Another Tristram Shandy-like effort from Woodward. Too long for journal publication. Posted because I think the orthodox views abou this matter are mistaken.

There is No Such Thing as Statistical Explanation[1]

James Woodward

I. 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 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. 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, with the probability p attaching, so to speak, directly to e. (The probability of heads on the next coin toss is 0.6) Thus the standard form of the proposed explanatory schema is something like

(S) : (1) Events of kind K have probability p of occurring
(2) e is an event of kind K
----

(3) e occurs

where the dotted line indicates that (1) and (2) if true, explain (3).

Alternatively, the proposed explanatory schema may be something like:

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

[1] Hard as it may be to believe, an ancestor of this paper was begun somewhere around 1983. (It formed the basis for a post-doc application at the University of Pittsburgh that year.) I worked on it a bit and then put it aside, in part to pursue other interest buts also because after a burst of attention to stastical explanation in the 1970s and 80s, the topic seemed to fade from interest. Recently, however, there has been a revival of interest (with virtually eveyone agreeing that there is such a thing as statistical explanation, but disagreeing about the requirements this must satisfy. This prompted me to revisit the ideas in this paper. I thought there was something to be said for getting my heterodox view (that there is no such thing as the explanation of individual outcomes by probability ascriptions) out on the table.

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, 2020 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 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 which at some points in my discussion will require separate treatment ; explicit treatment of the other possibilities will not be necessary since I am going to argue that there is no such thing as SE (understood as above) 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 bout the criteria SEs must meet. Defenders the claim that are SEs include Hempel, Salmon, and more recently Strevens, Clatterbuck , and others, Very few (e.g., Watkins, 1984, Kitcher, 1989, Woodward, 1989 have explicitly rejected 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 general 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 radioactive decay and similar phenomena by assigning these a probability but not 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 for the case of a fair coin) 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. Similarly the fact that if a fair coin is tossed a large number of times (in i.i.d trials) it is highly probable that the relative frequency of heads will be "close" to 0.5 does not explain why this outcome occurs.

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 many 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. This provides additional motivation for my discussion. In this section I allude briefly to some of these, 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 (that is how likely this is, where this can be more or less) this is congenial to an egalitarian version of SE such as Salmon's SR model. 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[2], then one will also be sympathetic to some form of SE. On the other hand, we will see that 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 the fact that there are a number of 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 reasons for skepticism about treatments of explanation that vindicate SEs.

The different views about explanation sketched above 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 the alternative explanandum not M? Those who hold that

---

[2] By this I mean roughly the assumption that (1) in S subsumes the conjunction of 2 and 3, with the latter being regarded as an "instance" of a generalization of form (1) (e.g. in the case of (R), a is a Ra atom and a decayed is an instance of 1*: all radium atoms have probability p of decaying within a certain time interval). Note, however, that the decay of a particular Ra atom is not literally an instance of (1*) , assuming we are thinking of an instance of (x) Fx-->Gx) as something of form Fa.Ga. Instead an instance of (1*) takes the form "radium atom a has probability p of decaying". On this understanding of "instance", and the assumption that instances are what is explained by generalizations, we don't get support for the existence of SEs.

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. I will argue for a negative answer to these questions. Despite frequent claims to the contrary in the philosophical literature, to understand and evaluate scientific theories invoking probabilities we do not need to interpret them as committed to the claim that there is such a thing as SEs. Thus we do not need to suppose that there are SEs to adequately describe scientific practice or for purposes of normative assessment. Moreover, 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. 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.

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). In particular a number of writers (e.g. Strevens, 2000, 2008, Clatterbuck, 2020, Emery, 2015, 2017) have appealed to IBE-like assumptions in support of the claim that there are SEs. One version goes roughly as follows. Consider the following two claims the conjunction of which I will label (2.3.1):

(2.3.1a) (EX/EV) e is evidence for h if h is the best explanation for e. (best explanation--> evidence)

(2.3.1b) (EV/EX) e is evidence for h only if h is the best explanation for e. (evidence--> best explanation)[3]

I will take IBE to be the claim that something like the following inference pattern is warranted:

(IBE) Suppose that 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

[3] Obviously one might endorse IBE without adopting 2.3.1b-- one might hold that when h if true would explain e, one may infer h, while agreeing that there are cases in which e provides evidence for h even though h does not explain e. However, classic discussions of IBE often seem to claim or assume that *all* inductive inference or judgments of inductive support involve the application of IBE, which gets us to 2.3.1b. (For example, something close to this claim is suggested in Harman, 1965). Moreover, as illustrated below, many of the philosophers who endorse SE seem to rely on something like (2.3.1b).

e, where e is known to obtain. Then we may infer that h is true or well confirmed or at least that e provides evidence for h.

(IBE) might be qualified in additional ways--for example, we might also require that h be a good or good enough explanation of e (rather than just better than alternatives.) However, for our purposes such qualifications won't matter.

To see how (2.3.1ab) might be used to motivate the claim that there are SEs, consider the probabilistic hypothesis h that a coin has a fixed bias of 0.7 toward heads. Suppose 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 (2.3.1b) 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 according to which this fact has low probability, why can't 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 (2.3.1b). 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). Now make the apparently natural assumption based on (2.3.1a)

(2.3.1c) if h explains e just as well as e' , then e and e' are equally good evidence for h

It then follows that (ii) is also evidence for h, contrary to what we have supposed, and hence that whether the outcome is (i) or (ii) has no evidential bearing on the correctness of h. Since this conclusion seems misguided, it seems that we have an argument for rejecting SE egalitarianism in favor of SE elitism.

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. My view is that both (2.3.1a) and (2.3.1b) are mistaken, as is the accompanying idea that evidential support for statistical hypotheses should be understood in terms of how well those hypotheses *explain* such evidence. To begin with the latter, information about frequencies can be evidence for a statistical hypothesis h or at least provide ground for accepting or rejecting that hypothesis (if

[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.

one is a classical statistician) without it being true that h explains that evidence. Neither classical statistics nor Bayesian treatments of evidential support need to be understood as tying evidence or grounds for acceptance/rejection to explanation along the lines of (2.3.1ab)-- a point that is reflected in the absence of any discussion of criteria for explanation in standard statistics texts devoted to hypothesis-testing and assessments of evidential support. Second, (as far as the argument hypothesized above for elitism goes) the claim that a single outcome (like heads on a single toss) can be serious evidence for or against a hypothesis about the bias of the coin does not seem generally accepted in scientific practice except perhaps in cases in which the bias is hypothesized to be very extreme[6]. But even if this claim about evidence is accepted, there is again no reason to tie it to claims about any version of SE in the manner described above. More generally, as I will argue, 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. Instead all that matters is whether those probabilities are objectively correct (empirically accurate, etc.) as shown by statistical tests and, to repeat, such assessments of correctness do not require assumptions about the existence of SEs. To the extent there is statistical evidence in favor of some hypothesis about the bias of a coin, we should accept or believe the hypothesis just on the basis of that evidence -- we don't need to appeal to considerations about whether high or low probability SEs explain or indeed to explanatory considerations at all. For example, if tails on a single toss is taken to be evidence against h (p=0.7) in the manner envisioned above (again not a position that I would recommend, but one that might be supported if, for example, one employed a likelihood ratio test comparing h against h' (p=0.5)), then we can reject h (or take there to be evidence against it) just on this basis-- we don't need elitist SE to reach this conclusion.

2.4. Although we don't need to invoke claims about SE (or IBE) 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 (2.3.1ab) and IBE than others. Suppose, for example, 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 (or perhaps just only if) h shows e to be highly probable, on the general grounds that explanation is a matter of showing that an outcome is to be expected. Then you naturally will think that there is a close connection between evidential support and explanation of the sort captured by (2.3.1ab). Showing that e provides evidential support for h will be pretty much the same thing as showing that h explains e. A similar observation holds for the DN model of explanation and an hypothetico-deductive (HD) account of explanation: If M is known to be true (the usual case when we are explaining) then the truth of M in conjunction with its DN-derivability/explainability from explanans S provides HD confirmation for S. In this respect, the close connection between DN explanation and HD confirmation licenses a kind of IBE: suppose that S is a potential DN explanation of M in the sense that S contain premises from which M is derivable and that would if true meet the other criteria for DN explanation. Since by assumption M is true, it will provide HD confirmation of S, hence boosting S from a potential explanation to something closer to an actual one (or at least a better supported one). A similar point holds if you make the "likelihoodist" assumption that P(e/h) is a measure of the evidential support that e provides for h and conjoin this with the corresponding assumption that P(e/h) is also a measure of how well h explains e, perhaps

[6] See also footnote 4.

making this additional assumption because you think that evidential support must track or mirror explanatory import in the manner suggested by 2.3.1ab.

By contrast, other accounts of explanation don't license such a close connection between explanation and confirmation or evidential support. In particular this is true for accounts that take explanation to have to do with exhibition of dependency relations between explanans and explanandum along the lines described above. Here providing evidence that the assumptions of a candidate explanation are true or correct is a very different matter-- subject to different requirements-- from showing that those assumptions if correct would successfully explain, in the sense of answering w-questions about a range of outcomes and showing what these depend on. If General Relativity is correct or nearly so, it would provide excellent explanations of a range of observed phenomena according to the w-condition/ dependency relation conception of explanation. GR is in this sense potentially highly explanatory. But on my view-- and the view that I take to fit with the dependency/w-condition conception of explanation -- this consideration about the potential explanatory power of GR does not itself provide good evidence that GR *is* correct or well -supported. For that independent evidence -- as provided by the various classic and more contemporary tests of GR (the deflection of starlight by the sun, the gravitational redshift, the advance of the perihelion of mercury etc.) not having to do with the merely potential explanatoriness of GR is needed.

In general on a dependency theory of explanation, evidential support and explanation don't necessarily line up or mirror one another in the way envisioned in 2.3.1ab. 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 this general fact (E) as evidence for GR. One reason for this, on my view, is that evidential support for a theory or hypothesis needs to distinguish in some way between it and at least some alternatives[7]. There are many alternative gravitational theories, including Newton's, that 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 whether theory (hypothesis etc.) T explains E has to do with the relationship between T and E and perhaps the relationship between T and other explananda E*. By contrast whether X is evidence for T depends in part on the relationship between X and alternatives to T.

One can also have evidential support without explanation. 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 even in conjunction with the gravitational force law, G does not explain why those masses have the value they do or

---

[7] As is apparent from careful studies such as Smith, 2014, Smith and Seth, 2020, assessments of evidential support in science have a very complicated structure and the most compelling support involves many mutually reinforcing considerations. I don't claim that discriminating between a candidate theory and alternatives is all that there is to evidential support, merely that this is one important component. As noted in Earman, 1992 this component is built into Bayesian treatments of evidential support as well as present in more straightforward examples of eliminative induction. As Earman and many others recognize, an important part of this undermining of alternatives methodology is a systematic effort at generating and searching among alternatives.

why they are a certain distance apart. (This may be fixed exogenously by an experimenter.) 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 (e) . Similarly, as suggested above, facts about frequencies can provide evidence for probability assignments without those assignments explaining frequency facts.

Consider an additional illustration of some of the subtleties associated with the relations between evidence and explanation. Suppose Jones develops lung cancer. If (hypothetically) it was the case that Jones' was a heavy smoker and this caused his cancer, this information would be a good explanation of his cancer. However we cannot conclude from this conditional that the correct explanation for Jones cancer was his heavy smoking. For starters we need independent evidence that he was a heavy smoker. And even this is not enough-- it is entirely possible, consistent with his smoking , that the explanation for is cancer is instead his exposure to asbestos, assuming that was also present. Indeed it is possible that the explanation for his cancer is asbestos exposure even if the probability of lung cancer conditional on heavy smoking is much higher than the probability of lung cancer conditional on exposure to asbestos[8]. Of course if I know only that Jones has lung cancer and must guess whether this was due to smoking or asbestos exposure and these are the only two possibilities, it is arguably reasonable to guess the former, but within a dependency conception of explanation, h's explaining e is different from whether e provides good grounds for guessing that h is present (or whether e is evidence for h) . According to the dependency conception, explaining Jones' cancer requires identifying what his

---

[8] If this contention seems puzzling, here are two supporting considerations. First, in some cases, it is possible to get evidence that is independent of the probabilistic information just described and that shows that cause c1 was not just present but efficacious in causing e and that alternative cause c2, although present, was not efficacious in causing e. If e is that Smith is dead and c1 the fact that he has been submerged under water without access to oxygen for several days, Pr (e/c1) =1. Suppose also that c2 -- that Jones has been shot in the head-- is also present. Pr(e/c2) is high but plausibly less than Pr(e/c1). Nonetheless, a pathologist may be able to establish (e.g., on the basis of whether there is water in Jones' lungs) that it was the gunshot rather than drowning that caused Jones' death. Thus the evidence required to discriminate between these two candidate causes for Jones' death goes well beyond the facts about the conditional probabilities Pr (e/c1) and Pr (e/c2) above and we can't infer that the explanation for Jones' death was the submerging rather than the gunshot merely on the basis of a comparison of these two conditional probabilities (or for that matter, a comparison of the unconditional probabilities Pr (c1) and Pr (C2). (Note also that Jones' prolonged submersion is excellent evidence that he is dead, but the latter does not explain the former-- another case in which 2.3.1b (Ev--> EX) is false.)

Second, and more generally, what is really of interest in this example is not probabilities like Pr (e/c1 is present) but rather Pr (e/ c1 caused (or causally explains) e). (These probabilities are *not* in general equal as the above example illustrates.) Since c1 can't cause e unless e occurs, this probability as well as Pr (e/c1 caused e) is equal to 1. We thus can't appeal to a difference between these two conditional probabilities to argue that one of these causal claims is better supported than the other. Similarly within a full Bayesian framework, what we are ultimately interested in is not Pr (c1 is present/e) but rather something more like Pr (c1 caused e/e*) where e* includes other evidence in addition to e such as absence of water in the lungs.

cancer depends on. Thus, in Jones' case and more generally, we need evidence that the proposed dependency relation *is* present and operating as claimed. Merely showing that *if* the dependency relation was present it would explain as well or better than alternatives does not accomplish this.

2.5. Distinctions among forms of explanation. Another issue raised by discussion of SE concerns an under -appreciated distinction between two possible forms of explanation: SE understood as above and *singular causal explanations* in which the outcome explained has probability of occurrence less than 1, given the cited cause. Consider Scriven's (1959) well-known example (call it EX1) , frequently cited in discussions of SE, in which Jones's paresis E is explained by his untreated syphilis S but in which only a minority (e.g. 0.25) of those with S develop E. EX1 is treated by many writers as a low probability SE and its apparent legitimacy as an explanation is thus taken to support an egalitarian version of SE. (The example is invoked not just by Salmon but more recently by Strevens and Clatterbuck, among others, as an example of an SE). I will argue below that this is a mistake: Although (EX1) is a legitimate explanation it is *not* an example of a SE as described above. In an SE, the probability assigned to the explanandum plays a crucial role in the explanation-- indeed in some recent versions this probability is understood as an explanatory factor (as indeed it must be if one thinks that probabilities are postulated to explain individual outcomes.) By contrast In EX1 what explains Jones paresis is his latent syphilis, not the fact that he has a certain (low) probability of developing paresis, given that he has syphilis. EX1 is a form of *causal* explanation in which the explanatory factor is a cause that operates probabilistically, not an explanation in which the explanatory factor is a probability. Moreover (EX1) satisfies the what-if -things had been different/dependency relation criterion while SEs do not and behaves differently from SEs in other respects as well. We thus cannot appeal to the legitimacy of EX1 in support of the claim that there is such a thing as SE.

3. Some More Background Assumptions

Before turning to details, some additional remarks about my background assumptions will be useful. 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 are real-valued, take values between in the interval [0, 1], are countably additive, are measures defined on a sigma-field and so on. One consequence of this is that probabilities cannot be identified with either actual relative frequencies, or with the "hypothetical" relative frequencies that would obtain under an infinite number of trials. 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 that the strongest connections between probabilities and relative frequencies are those given by the various laws of large numbers. Roughly what these laws tell us is that, under repeated i. i. d. trials characterized by a properly behaved random variable X, it is highly *probable* that the average value of X will be "close" to its expected value or *highly probable* that relative frequencies of values of X will be close to their assigned probabilities. The fact the connection between probabilities and frequencies is probabilistic in this way precludes any identification of probabilities with frequencies-- that is, "probability" is invoked in characterizing the relationship between probabilities and frequencies, precluding a reduction of the former to the latter .

Although probabilities cannot be *identified* with frequencies, I assume, as observed 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, for what I hope are obvious reasons, ascribe to them some of the features ascribed to "chances" in the recent philosophical literature. In particular, several recent accounts (e.g., Emery, 2015, 2017, Elliott, 2021) 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 against the position defended in this essay. For this reason, I will not talk about chances in what follows. Relatedly, I also deny that we should think of probabilities as postulated or "introduced" in order to play the role of explaining individual outcomes or facts about frequencies, as recently suggested by several philosophers (e.g. Emery, 2017)[9]. My probabilities are objective probabilities but without a connecton with explanation (or causation) built into them. In this respect I am simply following the standard mathematical treatments of probability which do not assume any kind of connection between probability and explanation and indeed do not appeal to explanatory considerations at all[10].

[9] To my mind, this is roughly like saying that we postulate the existence of the natural numbers in order to explain facts about the size of flocks of sheep, that we postulate the real numbers to explain facts about scalar quantities like mass, the existence of vectors to explain facts about electromagnetic fields and so on. Probabilities, real numbers, vectors and so on are parts of a general mathematical technology that we have available in formulating scientific theories but we don't postulate their existence to explain things in the way that, e.g., Pauli postulated the existence of the neutrino to save energy conservation in beta decay.

[10] In the literature on causal modeling and causal 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 it is thought important to keep probabilistic and causal claims distinct. See, e.g., Pearl, 2000 pp 38ff. For Pearl (and for me) causal notions should be understood in terms of responses to interventions-- this not something that can be defined in terms of statistical relationships. So-called probabilistic theories of causation of the sort that flourished in philosophy in the 1970s and 80s conflated probabilistic and statistical notions, roughly by construing Pr(E/C)> Pr (E) as the claim that C *causes* an increase in the probability of E.

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 real valued 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 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 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. I will note, however, that if accounts of SE apply only to quantum mechanical phenomena, this greatly limits their applicability -- advocates of this view are in effect claiming that there is a sui generis notion of explanation that applies only in quantum mechancis. My own view, which I will not try to defend here, is that ascriptions of objective probability to deterministic systems can be legitimate.

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 other 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, rather than something that describes the relationship between an explanans and explanandum. To use somewhat prejudical language , the idea is that the probability makes the explanandum happen or contributes to making it happen-- the probability itself (or some fact corresponding to it) has "umph" or "biff" . Philosophers who think of probabilities in the context of SE in a quasi-causal way-- as "powers" or perhaps "propensities" of some kind 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) the author has in mind. However, I think it plausible that earlier writers like Hempel and Salmon tend toward the first, relational picture-- the role of the probability ascription is to show that the relation between explanans and explanandum is an explanatory one. By contrast a number of more recent writers (e.g., Strevens, Emery) seem to have something more like the quasi-causal role for probabilities in mind. Arguably this is especially true of those writers who claim that probabilities are postulated to "explain" individual outcomes or facts about relative frequencies and/or contend that we infer the existence of probabilities via an inference to the best explanation[11]. Here it seems unavoidable that the probabilities themselves are thought of as explainers.

Finally, a brief remark about philosophical methodology: it is standard to compare philosophical accounts of explanation with "judgments" about particular cases, using the latter as a standard for assessing the former. When properly circumscribed I think that this procedure is legitimate and I will sometimes appeal to it[12]. However , I do not think that agreement with generally accepted judgments about particular cases is the only standard by which to judge an account of explanation or that the goal of such an account should just be to systematize such judgments. A theory of explanation should also make it clear why and what respects the discovery of explanations is a valuable goal in science and how explanatory considerations figure in theory assessment -- it should "do work" in giving us guidance about how to assess the explanatory credentials of different theories. Compare, for example, an account of explanation A that invokes feature X in order to fit certain judgments about particular cases with an alternative account A* which also fits judgments about particular cases just as well but which is not committed to X and is just as useful in comparing alternative theories with respect to their explanatory credentials. In such a case, I think it reasonable to conclude that there is no basis for preferring A to A*-- the feature X is dispensable in the sense that does no work other than capturing judgments some are inclined to make. The features introduced to capture SEs seem to me have exactly this status.

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 proposals in the philosophical literature about SE.

In this example (EX2) , the explanandum is the probability that a particle of mass m with kinetic energy E will penetrate a square potential barrier of width 2 a . 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 considerable calculation, to an explicit expression for the approximate probability of transmission through the barrier in terms of E, V and L.

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

(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

[11] As noted above, this is also true of those who construe the probabilities that figure in SEs as chances, building a connection with explanation into the notion of chance.
[12] For a (limited) defense of this procedure in connection with causal judgments, see Woodward, 2021.

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 explanans satisfies the what--if- things had been different criterion (w-condition) briefly described in Section **2. 1** with respect to its explanandum. 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 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 one dimensional potential well shows how changes in the length 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 of 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.

As another illustration consider the derivation of the expression for the mean free path L of a molecule in a gas conforming to the Maxwell- Boltzmann statistics

$$L = k_B T/\sqrt{2} p\sigma$$

This shows how a quantity L, -- the mean or expected value of the free path traveled by a molecule between collisions, thus a quantity defined with reference to a probability distribution -- depends on the temperature T of the gas, its pressure p, and the effective molecular cross-section σ (k_B is Boltzmann's constant). This shows us how L would change (would have been different) under changes in such variables as temperature and pressure.

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. In other words, the requirement is that there are true counterfactuals of the form

> (**W**) If X had been different in such and such a way, (e.g., X= x1 rather than x2) the value of E would have been different (where in this case E is the value of a probability or perhaps a particular probability density or some other quantity defined by reference to a probability distribution such as an expected value).

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

The w-condition captured by these counterfactuals is closely linked to a "difference-making" conception of causation and of explanation more generally. Focusing first on causation, we think that one mark of a causal relation is that, under the right circumstances, the cause makes a difference in some way to its effect[13]. "Making a difference" is naturally understood in terms of the idea that (in the right circumstances) variations in the state of the cause will lead to variations or changes in the state of the effect, which is just the w-criterion. It is an interesting question whether we should think of EX2 (construed as an explanation of a probability) as a causal explanation[14], but whatever one's view about this, it does identify factors (the kinetic energy of the particle encountering the potential barrier, the potential itself) which are difference-makers for its explanandum and (I claim) this is at least part of why we regard EX2 as genuinely explanatory.

As already intimated, yet another way of thinking about the w-criterion is that it captures the idea that a successful explanation involves the exhibition of dependency relations between explanans and explanandum-- a successful explanation should show that and how the explanandum phenomenon depends on the factors cited in explanans. Showing how the explanandum-phenomenon would have been different if the explanans factors had been different gives us information about such dependency relationships.

As will become clear below, satisfaction of this condition 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.)

By contrast, SEs do not satisfy the w-condition or cite dependency relations. Returning to R, (Section 1) that some radium atom a has a probability p of decaying in some time interval

---

[13] There are subtleties here involving cases of overdetermination and pre-emption, since these may seem to be counterexamples to the connection between causation (or explanation) and difference-making. One guiding idea in recent discussion, which I endorse, is that in such cases there is still a connection between causation and difference -making, but in order to "see" this connection we may need to modify and/or "hold fixed" various features of the system under discussion. For example in a case of symmetric overdetermination, the difference-making connection between one of the causes and the effect becomes apparent when the other cause is removed or held fixed at the value "absent". Preemption cases can be handled by holding fixed so-called "off path" variables at allowable values. Systematic rules for doing this are described in Hitchcock, 2001, Woodward, 2003 and many others. Once this is done, cases of overdetermination and pre-emption including those involving singular causal explanations may be seen to conform to a difference-making requirement.

[14] The account of causal explanation in Woodward, 2003 counts this as a case of causal explanation in a broad sense. As explained in Woodward, 2018, a more general notion of explanation as involving difference-making and the exhibition of dependency relationships can be extended to non-causal relationships.

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.

This is because 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 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. Again, note the contrast with explanations like EX 2-3 (with explananda understood as probabilities of barrier penetration and mean free paths) which do tell us what their explananda depend on.

Suppose, however, that you subscribe instead to some non-standard account of probability according to which an event of probability 1 "must" occur and an event of probability 0 must not. Might you then argue that the w-condition is satisfied after all with respect to the occurrence of the decay on the grounds that if p had been different in a very particular way (p=0) the decay would not have occurred? I don't think this works. A central commitment of the idea that there are SEs is that the accurate value of the ascribed probability p plays a central role in the explanation. Whether you are an elitist or an egalitarian, the value of that probability is supposed to do explanatory work in some way. The defense we are considering amounts to reinterpreting the SE so that the actual probability value p is not what plays any role in explaining the decay. Instead under this new interpretation, it is the fact that that the probability value of decay is non-zero that explains (or plays a role in explaining) the decay. In effect, the SE is reconstrued as claiming that the fact the decay was possible (because it has probability non-zero-- again we are assuming a non-standard treatment of probability) explains why the decay occurred. Quite apart from the point that this amounts to abandonment of core commitments of the SE framework, the idea that one can explain why an outcome is actual simply by citing the fact that it is possible is, to put it charitably, problematic. As far as I know, it is endorsed by no one, at least in the form described[15].

---

[15] On the w-condition or dependency theory of explanation, one can explain why an outcome is possible by identifying conditions that make a difference for whether it is possible. However, here the explanandum is the fact that the outcome is *possible*, rather than the *occurrence* of the outcome. Explaining why an outcome is possible is very different from explaining the occurrence of an outcome simply by appealing to the information that its occurrence is possible-- the latter of course is logically weaker than the fact that the outcome occurred.

It is worth emphasizing that these observations holds even if a high probability requirement is imposed on SE. Under such a requirement it will of course be true that given an explanans that endows explanandum E with a high probability p1, if that probability for E had instead taken some different and low value p2, E would not be explained. It is also true that if explanans X explains E in virtue of endowing it with a high probability, X will not explain not E since not E has a low probability. But in neither of these cases, does the explanation work by citing factors on which the occurrence of E depends.

Suppose that you are inclined to think that SEs do work by providing information about what their explananda depend on, but in some other way than what we have considered so far. What might this involve? Some may be tempted to invoke the picture briefly described above, according to which probabilities (or some of them -- perhaps those that qualify as "chances") are quasi-causal entities or powers or forces that contribute to the occurrence of the outcomes to which they are attached, by "probabilifying" them to different degrees. Perhaps one might then argue that if the probability p cited in schema S had been different, E would have been probabilified to a different degree-- a different level of probabilistic "umph" would have been delivered to E, so that in this respect the w-condition requirement is satisfied. I suggested above that this a highly problematic understanding of probability, but even putting that aside, we still face the problem that on this view what varies with different values of p are the degrees of probabilification, (these are what satisfy the -condition if anything does) rather than the target explanandum E.

Despite these difficulties, I expect that some readers will be inclined to search 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. As noted above, one might think instead of SEs as explanatory in virtue of conveying information about a pattern in the behavior of radium atoms and perhaps as showing us the decay of a particular atom is an “instance” of that pattern[16]. Alternatively, if you think that information about the degree to which an out come is expected explains that outcome you may conclude that elitist SEs are explanatory in virtue of providing that sort of information. The question then becomes whether one of these alternative views of explanation is acceptable or whether instead (as I will argue in Section 6) there are reasons to privilege difference-- making 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[17]. Suppose, following Scriven 1959, that Jones has paresis e and that this condition is caused by untreated syphilis s.

---

[16] But recall footnote 1.

[17] 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.

Assume, as is standard, that the probability of e, given that one has 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

5.1. (EX3) 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. Moreover, if (5.1) 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 (5.1) is a genuine explanation on a dependency or w-condition account, it wrong to treat it as an SE. Thus accepting (5.1) 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 (5. 1) and SEs. (5.1) 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[18]. Thus the w-condition requirement is satisfied by (5.2) .

---

[18] In view of the chancey nature of the connection between s and e, it is controversial whether the following counterfactual is true:

> 5.3. If Jones had untreated syphilis, the Jones would have developed paresis.

The Lewis-Stalnaker account of counterfactuals treats (5.3) as true but my own view is that when interpreted as an interventionist counterfactual (5.3) is false. (This is because I interpret 5.3 as the claim that if Jones did not have untreated syphilis, then under an intervention that causes him to have untreated syphilis, he would develop paresis, a claim which I take to be false.) Nonetheless the the truth of (5.2) and other associated counterfactuals is enough to ensure that (5.1) satisfies the w-criterion.

Note also that the assumption that there is a true counterfactual like (5.2) specifying conditions under which Jones' paresis would not have occurred is crucial to my account of why 5.1 is explanatory or at least why we are entitled to regard is as 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 c1 is present. Now suppose that s is present along with c1 and that the presence of s increases the probability of paresis to, say, 0.35. My view is that in such circumstances we have no basis for concluding that the counterfactual 5.1 is true, since, given what we know, we cannot exclude the possibility that paresis might have occurred even in the

In addition, there are a number of other true counterfactuals, also expressing relations of dependence, and answering w-questions in the vicinity of (5.1). For example, there is a related type-level causal generalization, known to and likely relied on by almost everyone employing (5.1) and arguably contributing to its explanatory status: untreated syphilis causes paresis. This is naturally understood as implying, for example, that in two otherwise similar groups in which no other causes of paresis are present, one of which has untreated syphilis and the other does not, the expected incidence of paresis will be higher in the former group. Variations in whether groups of subjects have or do not have untreated syphilis thus figures in answers to w-questions about the expected incidence of paresis in those groups.

We noted earlier that causal explanations work at least in part by citing difference-makers for their effects. (5.1) illustrates this idea. 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 with explanations of form S in which the cited explanatory factor --the probability p -- is not in fact a difference-maker

In addition , there are several other features of a singular causal explanation like (5.1) which 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. Indeed, both of these 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 is a reflection of the fact that untreated syphilis has an asymmetric explanatory relation with respect 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 e whether or not the probability of e, if s were to be present, is high, but we also think that s does not explain the non-occurrence of e even if that probability given s is high.

---

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 c1. Of course there might be additional evidence that allows us to definitely attribute Jones' paresis to s-- e.g. perhaps when s causes paresis it does so in accord with some characteristic modus operandi which is distinguished from the modus operandi by which c1 causes p and there is evidence that the former is present but not the latter, as in the shooting/drowning case described in footnote 7. In this case we have evidence that (5.1) is true. Note that if one observed Jones' paresis and needed to bet on whether it was caused by his untreated syphilis or c1, it would arguably be reasonable to choose s over c1. But this consideration does not show that s is what explains e and if in fact s explains e, this is not in virtue of the fact that Pr (e/s)> Pr (e/c1). Again to suppose otherwise is to conflate the explanatory relation with reasons for belief. For further discussion of and motivation for these claims, see Woodward 1990, 2021 .

Indeed, it seems that as long as it is true that s causes e and s is the only cause of e, and Jones has s and e , we can appeal to these facts to explain why Jones has e even if we don't know anything about the probability p of e in the presence of s and even if that probability is not cited in the explanation we provide. In other words, the value of the probability of e in the presence of s to seems not to figure in an explanation like (5.1) at all. By contrast the various philosophical accounts of SEs take a specification of that probability to be a crucial part of the explanation. 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[19].

Although the asymmetry present in 4.1 is not captured by either the egalitarian or elitist version of SE, the asymmetry is captured by the w-condition requirement. It is true that Jones' paretic condition would have been different if he had not had syphilis, but it is not true that Smith's non-paretic condition would have been different if he had not had syphilis. So again we see how the information about difference making, dependency relations, and w-questions seems different from (not captured by) the kind of information captured by the SE schema S.

There is yet another, related way in which a singular causal explanation like 5.1 differs from SEs as described above. (5.1) has a straightforward interventionist interpretation: intervening on the putative cause factor—e.g. preventing exposure to syphilis in the first place or by treating the syphilis effectively—is straightforwardly possible and the results of such an intervention would be to change -- in a reliable and systematic way-- whether Jones develops paresis. If we assume that there is a close connection between whether C causes E and whether it is possible to intervene on C and whether there is an associated change in E under such an intervention, then it is straightforward to conclude that syphilis causes paresis and in fact caused Jones' paresis. The same is true of the examples involving barrier penetration and mean free path discussed earlier.

By contrast the notion of intervening directly on a probability or on a probability distribution does not have a straightforward interpretation—it is not clear what this would amount to. Of course one can change a probability distribution by intervening to change *other* variables—for example, one can change the probability p of Jones developing paresis given that he has untreated syphilis by intervening to treat his syphilis but this not an intervention directly on p. Similarly, intervening to change the shape of a potential barrier V or the kinetic energy E of an approaching particle will change the probability of barrier penetration but it is unclear how one would go about intervening directly on this probability itself, independently of the variables V and E. Moreover, even if such interventions directly on probabilities are possible, most and arguably all such interventions would not change individual outcomes in any systematic away.

[19] My argument here is that both egalitarian and elitist SE are at variance with ordinary judgment about these cases. But-- to recall my earlier remarks-- I'm not claiming that egalitarian and elitist SE should be rejected merely because they are inconsistent with ordinary judgment. The dependency/w-condition account provides a principled reason for why we judge as we do in these cases, thus backing up these judgments with a more general account of what we are trying to accomplish when we provide explanations.

As argued above, changing the probability of barrier penetration from 0.25 to 0.50 is not a way of changing whether barrier penetration occurs. To the extent that we are willing to assume an interventionist framework, this is another consideration that puts pressure on the suggestion that that probabilities per se cause or provide causal explanations of individual outcomes. So this is another respect in SEs differ from explanations like (5.1) which are clearly causal[20].

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 (5.1) 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 (5.1) 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 (5.1) the explanandum-- occurrence of paresis - is not deducible from the factors cited in the explanans but (5.1) 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. (That is, expressed in terms of variables and values of variables, Salmon's claim is that C=1 can explain both E=1 and E=0). By contrast, in my view any plausible version of a dependency theory of explanation must accept Principle 1. It is hard to see what it 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 rejection of Principle 1 requires rejection of a dependency conception of explanation. Note also that 5.1/EX3 (as well as EX1 and EX2)

---

[20] At the risk of venturing too far into metaphysics land, let me note another relevant consideration. Probability is a modal notion. Probabilities are defined as measures over spaces of possible outcomes. Perhaps the impossibility of an outcome occurring can explain why the outcome does not occur, but as observed earlier it does not very natural to suppose that the possibility of an outcome (or some weight or measure over possibilities) is the sort of thing that can explain the occurrence of the outcome. This just seems like the wrong sort of thing to be an explainer. For a similar claim, see Hicks and Wilson, 2021.

respects Principle 1. Untreated syphilis explains 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.

6. Are Non-Dependency Accounts of Explanation Defensible?

So far one of my main arguments has been that if there are SEs, they must be understood in terms of a conception of explanation which does not take this to have to do with the tracing of dependency relations, but instead takes explanation to work by providing other kinds of information. This should not be surprising since the philosophers who first introduced issues about the structures of SE relied on just such non-dependency conceptions of explanation. For example, Hempel's IS model of SE relies on the more general idea that explanation is a matter of providing nomically based grounds for expecting that an explanandum will hold. Salmon's SR model relies on the idea that explanation is a matter of assembling statistically relevant information bearing on an outcome-- also a non-dependency notion, as noted below.

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 pattern subsumption are legion. On the assumption that untreated syphilis s is the only cause of paresis e, Jones' having e provides very strong grounds for expecting that he has s-- $Pr(s/e) = 1$-- but e does not explain s. (Note that arguably e is as strong grounds as it is possible to get for expecting s.) 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. Turning next to statistical relevance information of the sort appealed to in Salmon's SR model (understood as information about conditional and unconditional statistical independence relations) this often greatly underdetermines which causal and explanatory relations are present, even given information about temporal order and standard principles connecting statistical information and causal claims like the Causal Markov condition. Many different sets of causal relationships are often consistent with the same statistical relations[21] . In other words, statistical dependence does not amount to causal or explanatory dependence. X can be temporally prior to Y and statistically relevant to Y, both unconditionally and conditional on various other variables even though X does not cause Y (interventions on X do not change Y) and X can cause Y even though X and Y ae statistically independent.

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 to this line of argument is to invoke some variety of pluralism about explanation: it might be argued that even if there are there problems with non-dependency views of the sort just described (at least in connection with some examples), these 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. (As it might be put, "there are many explanatory virtues and many things that might be meant by a good explanation.") Hence, if

[21] The extent of this underdetermination is described by theorems in Pearl, 2000 and Spirtes, Glymour and Scheines, 1993

SEs don't fit well with a dependency framework, why not just understand them in terms of one of these alternative frameworks? Going further one might wonder whether 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 expanation 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).

I find this general line of argument unconvincing for several reasons. First, 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 immediately below). However, in order for this goal to provide useful guidance there must be non-trivial constraints on what counts as an explanation. That is, if there are a number of acceptable theories of explanation with different different and inconsistent implications for the evaluation of various examples, as the pluralist maintains , this threatens to undermine any principled basis for the assessment of the explanatory credentials of different hypotheses, which we've been assuming is one of the goals of an account of explanation. I agree of course the fact that 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. 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.

This issue about explanatory pluralism is particularly acute for those who subscribe to some form of IBE. If whether e is evidence for h depends on how well h explains e, then any indefiniteness in what counts as a good explanation will also infect assessments of evidential support. (Is there evidence for anthropogenic climate change? Well, evidential support is a matter of IBE and thus it all depends on what you understand by "explanation".) Whatever one thinks about *explanatory* pluralism, this sort of pluralism (or permissiveness) about *evidence* seems inconsistent with the way in which evidential considerations figure in science and indeed with what seems distinctive about science as opposed to other sorts of enterprises.

Another way of putting this is that the kind of pluralism just described threatens to trivialize the whole discussion around SE. The philosophers (Hempel, Salmon, Railton) who launched this discussion accepted certain ground rules -- in particular, they agreed on the need to motivate their models of SE by locating them within much more general frameworks for thinking about explanation. This implied in turn that problems for these more general frameworks had implications for their models of SE. Strong forms of pluralism seem to amount to dropping these ground rules with the result that the constraints on the whole discussion become unclear. Presumably it is not satisfactory to merely stipulate that such and such counts as an explanation and then "argue" that SEs are explanatory in virtue of satisfying this condition, but how exactly is this to be avoided under strong forms of pluralism?

7. IBE and the Role of Explanation in Theory Assessment

In this section I want to take a step back from the observations in the previous section and examine a more general issue which has implications for how we should think about SE, as well as being important in its own right. The general issue is this: in what sense is explanation a goal of inquiry and how does such a goal relate to what we are trying to accomplish when we develop philosophical accounts of explanation?

I begin with some assumptions, that I hope will seem relatively uncontroversial and that are shared by defenders of SE. Both in science and in everyday life we sometimes aim at providing explanations (in some sense or senses of that notion) and hypotheses, theories, models etc. (hereafter I will say hypotheses) that successfully do this have, to that extent, positive value. This is not to say that claims that fail to do this (because, for example, they are "merely descriptive" and not explanatory) are valueless, but rather that explanations are a good thing to have when we can get them. In particular, we can sometimes compare hypotheses with respect to whether they explain various candidate explananda. For example one might think that the Bohr model explains certain features of the emission spectrum of hydrogen but fails to explain the spectrum for helium. One important goal for a philosophical account of explanation is to provide standards or criteria for such comparisons. As argued previously one problem with pluralism about explanation is that unless it is constrained in some way it does not seem to serve this goal.

Now a more controversial claim: Recognition that finding explanations is a goal need not commit us to anything like "inference to the best explanation". Here, as before, I will understand IBE as the proposal that given a set of competing hypotheses, all consistent with presently available evidence e, we should infer that the hypothesis h which, if true, would best explain e *is* true (or at least is better confirmed than its competitors or at least that e is evidence for h.)

An alternative view, which I endorse instead and which consistent with the idea that explanation is an important goal in science is this: the fact that a hypothesis if true would explain some range of explananda can be a reason to investigate whether it is true, but this fact is not in itself reason to think that the hypothesis *is* true or better confirmed than competitors. Instead, to support the claim that such a hypothesis is true or well confirmed we must provide evidence that is independent of its explanatory potential (that is, whether if true it would explain well). The reason that explanatory potential can provide a reason to investigate whether a hypothesis is true is not that explanatory potential per se is evidence of truth but rather simply that we value explanation and we will have one if the hypothesis turns out to be true.

On this view, rather than thinking of the explanatory potential of a hypothesis as important because it provides evidence for truth, the discovery of explanations is instead regarded as, so to speak, an end in itself. Of course we want (or so I will assume) our candidate explanations to be true or evidentially well-supported[22] but it gets things the wrong way around to suppose that we value explanatoriness or potential explanatoriness because this indicates truth or evidential support. Instead we value successful explanation and this requires true or evidentially well supported hypotheses.

One way of motivating this view is to observe that truth per se is not a very good candidate for a goal of scientific inquiry. An exact catalog of the dimensions of each of the grains of sand

[22] For purposes of this paper I am going to put aside any discussion of recent claims that theories etc. that are "radically false" can none the less be used to explain. This claim is not directly relevant to discussions of SE. For what it is worth I think that these contentions about the role of falsehoods in explanation suffer from considerable unclarity about just what is claimed to be false and often rest on a kind of literalism in the interpretation of theories that has little to recommend it. See Woodward, 2023 for additional discussion.

on a large stretch of beach may contain lots of truths but for the most part these are not what science aims at discovering. Scientific inquiry instead has as a goal the discovery of more specific kinds of truths-- for example, truths that can be used to explain, among other possibilities[23]. In other words, a concern with explanation gives inquiry a focus -- a focus on the discovery of certain kinds of truth-- that it would not have if the goal was just the enumeration of truths of any sort whatsoever.

The structure of evidential support within real life science is often extremely complicated, with the connection between evidence and the hypothesis it supports taking many different forms[24]. But, as suggested above, one important component of support involves the elimination of competing alternatives to a hypothesis. We can think of this as connected to explanation in the following way: Suppose the context is such that it is plausible to make the working assumption that there is some explanation or other for a set of explananda. On my view, this assumption may turn out to be false-- for example, if the explananda are individual occurrences that are the result of some irreducibly stochastic process, as in QM --but it is often a plausible starting or default assumption. The goal of finding an explanation then helps to demarcate the class of alternative competing hypotheses we consider -- if this is our goal, we consider only those hypotheses that are potential explanations of the explananda. We then look for evidence that rules out or renders implausible as many as possible of these alternative explanations, so that in the ideal case only one remains. To oversimplify greatly, in a context in which the potentially explanatory hypotheses are deterministic we might provide evidential support for potential explanation h1 by showing that plausible alternative explanatory hypotheses h2..hn make predictions that are false and that the h1 makes predictions that are correct. Alternatively it may be that only h1 explains (and correctly implies) some target explanandum e; the alternatives h1..hn don't make false predictions in connection with e, but rather say nothing at all. In a context in which the explanatory hypotheses under consideration are probabilistic, we might proceed by showing that the alternatives h2... hn make predictions about probabilities that are rejected by appropriate statistical tests while this is not true for h1, leading us to adopt h1. We might describe these as inferences to the *only* explanation (IOE) (that is the only one that remains after empirical testing) rather than IBE. Inference to the only explanation is a two step process; the two steps being 1) the identification or construction of alternative potentially explanatory hypotheses (that is generation of alternative hypotheses that would if true explain some target explananda-- this is the step at which explanatory considerations are important), and 2) the empirical testing of these alternatives. By contrast IBE collapses these into a single step;

[23] Again, there are other sorts of goals that are important in scientific inquiry-- for example, the discovery of theories that can be used to predict successfully, even if these are not regarded as explanatory. Note also that this example illustrates another point relevant to our discussion: There are lots of features that a theory can possess that are valuable without those features providing evidence for the truth of the theory. That a theory can be used to successfully predict is a valuable feature and this of course requires that at least some parts of the theory - e.g. the predictions it makes-- be true. But it would be absurd to argue, in connection with theory T, from the premise that if T were true, it would predict well, to the conclusion that T is true or evidentially well-supported. That a theory would, if true, predict well (its potential predictiveness) is a reason for investigating whether it is true or whether it predicts well but its potential predictiveness is not evidence that it is true or predictive. Similarly for explanation.
[24] For a tour-de-force see Smith and Seth, 2020.

the step-1 fact that a hypothesis is potentially explanatory and best among rivals by itself provides evidential support of the sort sought in step 2[25].

I claim (Section 8) that it follows from these ideas about the role of explanation in theory evaluation 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.

8. The Assumption that there are SEs is Not Necessary for the Evaluation of Statistical Theories or the Correctness of Probability Ascriptions.

In arguing for the claims at the end of the previous section I begin 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. This allows us to consider, as I do immediately below, scenarios in which probability ascriptions are accurate and the only issue is the conditions under which they explain individual outcomes. I regard this as an extremely plausible assumption but 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 examples considered earlier.

---

[25] There is much more that might be said in support of this two step picture IOE than I am able to discuss here. But note an additional attraction of this picture in comparison with IBE: IBE makes much stronger demands for its legitimate use than IOE. First IBE requires account of explanation that allows one to determine not just (i) whether some set of alternative hypotheses are potentially explanatory of some target explananda but also (ii) to compare such hypotheses in a way that allows one to determine which explains "best". Use of IOE requires (i) but not (ii), since the task of discriminating among alternative explanatory hypotheses is left to empirical testing. Indeed use of IOE need not assume that we always have available well-defined notion of "best" explanation, as opposed to some conception that allows to assess whether or not some candidate hypothesis would if true explain (or explain well enough) some explanandum and thus that h explains explananda that h' does not. Second, an obvious question raised by use of IBE is why we should suppose that explanatory goodness is connected with truth (or with whatever we think is established when a hypothesis has strong evidential support.) Why should the best explanation be the true one? By contrast, IOE assumes no such connection between explanatory bestness (even assuming this is well-defined) and truth. Again, for IOE truth is established via a separate step.

Let us compare egalitarian SE and NO SE for a theory T that makes accurate predictions about the probabilities of individual outcomes. Suppose first that T explains those probabilities (in virtue of satisfying the w-condition and whatever addition criteria are thought appropriate for successful explanation). Strictly speaking SE and NO SE say nothing about how to assess theories that purport to explain probabilities 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 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 that imply probabilities or among the probability ascriptions they imply[26]. 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 or among the probability ascriptions themselves-- for example, in terms of their values, with higher probability assignments (or theories that make higher probability assignments) somehow being preferable. But this is some 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 or individual probability ascriptions beyond what is provided by NO SE-- both reach exactly the same conclusions about which probabilistic theories or probability assignments are good or acceptable. In other words the claim that individual outcomes in addition to probabilities of outcomes are explained in egalitarian SE does no real work in evaluation, either of theories that imply probabilistic claims or those claims themselves. One would get the same results if one were to drop egalitarian SE in favor of NO SE. So why assume SE, at least in its egalitarian version, especially if, as I have argued, its adoption generates other difficulties?

In response it might be argued that this argument fails to acknowledge the role played by IBE in the case for egalitarian SE: the judgment that certain probability assignments are accurate (it will be claimed) is itself based on explanatory considerations (we infer to the 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 SE. However, it is hard to see how this argument can be made to work. Since, according to egalitarian SE, 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 explain that outcome well and hence might be claimed to be well-supported on the basis of an IBE but of course the same consideration might be invoked

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

in support of any alternative probability assignment. It follows that 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 correct (where this is established in some alternative way) -- the claim that individual outcomes are explained by probability ascriptions do no work and can be dropped without loss[27].

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 O, corresponding to occurrence and non-occurrence and that the theory assigns a probability to this outcome. Assume (on the basis of the arguments presented 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) is sufficiently high, the occurrence of one these outcomes but not the other will be explained. On the other hand, if it should happen that Pr(E=1) = 0.5, neither outcome will be explained, so a T predicting this probability will be unexplanatory in connection with this individual event. In this respect the theory predicting intermediate probabilities, even if these are empirically accurate will be explanatorily inferior 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. A theory that accurately predicts that a six-sided die will be fair will be completely unexplanatory at the level of individual events. Similar conclusions follow if we just consider these probability ascriptions by themselves-- e. g. , the assignment Pr (side 1) =1/6 , rather than theories that imply them.

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 has probability 0.5 and also implies that in system 2 some different outcome 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 these probabilities and so on.

These explanatory assessments seem odd, but there are more fundamental problems. Suppose that we extract the following general advice from elitist SE. We should prefer probability ascriptions or theories that imply probability ascriptions, some of which are high (of course that implies that other probability values are low) to ascriptions that assign intermediate

---

[27] 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. (These are what we are trying to infer to via the IBE.) 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 (or way of determining) 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 SE/IBE appears redundant. In the second case, if we rely on SE/IBE we will make a mistaken ascription.

probabilities or all low probabilities. In other words, theories that make extremal probability assignments should be preferred.

So far we have been assuming that there is some way of evaluating the empirical accuracy of probability assignments in addition to or independent of IBE/elitist SE-based assignments. We thus face the same problem as before: If there is such an independent standard and it mandates assignments that are different from those mandated by IBE/elitist SE, then it is unclear why should not prefer the former assignments. Moreover if the recommendations of the independent standard always agree with those of IBE/elitist SE, we no longer have an argument that we need the latter to make sense of the evidential support for probability assignments-- the independent standard suffices.

Can the advocate of IBE-based arguments for elitist SE respond by contending that there is no independent standard-- all assessments of evidential support must be based on IBE? As noted previously this seems to imply that we could never be in a position to determine on empirical grounds that the IBE/elitist SE based assignments are mistaken, since the correct assignments are just what the IBE/elitist SE argument says. In addition we face the obvious problem of how to understand the evidential support for assignments involving low or intermediate probabilities. Presumably we can get empirical evidence that a coin has probability 0.5 of landing heads or a tosses of a die are fair, even though these probability assignments do not, assuming elitist SE, explain individual outcomes and thus apparently cannot be justified via an IBE based just on such outcomes. How is that supposed to work?

The obvious strategy for the advocate of elitist SE at this point is to consider outcomes (other than individual events) that *are* highly probable and hence explained by the above probability assignments, with this serving to provide evidential support for the assignments. In the case of the fair coin, these might be taken to be, e.g., a relative frequency of heads in a substantial repeated number of repeated tosses that falls within an interval of such frequencies that is "close" to 0.5 where the interval is chosen in such a way that some outcome within that interval is highly probable, given the probability assignment. 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". The argument would then be that the probability assignment Pr (Heads)= 0.5 explains, in accordance with elitist SE requirements, why the relative frequency falls within the chosen interval, since that outcome is highly probable, given the probability assignment, and this implies that the fact that the relative frequency falls within the interval is evidence for the probability assignment.

With this strategy we have moved a considerable distance away from the intuitions/ assumptions that originally motivated elitist SE (or for that matter, any 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 but rather facts about relative frequencies falling within intervals. Similarly it is such facts (and not individual outcomes or particular sequences of these) that serve as evidence for probability assignments[28]. Second, because, as illustrated, it is always possible to find explananda that are highly probable even given low probability assignments, it is no longer clear how elitist SE can be used as a basis

---

[28] Note also that adopting this strategy for elitist SE leads to a kind of bifurcation in what SEs explain. On the one hand, a high probability assignment to a particular outcome can be used o explain it. On the other hand, the explananda associated with low probability ascriptions are something quite different--- the falls-within-a probable-interval-facts described above.

for choice among such assignments or the theories that imply them. If an assignment of probability 1/6 to each of the faces of a die can explain some explananda (again 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 or the probabilities they assign? Third, with the modifications just described, the resulting picture of elitist SE now looks 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. The claims that are distinctive to SE and IBE-based arguments for SE seem to have been lost. Again, 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 SE or the role of explanatory considerations as unnecessary.

The underlying difficulty here derives, in my assessment, from the assumption that probability assignments are based on some form of IBE. This makes it seem as if we have to find some explanatory role for those assignments-- hence to some version of SE. But attempts to develop this line of thought lead to paradox: On egalitarian SE there can be no explanation-based inference supporting one probability assignment over another since all such assignments furnish equally good explanations[29]. Thus advocates of egalitarian SE must agree that there is some other basis for assigning probabilities besides IBE-based considerations. On elitist SE, probability assignments which are high provide better explanations of individual outcome than low or intermediate assignments and thus high assignments may seem to be preferentially warranted via an IBE. On the other hand, it seems undeniable that intermediate or low assignments are sometimes empirically warranted. This again seems to suggest either (i) that there is some other basis for probability assignments besides IBE and elitist SE or else (ii) that the explananda and evidence in such assignments are very different from individual events, instead being something like the fact F that observed frequencies fall within an interval. Acknowledging (i) undercuts IBE-based arguments for SE, since we don't need to appeal to explanatory considerations to warrant probability assignments. Adopting (ii) implies that low probability assignments can explain some outcomes involving frequencies (facts like F) and seems to undercut one of the primary arguments for elitist SE, which is that we need to to discriminate among different statistical theories and probability assignments in terms of how well they explain and elitist SE is required for this purpose.

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 19th century. As I understand his argument it assumes that there is a connection between explanation and evidential support that looks something like the assumptions (2.3.1ab) identified in Section 2:

---

[29] This is in effect Strevens' (2000) argument against egalitarian SE. I agree with this argument but not with Strevens' argument for elitist SE-- see Section 9.

2.3.1a) (EX/EV) e is evidence for h if h is the best explanation for e. (best explanation--> evidence)

(2.3.1b) (EV/EX) e is evidence for h only if h is the best explanation for e. (evidence--> best explanation

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 (2.3.1b), 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 (2.3.1a), (9.1*) would also be evidence for SM, in contrast to our judgment that occurrences like (1) but not (9.1*) provide evidence for SM and indeed 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 observed to occur and failure to diffuse is never observed. Thus, on this basis we have very good reason to accept

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

Note that we need not think of acceptance of (E) as the result of an IBE-- or at least we need not view matters in this way unless we are already committed to the view that inference to probabilities always involves IBE, a view that, as noted earlier, faces a number of problems. In fact sophisticated statistical tests do not seem required to establish E; prima-facie we can instead conclude E just on the basis of what we observe without any detour through claims about statistical explanation. (In all known cases, the gas diffuses, we have no reason to suppose that our observations are unrepresentative, etc.)

To the extent that SM implies (and explains) E that is a point in favor of SM. Moreover, any alternative theory T that implies that (E*) uniform diffusion has a low probability of occurring should be rejected on empirical grounds-- that is, it should be rejected because E* is false. (Again we don't need to assume some version of SE to reach this conclusion.) Similarly a theory T* that fails to imply (or explain) E , perhaps, e.g., because it merely says that uniform diffusion is possible or fails to say anything at all about uniform diffusion is, in this respect, inferior to SM. If, as Strevens seems to assume, we are considering a rather generic version of SM according to which T and T* are the only alternative to it, it is clear why we should prefer SM to these. E is evidence for SM because it preferentially supports SM over the alternatives to it. Thus we can explain why it was reasonable to accept SM and to think that there was strong evidence in support of it without assuming that SM explains individual events like particular episodes of gas diffusion[30]. Still less do we need to assume an elitist version of SE to explain the acceptability of SM.

Finally the difficulties described in Section 8 remain: Although Strevens' argument for (moderate) elitist SE focuses on high probability events such as diffusions and has little to say about the treatment of low probability events, it is nonetheless true that low probability assignments are sometimes warranted. As argued above, to make sense of these, 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, if this strategy is adopted, it is hard to see how the high probability requirement can be used to discriminate between theories like SM and alternatives on the basis of the probability assignments they make. Instead on the strategy under consideration, both high and low probability assignments can explain (the latter in virtue of explaining facts like F) and can be used in an IBE (assuming one regards this as an acceptable inference form). This just reinforces the conclusion that elitist SE does no work in licensing the acceptability of statistical theories like SM.

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 according to 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 ae 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 (should be assigned) via an inference to the best explanation. 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 accord with IBE type considerations is apparent in the following passage:

---

[30] Of course this whole story about the relation between gas diffusion and SM is widly oversimplified, as Strevens acknowledges. In actual fact, the acceptance of SM was based on many difference considerations, both theoretical and empirical. But this does not affect the issue of what we should conclude from the example.

> 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][31]

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, R, with heterozygous (Bb) parents. Egalitarian SE, as we have been understanding it, implies that R's blue eyes (or possession of a BB genotype) can be explained by observing that the probability of this outcome is 0.25[32]. Note however that if we know R 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 R must have received a B from each parent, with R's blue eyes following deterministically from (ii). Thus we know the specific sequence of events (according to some, the steps of the mechanism) that deterministically produced R's blue eyes. Some may regard this as an explanation of why R has blue eyes -- hence that there is

---

[31] 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 the suggestion seems to amount to the likelihoodist assumption that the hypothesis hi for which Pr (e/hi) is highest (among the various alternatives hi) is best supported by e and that Pr(e/hi) is a measure of how well hi explains e. Of course this is compatible with Pr(e/hi) being low.

[32] Cf. Clatterbuck, 2020:

> So, to explain why Ray, the child of two brown-eyed parents, has blue eyes, the Mendelian cites the Law of Segregation, that the brown eye allele (B) is dominant to the blue eye allele (b) and that his parents are heterozygous (Bb Bb). The explanans entails that offspring of this pairing will have blue eyes (bb) with probability .25 and that the expected frequency of blue eyes among offspring will be 1/4.

Again the issue is why we should suppose that the entailment described in the second sentence suffices to explain Ray's blue eyes, as the first sentence asserts.

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 arises for both the elitist and egalitarian versions of SE. So far we considered mainly discrete probability distributions, which can be represented by a 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 well- behaved pdfs, the probability of *X* taking any particular 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 Section 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[33] to count as "outcomes" in the relevant sense, taking these claims to be the explananda of SEs provided by pdfs. But without some further restrictions on what intervals are acceptable, triviality threatens for this version of elitism. This is because for any value x taken by the random variable X , there will always be some interval such that the probability that x falls within that interval is high enough to exceed whatever threshold is imposed by elitist SE for successful statistical explanation. Indeed, regardless of the pdf involved integration over the full range of possible values of X will of course yield the result that X falls within that range with probability 1, so that we can always achieve the highest possible threshold.

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 achieve 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 probable that values of the random variable will fall within this interval) are 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

[33] Presumably the "interval" needs to satisfy some "connectedness" requirement (that is, the target explanandum should not be something like X falls with in (1,2) U ((4,6) but I will not pursue how this might be formulated or what the rationale for the requirement might be

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. For non-peaked densities we also might consider invoking the "falls within a probable interval strategy" -- a non-peaked density is evidentially supported if it assigns a high probability to observed outcomes within some appropriate interval but this faces the difficulties described previously.

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 is what 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. But I think that the deeper issue is that it is unclear how to settle this question in a principled way. After all, 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 a number of 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. Indeed I claim that even if were true universally that all ice cubes melt when placed in warm water, this generalization would not explain why some particular episode of melting. Moreover I don't think that there is anything particularly counterintuitive or incredulity inducing about either of these claims[34], although I focus in what follows on the high probability claim.

Suppose that I know nothing about SM and observe an ice cube melting in warm water. I ask why

(12. 1) this ice cube is melting.

[34] Alternatively, even if what I say about counterintuitiveness is not fully accepted, I hope that it will suggest how problematic it is to frame one's account of explanation just around what seems to be intuitive judgments about particular cases.

Is it an explanation to be told that

(12.2) (12.1) 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? (Why do all or almost all of them behave this way?) Note also that unless I have had very unusual life experiences I already have observed lots of ice cubes in warm water, all of which have melted, so in asking for the above explanation, I'm not asking whether such melting behavior is common or regular or probable. Instead I'm asking for some other sort of information -- very likely about what the melting behavior depends on.

Note also that the information that ice cubes always or virtually always melt in warm water was widely known before the discovery of SM. The basic ideas of probability theory were also developed before SM, so that the high probability of ice cube melting was also well known before the advent of SM. Thus if (12. 2) explains (12.1) it follows that many people were in possession of an explanation for why individual cubes melted long before SM was known. Indeed even some one who knew nothing about the existence of molecules or the relation between molecular kinetic energy and temperature could possess such an explanation. Similar conclusions hold for other explananda discussed by Strevens and others such as gas diffusion. Given elitist SE, one has an explanation of individual episodes of gas diffusion based on the known truth that diffusion is highly probable, even if one does not know that gases are composed of molecules, that these are what diffuse, that possible microstates of the gas have equal probability and so on. This does not seem to me to capture explanatory judgments that are widely endorsed in science. I don't claim that by themselves these considerations show that the claim that (12.2) explains (12.1) is wrong but they do seem to me to undercut the contention that there is something particularly counterintuitive about the denial of this claim.

At this point I anticipate 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). In the same way, in a deterministic context, we need to distinguish an explanation for why this (iii) particular a which is F is also G (why this particular emerald is green) from an explanation of why (iv) all Fs are Gs. (why all emeralds are green.) (iv) can be used to explain (iii) even though some deeper theory is needed to explain (iv).

This two -level picture of explanation is deeply entrenched in philosophical discussion but I don't think it fits scientific practice. Although there are exceptions, particularly in the historical sciences, science is not for the most part concerned with explaining individual outcomes or events but rather with explaining generic repeatable robust patterns -- phenomena in the sense of Bogen and Woodward, 1988. That is, the explananda of interest are taken to be , e.g., why all samples of pure iron have a melting point of 1538 C rather than why this particular sample of iron has melting point 1538 C , and why the intensity of an electromagnetic field due to current passing through a long straight wire is perpendicular to wire and falls off as the first power of the wire rather than why some particular wire has this property. In cases like these, to the extent that there is any concern with the explanation of particular outcomes, these make use of the same information that is used to explain the generic outcome, although specialized to the particular outcome. For example, to the extent it is true that all emeralds are green, this has to do with the presence of trace amounts of chromium and vanadium in these gems. It is the presence of these elements which also explains why some particular emerald is green. There is no separate level of explanation according to which the greenness of some individual emerald is explained by reference to all emeralds being green, with the latter then being explained by the presence of the trace elements.

Similarly 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. However, the analogy with the emerald case is only partial: while SM does indeed explain why ice cube melting is highly probable, if my arguments above are correct, there is no need to suppose that SM explains individual episodes of ice cube melting-- at least via any form of SE. The disanalogy arises because of the probabilistic character of the melting process in comparison with what we are taking to be deterministic way in which the presence of the appropriate trace elements leads to the greenness of emeralds[35].

I said above the targets of scientific explanation are robust repeatable phenomena. Emery (2017) adopts a similar position but she holds that for SEs the robust phenomena that are explained are facts about relative frequencies being close to what they would be if those frequencies exactly mirrored the probability value doing the explaining (that is close to the frequency that has the maximum likelihood-- e.g., relative frequencies that are close to 50% heads in the case of tosses of a fair coin). But such phenomena are, in an obvious sense, non-robust or at least less than fully robust. However closeness is understood, frequencies that are not close in the sense described above can be expected to occur with predicable probabilities What is robust (or considerably more robust) are the probabilities themselves, assuming they are empirically accurate. Thus in the case of theories that predict probabilities of outcomes, these are the natural candidates for the stable phenomena that serve as explananda. Other candidates -- individual outcomes, frequencies of outcomes, frequencies that fall within some interval -- are, to varying degrees, not robust in the sense that these may fail to occur, even when the theory in question applies. In other words the stable patterns that stands in need of explanation from a scientific perspective are the probabilities, rather the other explananda described above.

13. The Nothing More to Be Said Argument for the Existence of SEs

---

[35] This is not to say that there is no explanation of why an individual ice cube has melted but rather that whatever form this takes it will not appeal to a probabilistic claim like 12.2. For example, one might hold that there is a possible-in-principle deterministic explanation that takes the form of deducing the trajectories of the individual molecules composing the cube as melting occurs from the initial conditions characterizing these and their laws of motion. There is also a close by singular causal explanation that appeals to the environmental factors that causally contribute to the melting-- such as the fact the cube was removed from the fridge and placed in a room temperature environment. Here the explanatory factor is this environmental cause rather than a claim about high probability.

There is yet another line of argument that may seem to support the existence of SEs. The basic idea is this: Once one has described all of the causal or other explanatory factors on which an outcome depends, so that there is nothing more to be said (NMTBS) about these, one has explained the outcome. One way of motivating this idea appeals to David Lewis' claim that to provide a causal explanation of an event is to provide information about its causal history. Lewis claims that what he calls "negative information" -- information that certain causes were absent-- still counts as information about causal history and thus can be part of a part of a causal explanation. In the extreme case in which there are no causes for an outcome -- Lewis' illustration is the limit on of stellar collapse provided by the Pauli exclusion principle-- we can nonetheless explain the outcome by citing the absence of causes. Applied to SE, the implication is that once we have described all of the factors that are relevant to the occurrence of the event -- either by affecting its probability or in some other way--so that there is nothing more to be said about why it occurred, then we have explained it. A more general version of this line of argument appeals to the idea that to explain an outcome is to provide information about that portion of the "ideal explanatory text" (as in Railton, 1978, Salmon, 1984) that concerns the outcome. If the ideal text lists no causes for the outcome or no factors that determine the outcome but does describe the probability of the outcome, given various antecedent factors , then that information explains the outcome, on the assumption that we have provided all of the information there is that bears on the occurrence of the outcome. Indeed, according to this line of thought, we have provided the best possible explanation for the outcome, since we have omitted no information relevant to its occurrence.

Of course this argument in defense of SEs is prima-facie problematic when applied to deterministic systems whose behavior is described in terms of probabilities, since in such cases, there *is* more to be said beyond the probability ascriptions-- the deterministic story. However, in quantum mechanical cases, the various no hidden variable theorems do seem to provide assurance that there is nothing more to be said about individual outcomes beyond the probability ascriptions themselves. Appeal to the NMTBS argument thus seems to support the conclusion that SEs may be found in quantum mechanics but perhaps not elsewhere.

The idea that one can explain or, more specifically, provide a causal explanation for an outcome by appealing to the fact it has no causes is , to understate matters, not how we usually think about explanation and certainly not how we think about causal explanation. Instead, our usual practice, both in science and in common sense, is to think that causal explanation requires citing true claims describing existing causes of the outcome, not just the information that these are absent. Even if one thinks that there are non-causal explanations, as long as these are a matter of citing factors on which an explanandum depends (the view I have been defending), one won't think that task is accomplished by saying that there are no factors on which the explanandum depends even if that information is correct.

Of course the idea that there are SEs does not require the assumption that one can provide a causal explanation of an outcome by providing only negative information. SEs do provide some "positive" information about outcomes-- their probabilities. However, if one has misgivings about the general NMTBS argument, one presumably will not find the claim that probability ascriptions explain individual outcomes because they say all that can be said very convincing, even when this claim is true.

Putting this aside, reflection on the NMTBS argument does raise the following question: where, so to speak, does one draw the line about the kind and amount of information that needs to be provided for explanation, if one assumes that a putative explanans consisting of only negative information is insufficient to explain? Why say, as the defender of SE does, that the provision of information about probabilities of outcomes is sufficient to explain but anything less is not? Consider a theory that correctly tells us only that under certain conditions (but not others), various explananda are possible, but without specifying their probabilities of occurrence. Suppose that there is nothing more to be said about why these outcomes occur other than these facts about possibility. Suppose one of these outcomes occurs. Can we explain its occurrence by citing the conditions that led to its being possible? Can we explain why the outcome occurs just by citing the fact that it was possible (leaving out information about the conditions relevant to its possibility) and that there is nothing more to be said about why it is possible? (Talk about the power of possibility!) Alternatively, how about a theory according to which the probability that X takes some value x falls within a certain interval-- e.g Pr (X=x) falls within (0.4-- 0.7), again with nothing more to be said about why x takes some particular value within this interval? Is that sufficient for an SE explanation of X=x? How about cases in which the values of X comprise a non-measurable set, so that no probability assignments are possible but in which X takes some particular value within this set? If it is appropriate to think of probabilities (or at least specific probability assignments) as quasi-causal-- they make things happen-- and other things that might go into an explanans (e.g., information about mere possibilities) as causally inert, we would have answer to the question of why information about probabilities has a special status in explanation. However, as we have argued, this interpretation of probabilities is problematic. Indeed, as suggested earlier, probabilities seem more akin to possibilities (they are measures over possibilities) than to causal agents and possibilities seem causally inert.

The NMBS argument raises an interesting and relatively unexplored question in the theory of explanation. Should that theory allow for the possibility that there are certain explananda that cannot be explained-- not in the sense that we might not be able to discover the needed information relevant to explaining their occurrence but in the sense that even if we had all relevant information (the full ideal text) that would not be enough to explain it? The NMBTS argument in effect claims that there are a priori reasons why this cannot happen while a dependency-based account of explanation takes it to be an empirical, a posteriori issue whether explanations for certain kinds of outcomes are possible even in principle-- it may be that nature does not cooperate in our efforts to explain certain explananda in the sense that the required dependency relations do not exist. Arguably this is the case for certain particular quantum mechanical outcomes. On this view, while science looks for explanations, it does not follow that it must always succeed in finding them.

Still despite its arguably counterintuitive and too-good- to-be-true character[36], it is nonetheless worth asking whether there is some functional rationale for accepting the NMTBS argument. Would it be good or bad, methodologically speaking, if we were adopt this way of thinking about explanation? One obvious observation is that we do not have access to the ideal

[36] Too good to be true in the sense that an uncharitable description of the NMTBS argument is that it is a proposal to declare success in finding something on the basis of a demonstration that there is nothing of that sort to be found. Compare "the evidence for p is that there is no possible evidence for p-- indeed this is the best possible evidence since there is none better." In general a demonstration that we cannot fully succeed in reaching a goal is not a good reason for reconceptualizing "success" in a way that implies that we have reached the goal after all. Perhaps the goal of finding an explanation behaves differently but if so, some additional argument for this conclusion is needed.

explanatory text for most explananda and hence are in no position to say whether there is nothing more to be said beyond about them beyond the explanations we currently possess. Even in the most favorable cases, it seems that the most that we have reason to believe is that the ideal text will not contain explanations meeting certain specified conditions rather than that there is nothing more to be said of any kind about various explananda -- e.g., that the ideal text will not contain some local hidden variable theory that explains various quantum mechanical phenomena. But even in this case it is hard to see how we can be confident that there are no possible theoretical advances of a different character that might cast explanatory light on the explananda in question. The worry is thus that in realistic cases we are either in no position to appeal to the NMTBS argument (since we can't tell whether there is nothing more to be said) or else, alternatively, we may apply it prematurely in a way that shuts down inquiry, concluding on the basis of NMTBS considerations that we have fully explained certain phenomena and a search for some further or deeper explanation is unnecessary.

References

Bogen, J. and Woodward, J. (1988) "Saving the Phenomena" *Philosphical Review* 97:303-352

Clatterbuck, H. (2020) "A Defense of Low Probability Scientific Explanations" *Philosophy of Science* 87: 91-112.

Earman, J. (1992) *Bayes or Bust.* Cambridge: MIT Press

Elliott, K. (2021) "Where are the Chances?" *Synthese* 199: 6761-6783

Emery, N. (2015) “Chance, Possibility, and Explanation.” *British Journal for the Philosophy of Science* 66 (1): 95–120.

Emery, N. (2017) "A Naturalist's Guide to Objective Chance" *Philosophy of Science*

Engel, E. (1992) *A Road to Randomness in Physical Systems.* Berlin: Springer- Verlag

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. New York: Free Press.

Hicks, M. and Wilson, A. (2021) "How Chance Explains" *Nous*

Hitchcock, C. (2001) "The Intransitivity of Causation Revealed in Equations and Graphs" *Journal of Philosophy* 98: 273-299.

Keller, J. (1986) "The Probability of Heads" *American Mathematical Monthly* 93: 191-6.

Kitcher, P. (1989) "Explanatory Unification and the Causal Structure of the World" in *Scientific Explanation*. Minnesota Studies in the Philosophy of Science, volume 13 (1989), eds. Philip Kitcher and Wesley Salmon, 410-505.

Lewis, D. (1986) "Causal Explanation" in Lewis, D. , *Philosophical Papers, Volume 2*. Oxford: Oxford University Press.

Pearl, J. (2000) *Causality*. Cambridge: 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 *Statistical Explanation and Statistical Relevance,* ed. W. Salmon, 29–87. Pittsburgh: University of Pittsburgh Press.

Salmon, W. (1984) *Scientific Explanation and the Causal Structure of the World*. Princeton. Princeton University Press.

Salmon, W. (1989) *Four Decades of Scientific Explanation*. Minneapolis: University of Minnesota Press.

Scriven, M. (1959) "Explanation and Prediction in Evolutionary Theory" *Science* 30: 477-82.

Smith, G and Seth, R. (2020) *Brownian Motion and Molecular Reality.* Oxford: Oxford University Press.

Smith, G. (2014) "Closing the Loop: Testing Newtonian Gravity, Then and Now," in *Newton and Empiricism* ,ed. Zvi Beiner and Eric Schliesser, Oxford University Press, 2014, pp. 262-351.

Spirtes, P., Glymour, C. and Scheines, R. (1992) *Causation, Prediction and Search*. Berlin: Springer- Verlag.

Strevens, M. (2000) "Do Large Probabilities Explain Better?" *Philosophy of Science,* 67: 366–90.

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

Watkins. J. (1984) *Science and Scepticism.* Princeton: Princeton University Press.

Woodward, J. (1989) "The Causal/Mechanical Model of Explanation," *Scientific Explanation*. Minnesota Studies in the Philosophy of Science, volume 13 (1989), eds. Philip Kitcher and Wesley Salmon, 357-383.

Woodward, J. (1990) Supervenience and Singular Causal Claims." In *Explanation and Its Limits*. (Royal Institute of Philosophy Conference), ed. Dudley Knowles. Cambridge University Press.

211-246.

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

Woodward, J. (2018) "Some Varieties of Non-Causal Explanation" In *Explanation Beyond Causation* ed. Juha Saatsi and Alex Reutlinger. Oxford University Press, pp 117-137.

Woodward. J. (2021) *Causation with a Human Face: Normative Theory and Descriptive Psychology*. New York: Oxford University Press.

Woodward, J. and Ross, L. (2022) "Scientific Explanation" Stanford Encyclopedia of Philosophy.

Woodward, J. (2023) “Sketch of Some Themes for a Pragmatic Philosophy of Science” in *The Pragmatists Challenge*. (ed. Andersen and Mitchell) Oxford: Oxford University Press.