist version of SE according to which only outcomes that are assigned high probabilities are explained. Consider first a theory T in which the explanandum is a binary outcome variable E which takes values 1 and 0, corresponding to occurrence and non-occurrence and that the theory assigns a probability to such outcomes. Assume (as above) that we have some standard for judging whether the assigned probabilities are empirically accurate. If the probability T assigns to an outcome is empirically inaccurate, T is presumably not explanatory of that outcome, so let's suppose that T's probability assignments are accurate. Then if either $Pr\ (E=1)$ or $Pr\ (E=0)$ as assigned by T is sufficiently high, the occurrence of one these outcomes but not the other will be explained. On the other hand, if it happens that $Pr(E=1)=0.5$, neither outcome will be explained (or at least not explained well). Thus, a T predicting this probability will be unexplanatory of this explanandum or will fail to explain it well, even if this probability is accurate. In this respect the theory predicting intermediate probabilities will be explanatorily inferior (at the level of explaining individual outcomes) to one that accurately assigns more extreme values such as $Pr\ (E=1) = 1$. The latter will at least explain some outcomes, even if not all. Similarly, on elitist SE, a theory that accurately predicts that a six-sided die is fair will be completely unexplanatory at the level of individual events. More generally, to the extent that explanatory success or goodness is a reason for valuing a theory and we count explanations of individual events as contributing to explanatory success, elitist SE seems to imply that we should prefer theories that make correct extremal probability ascriptions as providing better explanations.

Going further, consider a theory like QM that makes different (and accurate) probabilistic predictions about a range of different systems. Suppose that when applied to system 1, QM correctly implies that some binary outcome such as tunneling through

[19] This one is Strevens' points in his 2000. But he takes this to be an argument for elitist SE while I disagree.

a potential barrier has probability 0.5 and also implies that in system 2 with different initial conditions tunneling has probability 0.8. According to elitist SE we apparently should conclude that QM is less explanatorily successful with respect to system 1 than to system 2, even if in both cases QM tells us exactly what these probabilities depend on—how the Hamiltonians of the systems in conjunction with initial and boundary conditions determine (via the same general inference pattern) these probabilities and so on. These assessments seem odd but more importantly they seem pointless—in assessing a theory that makes extremal vs intermediate probability ascriptions it seems as though (as far as the ascriptions themselves go) the only thing that should matter is the accuracy of those ascriptions. For example, the fact that quantum mechanics predicts a probability of barrier penetration for system 2 of 0.5 is not a point against that theory, showing it to be explanatorily deficient, as long as that probability is accurate.

So far we have been assuming (for the purpose of assessing SE) that there is some standard for the evaluation of probability assignments which is independent of explanatory considerations of the sort associated with IBE. Can the defender of either egalitarian or elitist SE respond by arguing that there is no such standard–that instead the judgment that certain probability assignments are accurate must itself be based on explanatory considerations (we infer to those probability assignments based on IBE), so it is question-begging to suppose that we already know what the accurate assignments are in assessing the explanatory claims of egalitarian or elitist SE?

It is hard to see how this argument can be made to work. Let's focus first on egalitarian SE. Since, according to that doctrine, all probability assignments explain equally well, we cannot use IBE in conjunction with SE to argue that any particular probability assignment is supported over another on the basis of explanatory considerations. If egalitarian SE is correct, the assignment of probability = 0.1 to an outcome would (if true) explain that outcome well and hence might be claimed to be well-supported on the basis of IBE, but of course the same consideration might be invoked in support of any alternative probability assignment. So if the only basis for probability assignments is that we infer from how well such assignments would explain if correct to the conclusion that the best explaining assignments are correct, egalitarian SE provides no grounds for one assignment over another. Thus, the defender of egalitarian SE needs some other basis for probability assignments besides an appeal to IBE but this is just to acknowledge, as argued above, that what really matters is simply whether those assignments are accurate or correct (where this is established in some alternative way). Again, the claim that individual outcomes are explained by probability ascriptions does no work and can be dropped without loss.[20]

[20] Compare this with the use of IBE in assessing deterministic hypotheses. In such cases, we presumably have some independent grasp on what it is for a hypothesis to be true or correct—IBE is thought to be warranted because it leads to true or correct hypotheses or at least provides reasons in support of such a hypothesis. If the use of IBE in connection with ascriptions of probability is understood along similar lines, then there also must be some independent standard for whether the probability ascriptions are correct, presumably involving standard statistical tests. But then the question becomes: why not just use that standard? It will either agree with what is recommended by the model of SE in combination with IBE that is adopted or not. In the first case, the appeal to IBE cum a version of SE appears redundant. In the second case, if we rely on SE/IBE we will make a mistaken ascription.

Parallel objections apply to any attempt by the defender of elitist SE to argue that there is no independent standard for probability assignments besides that based on IBE. This too seems to imply that we could never be in a position to determine on empirical grounds that probability assignments are mistaken, since the correct assignments must be just what the IBE/elitist SE argument says. To illustrate, suppose a coin is tossed ten times, resulting in six heads. Then an assignment of *Pr*(heads) = 0.9 if correct would provide a better explanation of the individual outcomes that come up heads than an assignment of *Pr* (heads) = 0.5 which, depending on the version of elitist SE adopted, either explains these individual occurrences less well or not at all. If we base probability ascriptions just on IBE (no other standard) we seem led to reason as follows: since the former *Pr*(heads) = 0.9 ascription would provide a better explanation of the individual outcomes if correct, we infer that it *is* correct. Of course, this seems completely unwarranted. In addition, we face the obvious problem of how to understand the fact that we can get evidence for low or intermediate probabilities. That is, presumably we can get empirical evidence that the tosses of a die are fair, even though this probability assignments does not, assuming elitist SE, explain individual outcomes and thus apparently cannot be justified via an IBE based just on individual outcomes.

8 Relative frequencies as explananda?

An obvious alternative strategy for the advocate of elitist SE at this point is to consider outcomes (other than individual events) that *are* highly probable and to claim that these are explained by probability assignments, with this, via some assumption like EX/EV, serving to provide evidential support for the assignments. In the case of the fair coin, such an explanandum might be taken to be, e.g., the fact that the relative frequency of heads in a substantial number of repeated tosses falls within an interval such that some relative frequency outcome within that interval is highly probable. For example, given a fair coin tossed 20 times in i.i.d trials, the probability that the number of heads will be between 2 and 18 is "high".[21] The argument would then be that the probability assignment *Pr* (Heads) = 0.5 explains, in accordance with elitist SE requirements, why the above relative frequency falls within the chosen interval. (As noted earlier, this only makes sense if there is some independent standard for the correctness of probability assignments.)

With this strategy we have moved a considerable distance away from the intuitions/assumptions that originally motivated elitist SE (or for that matter, any other form of SE). The relevant explananda for the probability assignments, at least when these are non-high, are no longer individual outcomes or even particular relative frequencies for such outcomes (which will usually be low for frequent trials) but rather facts about relative frequencies falling within intervals. Moreover, we face the problem of how these intervals are to be chosen. Do we say that the hypothesis *h*: *Pr* (H) = 0.5 explains very well why the frequency of heads in 20 tosses falls in the inter-

[21] Recall that we need to consider an interval because the probability of any particular relative frequency, even one that is "close" to 0.5 will be low for any substantial number of trials.

val [2,18], since that probability is high, that *h* explains less well why the frequency falls in the interval [6,14] and so on? And if we still want to use explanatory considerations to warrant probability assignments how exactly is that supposed to work, given this new conception of the explananda of SEs? If an assignment of probability 1/6 to each of the faces of a die can explain facts about frequencies falling within an interval by endowing these with high probability, how can elitist SE be used to distinguish (as some of its advocates clearly wish to) among statistical theories on the basis of the probabilities they assign? Any probability distribution will assign high probability to *some* interval of frequency outcomes. The resulting picture of elitist SE looks, insofar as it has implications for probability assignments, very much like what one would get from methods used in conventional statistical testing—significance tests, use of measure of goodness of fit and so on. Again, it seems as though one could just as well say that probability assignments are to be made on the basis of conventional statistical procedures and drop any assumptions about the role of explanatory considerations or SE as unnecessary.

9 Strevens on statistical explanation

In the following two sections I want to use two relatively recent discussions—one of which (Strevens, 2000) defends elitist SE and the other of which (Clatterbuck, 2020) defends egalitarian SE—to further illustrate some of the claims made in previous sections.

Strevens claims that we need to assume some version of elitist SE to make sense of scientific practice and in particular the acceptance of statistical mechanics (SM) in the latter part of the nineteenth century. As I understand his argument, it assumes that there is a connection between explanation and evidential support that looks something like the assumptions EV/EX and EX/EV identified in Sect. 2.

Strevens asks us to consider observations like the following:

(9.1) A sample of gas is confined by a partition to one part of a box. The partition is removed, and the gas diffuses uniformly throughout the box.

Strevens notes that SM confers a high probability on events like (9.1) and he also suggests that we think of (9.1) as evidence for SM. I take him to then argue as follows: By assumption EV/EX we can conclude that SM must explain (9.1). Since (9.1) describes an individual outcome, this shows that some form of SE is correct. Moreover, suppose (what is overwhelmingly unlikely) that we instead observe

(9.1*) the gas failing to diffuse.

If egalitarian SE is correct, SM would also explain (9.1*). But then, in virtue of the evidence-explanation link (EX/EV), (9.1*) would also be evidence for SM, in contrast to our judgment that (9.1*) would provide evidence against SM. Thus, we should conclude that some version of elitist SM is correct.

Some of the problems with this line of argument have already been noted. It is uncontroversial (and I take Strevens to agree) that, under the conditions described in (9.1) uniform diffusion is always (or virtually always) observed to occur and failure to diffuse is never (or virtually never) observed. Thus, on this basis we have very good reason to accept

(9.3) Uniform diffusion has a very high probability of occurring.

We need not think of acceptance of (9.3) as the result some explanation-based inference such as EX/EV, relying on the claim that there are SEs. In fact, sophisticated statistical tests do not seem required to establish (9.3); we can conclude (9.3) is correct just on the basis of what we observe without any detour through claims about whether there are SEs. (In virtually all known cases, the gas diffuses, we have no reason to suppose that our observations are unrepresentative, etc.) Note also that we may establish (9.3) through observation without having any idea what explains it (at least on a dependency conception of explanation), either in an individual case or in general.

To the extent that SM implies (9.3), that is a point in favor of SM. Moreover, any alternative theory *T* that implies that

(9.4) uniform diffusion has a low probability of occurring

should be rejected because (9.4) is false. If, as Strevens seems to assume, we are considering a rather generic version of SM according to which it and *T* are the only alternatives, it is clear why we should prefer SM to *T*. (9.3) is evidence for SM because it preferentially supports SM over the only alternative (*T*). Thus, we can explain why it was reasonable to accept SM without assuming that SM explains individual events like particular episodes of gas diffusion.

In addition, the difficulties described in Sects. 7 and 8 remain: in cases in which outcomes are not binary, to make sense of low probability assignments, consistently with elitist SE, one seems forced to take the relevant evidence (and what is explained) to be something like the fact F that some observed relative frequency falls with an appropriate interval, where F is rendered highly probable by the probability ascription. In addition to being conceptually awkward (since we have shifted away from individual outcomes), this strategy seems to collapse into the uncontroversial claim that the acceptability of SM turns on whether it makes empirically correct predictions about probabilities, a claim that has nothing to do with whether SEs are possible.

10 Clatterbuck on statistical explanation

Clatterbuck (2020) responds to Strevens by claiming that egalitarian SE provides a better account than elitist SE of scientific practice, focusing on Mendelian genetics. She points out that, assuming Mendel's law of segregation, if the brown eye allele (B) is dominant to the blue- eyed allele (b), a child of two brown-eyed parents who are heterozygous (Bb, Bb) has a probability 0.25 of having blue eyes. This probability assignment is well confirmed but, as she notes, according to elitist SE, it does not explain or (on a moderate version of elitist SE) does not explain very well the fact that the child has blue eyes. She argues, as I have, that *if* probability assignments are to be accepted or regarded as confirmed on the basis of how well they explain individual outcomes, elitist SE appears to imply that this probability assignment is not well supported, contrary to what is assumed in scientific practice. She infers from this that (i) egalitarian SE is correct and that (ii) in general, the standard for evaluating probability assignments is simply whether they are empirically correct, with no preference given to assignments that are high, independently of whether they are correct.

Of course, I agree with (ii) but, as suggested above, don't think that we need to assume (i) to reach this conclusion. As I understand her, Clatterbuck views (i) and (ii) as closely linked because she agrees with Strevens and many others who have discussed SE that probability values are (or should be) assigned on the basis of explanatory considerations. That is, she assumes that if the assignment of low probabilities to outcomes is warranted, this must be because those low value assignments *explain* the outcomes in question—hence that egalitarian SE must be correct. Her assumption that correct probability assignments track explanation (in her case in accord with IBE-type considerations) is apparent in the following passage:

> my egalitarian proposal, that what we are trying to do in IBE is to find the theory that will ultimately assign the true objective probabilities to our observations, can explain the explanatory advantage [of Mendelian genetics][22]

The idea that (i) we should adopt the theory that "assign true objective probabilities" is uncontroversial but this is not, contrary to what the passage quoted above seems to suggest, the same as (ii) egalitarian SE which is the claim that low probability ascriptions explain individual outcomes. As I have argued, we can accept (i) without accepting (ii) and the fact that Mendelian genetics assigns the true objective probabilities in this case is sufficient to explain its advantage over alternatives, without any need for (ii).

Before leaving this discussion of Clatterbuck it is worth noting a subtlety that may mislead us. Return to the example of the blue-eyed child, X, with heterozygous (Bb) parents. Egalitarian SE, as we have been understanding it, implies that X's blue eyes (or possession of a BB genotype) can be explained by observing that the probability of this outcome is 0.25. Note, however, that if we know X has blue eyes and we know her parents both have brown eyes, we know, assuming Mendel's laws, much more than the above probability ascription—we know (i) that both parents must have been Bb and (ii) that X must have received a B from each parent, with X's blue eyes

[22] Elsewhere she describes the following as "an independently plausible … connection between high probabilities and theory confirmation":

> if we pay attention to our total evidence and it is sufficiently large and probative, then the theory that makes the evidence most probable best explains it and is most strongly favored by it.

This also assumes a connection between explanation and evidential support but here the assumption seems to amount to the combination of (i) the likelihoodist assumption that the hypothesis h_i for which $Pr\ (e/h_i)$ is highest (among the various alternatives h_i) is best supported by *e* and (ii) the best supported hypothesis explains *e*. Again, one can accept (i) without accepting (ii).

following deterministically from (ii). Thus, we know the specific sequence of events that deterministically produced X's blue eyes. Arguably this *is* an explanation of why X has blue eyes, because the **W** condition is satisfied. Hence there is information in the example that provides an explanation of an individual outcome. However, that this is not an explanation of an individual outcome by reference to a low probability which is what is at issue with egalitarian SE.

11 Probability densities?

There is yet another puzzle which deserves mention and that arises for both the elitist and egalitarian versions of SE. So far, we considered mainly discrete probability distributions, involving random variables that take only a finite (or at least countable) number of values. But many theories that make statistical predictions do not take this form. Instead, they make use of integrable real-valued random variables that can be represented by probability density functions (pdfs). This is the case, for example, for many but not all of the statistical predictions of quantum mechanics. In general, a probability density function $f(x)$ for the random variable X will give the probability that the value of X falls within a certain interval of values for X: $Pr\,(a<X<b)=\int_a^b f(x)dx$. In the case of quantum mechanics, the probability might have to do with the probability that a decay occurs within a certain temporal interval or that the probability of the results of a measurement of an observable like position or momentum falls within some interval. This leads to some perplexities if one thinks that pdfs can figure in statistical explanations, as one should if one thinks there is such a thing as statistical explanation at all. For typical pdfs, the probability of X taking any value x (where x is some real number) must be zero. On the elitist version of SE, it follows that one cannot explain why X takes that particular value—so events or occurrences of this sort can *never* be explained. The most natural response to this difficulty is to follow a strategy like that sketched in Sect. 8 and to broaden one's conception of what counts as an individual outcome or event and allow claims that the value of X falls within some interval[23] to count as "outcomes" in the relevant sense, taking these claims to be the explananda of SEs provided by pdfs. But (again) without some further restrictions on what intervals are acceptable, triviality threatens for this version of elitism. For any value x taken by the random variable X, there will always be some interval such that the probability that $X=x$ falls within that interval is high enough to exceed whatever threshold is imposed by elitist SE for successful statistical explanation.

Since the underlying motivation for elitist SE is to discriminate among probability assignments with ascriptions of high probabilities being better from the point of explanation, the strategy under consideration obviously requires further restrictions if it is to fulfill this motivation. One apparently natural possibility would be to regard pdfs (or theories entailing these) that are very narrowly peaked (with most of the probability mass piled up in some narrow interval, so that it is highly prob-

[23] Presumably the "interval" needs to satisfy some "connectedness" requirement (that is, the target explanandum should not be something like X falls within the union of (1,2) and (4,6) but I will not pursue how this might be formulated or what its rationale might be.

able that values of the random variable will fall within this interval) as preferable on explanatory grounds. However, as far as I know, no one has explicitly proposed this and the difficulties described in previous sections remain: It is entirely possible, as an empirical matter, for the outcomes not to be distributed in this way, and if so, it seems clear that we should prefer the empirically correct pdfs to those that satisfy the sharp-peaked criterion And if some sharply peaked density is empirically correct, we should prefer it just on these grounds, so that again appeals to explanatory considerations are superfluous.

Suppose on the other hand, that one favors egalitarian SE. One then faces the issue (as it might be put) of how low to go in contexts involving pdfs. If low probabilities can explain, can probability zero ascriptions to individual outcomes resulting from a pdf explain or do explanatory probabilities have to be greater than zero, although they are allowed to be arbitrarily small? Consider the following exchange:

A: Why did *e* occur?

B: It had probability zero of occurring. That explains why it happened.

B's response doesn't exactly trip off the tongue, and this may lead the egalitarian to opt for the small but non-zero alternative. One then faces the question of what "small" means. Probability zero ascriptions to what are naturally regarded as individual outcomes are often empirically warranted (especially when this occurs in the context of a pdf) and are consistent with those outcomes being possible. On what basis do we decide that such outcomes are unexplained while other outcomes of low probabilities (perhaps including those within epsilon of zero for any epsilon you choose) can be explained? The fact that a commitment to egalitarian SE embroils us in such questions is one more reason to avoid it (as well as a commitment to its elitist cousin).

12 The denial that there are SEs is not counterintuitive

Despite the considerations advanced in this essay, I expect that some readers will respond to my rejection of SE (perhaps particularly its elitist version) with incredulity. Am I really claiming that one cannot explain why, e.g., an ice cube melts in warm water by appealing to the fact that this has an extremely high probability of occurring? Yes, that is exactly what I am claiming.

Here are a couple of additional considerations that help to make such claims seem less counterintuitive. First, I'm not denying that statistical theories like SM apply to individual outcomes. On my view, SM does explain something about individual outcomes—for example it explains

(12.1) why this individual ice cube has a high probability of melting.

My challenge for the defender of SE asks why, in view of the difficulties discussed above, we need to go beyond (12.1) in taking SM to also explain

(12.2) why this ice cube melts.

A second consideration derives from the difference between thinking of explanation merely as a matter of subsumption under a repeatable pattern and a dependency or difference-making conception of explanation. Suppose I observe an ice cube melting in warm water. I ask why

(12.3) this ice cube is melting.

Is it an explanation to be told that

(12.4) this melting has a very high probability of occurring.

My strong inclination is to think that this is no explanation at all. If I'm puzzled about why the cube melts, how does it help to tell me that this is overwhelmingly likely behavior for all similarly situated cubes? Why doesn't this just generalize my puzzlement? It seems to me that in asking for an explanation of (12.3), I'm not asking whether such melting behavior is common or regular or probable or even universal. Instead, I'm asking for information about what the melting behavior depends on. Some version of the standard SM story about almost all microstates of the ice cube consistent with its observed macrostate leading via any reasonable dynamics to a macrostate in which the cube melts (the latter hence being overwhelmingly probable) provides an outline of an answer to this question. Note, however, that this answer tells us what the high probability of melting depends on. It is true that if one thinks that explanation just involves subsumption under a repeatable pattern it may seem reasonable to suppose that (12.2) explains (12.1) but, as argued above, explanation requires more than this.[24]

13 Conclusion

I have argued that there is no such thing as a statistical explanation of individual outcomes. The assumption that there are SEs is not needed for the evaluation of statistical theories or hypotheses—not needed either for evaluation of their explanatory merits or their empirical accuracy. I have also argued that a recent strategy of arguing for the existence of SEs via an appeal to explanation-based considerations like EV/EX is unsuccessful. Issues concerning the existence of SEs are closely bound up with

[24] These remarks may elicit the following response: One needs to distinguish the (i) explanation of individual episodes of ice cube melting from (ii) the explanation of why ice cubes have an overwhelmingly high probability of melting. Of course, SM is required to explain (ii) but the high probability ascription by itself can explain (i). The argument given above neglects this distinction, mistakenly assuming that the obvious failure of the high probability ascription to explain (ii) is a reason to think that it also does not explain (i). For reasons of space, I cannot fully address this objection here. I will only say that although this two—level picture of explanation is deeply entrenched in philosophical discussion, I don't think it fits scientific practice. On my view there is no separate level of explanation in which individual events of ice melting are explained by reference to these being highly probable with their high probability being explained in turn by SM.

the merits of alternative theories of explanation and thus serve as one important test case for such theories. The NO SE position is supported by a dependency or what-if-things-had-been-different account of explanation and, to the extent that the NO SE position is independently plausible, this lends support to dependency accounts. At least in probabilistic contexts, and perhaps more generally, whether outcomes provide evidential support for a hypothesis is a very different matter from whether the hypothesis explains those outcomes.

Acknowledgments: Thanks to Nina Emery and Porter Williams for helpful comments on an earlier comments. Conflict of Interest: The author has no conflicts of interest, financial or otherwise. Funding: No sources of funding to acknowledge for this article.

Acknowledgments Thanks to Nina Emery and Porter Williams for helpful comments on an earlier comments.

Conflict of Interest The author has no conflicts of interest, financial or otherwise.

Funding No sources of funding to acknowledge for this article.

References

Clatterbuck, H. (2020). A defense of low probability scientific explanations. *Philosophy of Science, 87*, 91–112.

Elliott, K. (2021). Where are the chances? *Synthese, 199*, 6761–6783.

Emery, N. (2015). Chance, possibility, and explanation. *The British Journal for the Philosophy of Science, 66*(1), 95–120.

Emery, N. (2017). A naturalist's guide to objective chance. *Philosophy of Science, 84*, 480–499.

Engel, E. (1992). *A road to randomness in physical systems*. Springer.

Harman, G. (1965). The inference to the best explanation. *The Philosophical Review, 74*, 88–95.

Hempel, C. (1965). *"Aspects of Scientific Explanation" in aspects of scientific explanation*. Free Press.

Keller, J. (1986). The probability of heads. *American Mathematical Monthly, 93*, 191–196.

Kitcher, P. (1989) Explanatory unification and the causal structure of the world. In P. Kitcher & W. Salmon (Eds.), *Scientific explanation*. Minnesota studies in the philosophy of science (Vol. 13, pp. 410–505). University of Minnesota Press.

Pearl, J. (2000). *Causality*. Cambridge University Press.

Railton, P. (1978). A deductive nomological model of probabilistic explanation. *Philosophy of Science, 45*, 206–226.

Salmon, W. (1971). Statistical explanation. In W. Salmon (Ed.), *Statistical explanation and statistical relevance* (pp. 29–87). University of Pittsburgh Press.

Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton University Press.

Scriven, M. (1959). Explanation and prediction in evolutionary theory. *Science, 30*, 477–482.

Smith, G. (2014). Closing the loop: Testing Newtonian gravity, then and now. In Z. Beiner & E. Schliesser (Eds.), *Newton and empiricism* (pp. 262–351). Oxford University Press.

Strevens, M. (2000). Do large probabilities explain better? *Philosophy of Science, 67*, 366–390.

Strevens, M. (2008). *Depth: An account of scientific explanation*. Harvard University Press.

van Fraassen, B. (1980). *The scientific image*. Oxford University Press.

Watkins, J. (1984). *Science and scepticism*. Princeton University Press.

Woodward, J. (1989). The causal/mechanical model of explanation. In P. Kitcher & W. Salmon (Eds.), *Scientific explanation*. Minnesota studies in the philosophy of science (Vol. 13, pp. 357–383). University of Minnesota Press.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford University Press.
Woodward, J. (2018). Some varieties of non-causal explanation. In J. Saatsi & A. Reutlinger (Eds.), *Explanation beyond causation* (pp. 117–137). Oxford University Press.
Woodward, J. (2021). *Causation with a human face: Normative theory and descriptive psychology*. Oxford University Press.
Woodward, J. (2025). The place of explanation in scientific inquiry: Inference to the best explanation vs inference to the only explanation. *European Journal for the Philosophy of Science, 15*, 21. https://doi.org/10.1007/s13194-025-00649-1