

# Inference to the Best Explanation and the New Size Elitism<sup>1</sup>

Katrina Elliott

Inference to the best explanation supposedly directs us to be most confident in whichever hypothesis *would, were it true*, provide the best explanation of our evidence. Call this the “subjunctive conception” of IBE. Inference to the best explanation also supposedly directs us to be most confident in whichever hypothesis best exemplifies *the theoretical virtues*, such as antecedent plausibility, lack of *ad hoc* elements, simplicity, and breadth of scope/unification. Call this the “theoretical conception” of IBE.

I argue that inference to the best explanation cannot do both. The problem is that the theoretical virtues are not indicative of how well a theory would explain, were it true. Hypotheses that score poorly with respect to the theoretical virtues would, were they true, explain just as well as would their more theoretically virtuous competitors.

The conservative solution is to abandon the subjunctive conception of IBE, while retaining the idea that we ought to choose theories on the basis of the theoretical virtues. Rejecting the subjective conception of IBE makes little difference to the most well-researched debates surrounding IBE. The subjunctive conception of IBE has not been entirely harmless, however. I close by considering the consequences its rejection has for a debate concerning the nature of indeterministic scientific explanation.

---

<sup>1</sup> Thanks go to Michaela McSweeney, as well as to audiences at Princeton University and Syracuse University.

## 0. Introduction

Gilbert Harman's seminal essay "The Inference to the Best Explanation" characterizes the titular inference twice over. Harman describes inference to the best explanation ("IBE") as recommending that we believe theories that would, were they true, do a particularly good job of explaining our evidence. He writes, "... one infers, from the premise that a given hypothesis would provide a "better" explanation for the evidence than would any other hypothesis, to the conclusion that the given hypothesis is true." (Harman 1965, p. 89) I'll call this conception of IBE "subjunctive", since it describes IBE as involving judgements about how well various competing hypotheses *would* play the role of explainers.

Harman then notes that his description of IBE raises the question of how to rank hypotheses in terms of their potential explanatory quality. He writes, "Presumably such a judgement will be based on considerations such as which hypothesis is simpler, which is more plausible, which explains more, which is less *ad hoc*." (p. 89) I'll call this conception of IBE "theoretical", since it describes IBE as involving judgements about how well various competing hypotheses exemplify the theoretical virtues.

Harman assumes that these two conceptions are compatible descriptions of a single inference, and so must also assume that the theoretical virtues are guides to which hypotheses would better explain our evidence were they true. But they are not. After further clarifying the subjunctive and theoretical conceptions of IBE, I show that choosing theories on the basis of commonly cited theoretical virtues systematically conflicts with choosing theories on the basis of how well they would explain our evidence were they true. Moreover, when the theoretical and subjunctive conceptions conflict, the theoretical conception's verdict is the more plausible. So, faced with a choice between inferring on the basis of theoretical considerations or inferring on

the basis of otherworldly explanatory potential, we should make the inferences endorsed by the theoretical conception of IBE.

Tellingly, rejecting the subjunctive conception would leave most philosophical discussions of IBE uninterrupted. For example, much attention has been given to the question of whether Bayesianism and IBE are compatible.<sup>2</sup> However, the subjunctive conception of IBE does no work in this debate, which concerns whether and how using theoretical virtues as guides to theory choice is compatible with Bayesianism. Philosophers also argue over which theoretical virtues should guide theory choice and why (see e.g., Kuhn 1977, Callebaut 1993, Longino 1995, Sober 2015). Here again, considerations about which hypotheses would better explain were they true do no work at all. In particular, no arguments are given that the theoretical virtues should guide theory choice *because* they track which hypotheses would best explain were they true. Though the subjunctive conception of IBE is regularly voiced in discussions of IBE, it is usually the theoretical conception that does the lifting, heavy or otherwise.

There is at least one debate, however, to which the subjunctive conception is central. Michael Strevens (2000) and Bradford Skow (2013) have argued for a view about scientific explanation in indeterministic contexts that Strevens named “size elitism”: all else being equal, explanations of likely occurrences are better than are explanations of unlikely occurrences. Elsewhere, I have argued for a view of scientific explanation that contradicts size elitism. (Elliott, *forthcoming*.) Here, I argue that a central motivation for size elitism dissolves once the subjunctive conception of IBE is rejected.

---

<sup>2</sup> Van Fraassen 1989 notably argues that Bayesianism and IBE are incompatible, but many philosophers disagree: perhaps IBE puts constraints on our priors (see e.g., Weisberg 2009), or perhaps IBE is a heuristic guide to Bayesian inference (see e.g., Okasha 2000, Lipton 2004, Dellsén 2018).

## 1. The Subjunctive Conception of IBE

The subjunctive conception of IBE has received only modest revisions since its articulation in Harman 1965. In a highly influential discussion nearly forty years later, Peter Lipton lands on a similarly subjunctive conception of IBE: “[Inference to the best explanation] claims that the explanation that *would, if true*, provide the deepest understanding is the explanation that is likeliest to be true”. (Lipton 2004, p. 77, my emphasis)

There are two improvements made by Lipton’s preferred gloss.

First, Lipton’s characterization of IBE reflects deference to van Fraassen’s well-known argument that IBE cannot warrant absolute confidence in the hypotheses it recommends. (van Fraassen 1989) That a particular hypothesis is best among a list of historically given hypotheses does not imply that it is good; it might instead be the best of a “bad lot”. At most, then, it seems that IBE warrants the judgement that the hypothesis that best explains is more likely to be true than its (considered) competitors, rather than the judgement that the hypothesis that best explains is likely to be true to any particular degree.

Second, Lipton offers a new gloss on the sense of “better” explanation that’s meant to be operative in IBE. For Lipton, IBE counsels us to prefer candidate explanations that would offer “deeper understanding” rather than “better explanation”. One nice feature of this gloss is that it helps to ward off various misunderstandings of IBE. For example, notice that, in calling one explanation “better” than another along some dimension, we need not mean that the explanation is better than its rival *at playing an explanatory role*. Compare: it’s good for NFL quarterbacks to be socially active and aware. Along this measure, Colin Kaepernick is a better quarterback than is Patrick Mahomes. It does not follow, however, that being socially active and aware

makes one better *at the role of quarterback*. Mahomes is better than Kaepernick at playing the role of quarterback. Similarly, there are endless ways in which one candidate explanation might be better than another that have nothing to do with how well those candidate explanations would perform the role of explainer were they true.

Here's one such way. There is a perfectly good sense in which candidate explanations are "better" the more likely they are to be true. However, this cannot be the sense of "better" invoked by IBE. While it is certainly good advice to be more confident in more likely hypotheses, the advice given by IBE is meant to be more substantive. As Lipton puts it, "we want a model of inductive inference to describe what principles we use to judge one inference more likely than another, so to say that we infer the likeliest explanation is not helpful." (Lipton 2004, pg. 76) Since there seems not to be a perfectly good sense in which candidate explanations convey "deeper understanding" the more likely they are to be true, Lipton's "deeper understanding" is perhaps less apt to be confused with "more likely to be true" than is Harman's "better explanation".

Here's another way. One role that scientific explanations play is in serving as the basis for appropriate answers to context-sensitive why-questions. (See e.g., Bromberger 1966, van Fraassen 1983, Lewis 1986.) To illustrate, it is helpful to invoke Peter Railton's distinction between an ideal explanatory text and an act of explaining. (Railton 1981) Ideal explanatory texts contain all information that is explanatorily relevant to an explanandum. Depending on one's theory of scientific explanation, an ideal explanatory text for a particular explanandum might be so vast and complex that no actual scientific community would or could produce such a thing. In contrast, individual acts of explaining (in which scientists and laypeople both

participate) aim at answering why-questions by providing relevant information about portions of the ideal explanatory text that are selected by context.

What portions of an ideal explanation are best invoked to answer particular questions varies from context to context, and depends on factors such as the background information, interests, and abilities of the relevant audience. One way an *act of explaining* might be better or worse, then, is by doing a better or worse job of conveying contextually appropriate information about an ideal explanatory text.

Examples are easy to generate. Philosophers like David Lewis think that the entire causal network leading up to an event's occurrence appears in the ideal explanatory text for that event. (Lewis 1986) Suppose I've recently been promoted at work and we want to know why. My recent award for punctuality is a node in the causal chain leading to my promotion at work, but so too are many distant events, such as my grandparents' marriage (which is a cause of my parents' births, which were causes of my birth, which is a cause of my being so punctual, which...). For Lewis, then, citing my grandparents' marriage as a (partial) explanation of my promotion at work is strictly speaking correct (since the marriage is a cause of the promotion), but there is nevertheless a sense on which it is "better" to cite my recent punctuality award (given the context). Similarly, it is better to explain the world in macrophysical rather than microphysical terms to those who cannot wrap their heads around modern physics. And, it is better to explain the brush fire by telling me about the discarded cigarette rather than about the oxygen-rich atmosphere, which I already knew was present.

Importantly, IBE does not invoke this context-sensitive sense of "better" explanation. After all, IBE is a method for choosing between incompatible candidate explanations, whereas more or less contextually-appropriate answers to why questions need not be competitors.

Though it is better to explain my recent promotion by citing my award rather than by citing my grandparents' marriage, it is nevertheless true that my grandparents were married and that their marriage was a cause of my recent award. Obviously, IBE does not council us to infer from my promotion that I receive an award *rather than* that my grandparents were married. So, the sense of "better" that is operative in the subjunctive conception of IBE is not the sense in which some acts of explaining do a better job of conveying appropriate portions of the ideal explanatory text than do others. To my ear, at least, it is harder to mistake Lipton's "deeper understanding" than Harman's "best explanation" for being a concept that picks out only the most contextually appropriate acts of explaining.

So far we have been discussing what is not meant by Harman's "better explanation" or Lipton's "deeper understanding". Can we say something both substantive and positive about how to understand these notions? We might try turning to philosophical theories of scientific explanation for guidance about the particular explanatory role invoked in IBE and what it takes to perform better or worse in that role.

The problem is that IBE must invoke a generic and theory-neutral conception of the explanatory role. A subject using IBE is supposed to arrive at the same inferential conclusions whether or not she is, say, a Kitcher-style unificationist (Kitcher 1989), a Railton-style mechanist (Railton 1981), a Woodwardian causal theorist (Woodward 2003), or someone who has simply not thought much about the specific roles scientific explanations play. So, IBE cannot invoke a sense of "better" or "deeper" explanation that is sensitive to one's preferred theory of scientific explanation. Relatedly, advocates of IBE (including Harman and Lipton) claim that the inference vindicates both widely-endorsed commonsense and scientific inferences. I assume that lay people and scientists often have quite different conceptions of explanation in mind when

picking between candidate explanations. So, we have yet another reason to interpret IBE as invoking a broad and theory-neutral conception of the role of explanation.

Harman, Lipton, and others are happy to rely on their audience's intuitive pre-theoretic understanding of "better explanation" and "deeper understanding" when formulating and defending IBE. It seems I won't be playing fair, then, unless I also leave the meaning of "better explanation" and "deeper understanding" as pre-theoretic and vague. However, should you start to feel as though you do not have a firm enough grip on what is meant by "better explanation" or "deeper understanding" to evaluate the cases that follow, I suggest that you take this as further evidence against the subjunctive conception. As we will see in the next section, the theoretical conception of IBE does not involve making any judgements about "better explanations" or "deeper understanding".

Like Lipton, I am convinced by van Fraassen's "best of a bad lot" argument, and so happily take on his first improvement on Harman's initial formulation of IBE. Unlike Lipton, I am unsure whether glossing the relevant sense of "better explanation" in terms of "depth of understanding" is the best option for defenders of the subjunctive conception of IBE, but it might be an improvement so I'll take it on as well. Amending Harman's initial subjunctive characterization of IBE in light of Lipton's two improvements gives us something like the following:

**Subjunctive Conception of IBE:** When considering, in light of your evidence, which of several hypotheses to believe, be most confident in whichever hypothesis would, were it true, provide the deepest understanding of your evidence.



## 2. The Theoretical Conception of IBE

Harman's presumption that items off the familiar list of theoretical virtues also serve as guides to which hypotheses would provide better explanations (in the relevant sense of "better") has gone largely unchallenged.<sup>3</sup> Thagard (1978, p. 89) writes, "Inference to the best explanation is inference to the theory that best satisfies the criteria of consilience and simplicity, as well as a third: analogy." Lipton (2004, p. 138) writes, "But there is considerable agreement over the identity if not the analysis of many inferential virtues, and it is striking how many of these appear also to be explanatory virtues... Thus among the inferential virtues commonly cited are mechanism, precision, scope, simplicity, fertility or fruitfulness, and fit with background belief. All of these are also plausibly seen as explanatory virtues." Henderson (2014, p. 690) writes, "It is often suggested that a theory provides the best explanation if it possesses the optimal combination of explanatory virtues in relation to the phenomena, where explanatory virtues include considerations such as simplicity, unification, scope, and fruitfulness." Dellsén (2018, p. 1747) writes, "What makes one explanation better than another is seldom spelled out in detail, but is generally taken to depend on various factors known as *explanatory considerations*. These include: explanatory power... antecedent plausibility... simplicity, fecundity, testability, avoidance of ad hoc elements, and explanatory depth..."

Thus, advocates for IBE commonly assume that familiar theoretical virtues are guides to which explanations are better or worse (in the relevant sense), and so commonly conceive of IBE as an inference to theories that best exemplify familiar theoretical virtues. In other words,

---

<sup>3</sup> One notable exception is Barnes 1995, who expresses skepticism about whether *all* traditional theoretical virtues track how theories would explain were they true. I use different arguments than does Barnes, and draw some different conclusions.

advocates for IBE commonly voice the following conception of IBE alongside the subjunctive conception:

**Theoretical Conception of IBE:** When considering, in light of your evidence, which of several hypotheses to believe, be most confident in whichever hypothesis best exemplifies the theoretical virtues.

Versions of the theoretical conception of IBE have also been voiced (though differently named!) by philosophers who do not explicitly endorse *any* connection between explanation and inference. The idea that theoretical considerations other than empirical adequacy do, and should, drive theory choice is deeply entrenched within analytic philosophy of science, dating back at least to Pierre Duhem. I do not intend to challenge, in any way, the venerable thesis that extra-empirical theses play an important and even truth-tracking role in scientific theorizing. Instead, I aim to argue that it was a mistake to entangle that thesis with vague talk of how well various theses would explain were they true.

### 3. Incompatibility Between the Subjunctive and Theoretical Conceptions

To argue that the subjunctive and theoretical conceptions of IBE are mutually incompatible, I'll present four cases. In each case, I argue that the theoretical conception ranks one of two hypotheses under consideration as being more likely, while the subjunctive conception of IBE ranks both hypotheses as being equally likely. Before turning to these cases, however, a few more clarificatory remarks are in order.

First, recall that applying the subjunctive conception of IBE requires making a comparative judgement about how “deeply” we understand our evidence in light of incompatible hypotheses. And, recall that you’ve been instructed to use whatever pre-theoretical understanding of “depth of understanding” you came in with, subject to a few qualifications (see section 2). Here is one more qualification: the sense of “depth of understanding” invoked by the subjunctive conception of IBE must be idealized in such a way that it need not track the actual understanding of any actual individual. For example, I speak a little Spanish but no German; I’ll have a much deeper understanding of a phenomenon if you explain it to me in Spanish than if you explain it to me in German. Clearly, however, IBE does not recommend that I believe Spanish sentences over German ones. Instead, the subjunctive conception of IBE should be understood as recommending hypotheses on the basis of how much understanding each would provide to an ideal agent who fully understands these hypotheses and their implications.

Second, the theoretical conception of IBE is more of a thesis-schematic than an actual thesis. Turning the thesis-schematic into a thesis requires deciding which are the relevant theoretical virtues. For brevity’s sake, I’ll focus on four of the most commonly cited: antecedent plausibility, avoidance of *ad hoc* elements, simplicity, and broad scope/unification. Once we get some cases on the table, it should become clear that the disagreement between the subjunctive conception and the theoretical conception is not an artifact of my choice of theoretical virtues. While one could surely cook up new “virtues” that do successfully track our pre-theoretical judgements about which hypotheses would best explain were they true, I argue that my commonly cited theoretical virtues would not be among them and that the new “virtues” would be poor guides to theory choice.

Finally, my goal is to convince you that the theoretical virtues do not track our pre-theoretic judgements about how well various hypotheses would explain were they true. However, if I were to just show that the theoretical virtues *considered jointly* favor different hypotheses than do subjunctive explanatory considerations, you might wonder whether some slightly different version of the theoretical conception (formulated, perhaps, by leaving out some of the virtues I consider) might yet align with the subjunctive conception. Accordingly, each of the cases that follow isolates a single theoretical virtue and presents two hypotheses that differ with respect to how well each exemplifies *that* virtue. In other words, if each of my four cases convinces you, then you'll also have been convinced that the subjunctive and theoretical conceptions of IBE cannot be brought into alignment by invoking *any* of the four theoretical virtues I consider.

### 3.1 Antecedent Plausibility

Creationists, who maintain that the Book of Genesis is a literally true description of the creation of the universe, argue that the universe is relatively young (i.e., only 6,000-10,000 years old). To establish the empirical adequacy of creationism, they offer non-standard explanations of a wide variety of physical phenomena. For example, if the universe is very young, why can we see stars that are separated from us by such massive distances? In a young universe, light will seemingly not have had time to travel from the distant stars to us. Creationist Barry Setterfield (Setterfield 2013) has an answer: we can see the distant stars despite the young age of the universe because the speed of light has been slowing since the creation of the universe.

Let's call Setterfield's answer "Hypothesis 1", or "H1" for short:

**H1:** The universe is 6,000-10,000 years old and the speed of light has been slowing since the creation of the universe.

Of course, Setterfield's view has not been widely accepted. Instead, most of us are very confident that the speed of light is constant in a vacuum and that universe is around 13.8 billion years old. Let's call this rival hypothesis "H2":

**H2:** The universe is around 13.8 billion years old and the speed of light is constant in a vacuum.

Now, imagine that someone learns for the first time that some of the stars visible in the night sky are millions of lightyears away.<sup>4</sup> In other words, imagine that someone acquires the following new evidence, which I'll call "E1" for short:

**E1:** Some visible stars are millions of lightyears from Earth.

Suppose that, prior to learning the distances of various visible stars, our individual was skeptical of H1 and confident in H2. In other words, suppose that H2 is more antecedently plausible for our agent than is H1. In light of their new evidence, which should our subject regard as more likely: the hypothesis that the universe is young and the speed of light has been

---

<sup>4</sup> Setterfield will think that the label "lightyear" is misleading, but we could easily use some non-question begging standard of measurement to pick out that same distance.

slowing down since the creation of the universe, or the hypothesis that the universe is old and the speed of light in a vacuum is constant?

Let's first answer by applying the theoretical conception of IBE. I have already stipulated that, for our agent, H2 is antecedently more plausible than H1. I further assume that the remaining theoretical virtues are equally possessed by both hypotheses. In other words, I assume that there are no relevant theoretical considerations I have failed to take into account that, on balance, favor H1.

The theoretical conception of IBE tells us to be more confident in more virtuous hypotheses and less confident in less virtuous hypotheses. But I take it that H1, i.e., the hypothesis that the speed of light has been decreasing, is not a very theoretically virtuous hypothesis because it fits so poorly with other things we, and our hypothetical agent, believe. In other words, H1 scores miserably with respect to antecedent plausibility. Since I have assumed that H1 has no further hidden virtues that will rank it above H2, I conclude that the theoretical conception recommends H1.

What does the subjunctive conception of IBE recommend? To answer that question, we suppose, in turn, that each hypothesis under consideration is true. We then ask, still under that supposition of truth, how deep an understanding of our evidence is conveyed by each hypothesis.

Let's start with H2. Were I to know that the universe is old and the speed of light is constant in a vacuum, how deeply would I understand why I can see stars that are millions of lightyears away? Decently deeply, I'd say. Of course, H2 doesn't tell me anything about how vision works, how stars produce light, how much of space is a vacuum, how the universe began, why the speed of light is constant, etc. Still, I get the basic idea; the stars are very far away but

also very old, and the speed of light in a vacuum is constant but sufficiently fast that light has had time to travel from the distant stars to me.

Next, the creationist's hypothesis H1. Were I to know that the universe is very young and the speed of light has been decreasing since the creation of the universe, how deeply would I understand why I can see stars that are millions of lightyears away? Just as deeply as I would were H2 instead true, it seems to me. Like H2, H1 leaves a lot of questions unanswered. H1 does not tell me how vision works, how stars produce light, how much of space is a vacuum, how the universe began, why the speed of light has been slowing, etc. Nevertheless, I once again get the basic idea: the stars are very far away and very young, but light used to travel so much faster than it does now that light had plenty of time to reach Earth in 6,000-10,000 years. So, were I to know H1, I think I would understand why I can see the distant stars (E1) just as well as I would were I to know instead that the universe is very old and the speed of light is constant in a vacuum (H2). I conclude that the subjunctive conception of IBE recommends that our subject be equally confident in both hypotheses.

We have, then, our first example of the incompatibility of the two conceptions of IBE: the theoretical conception says our subject should be more confident in H2 than in H1 in light of evidence E1, and the subjunctive conception says our subject should be equally confident in H2 and H1 in light of evidence E1. We also have our first example of the theoretical conception giving better advice than does the subjunctive conception. Though H1 and H2 are both consistent with E1, our subject should take E1 to further reinforce her antecedent confidence that H1 is true and H2 is false—at least, that's what I would suggest. So, here we have our first case in which the theoretical and subjunctive conceptions of IBE give different advice and the subjunctive conception's advice is bad.

In retrospect, we should not be surprised that evaluations of a hypothesis's antecedent plausibility pull apart from evaluations of how deeply a hypothesis would explain were it true. Antecedent plausibility is simply a measure of how likely a subject finds a hypothesis prior to her new evidence. That we, who live in actuality, do not presently believe a particular theory is hardly a reason to doubt the theory's ability to explain if only it were true. The point is even more obvious when we consider examples from the history of science. For instance, classical mechanics is a theory we currently regard to be false. Nevertheless, had classical mechanics been true, presumably it would have provided deep explanations of a wide variety of physical phenomena. Commonsense examples offer still further support. Suppose that the detective suspected the butler from the start—even more so now that the butler's shoeprints were found near the body. Nevertheless, *were* we to know that the maid framed the butler by committing the murder in his shoes, then we would just as well understand the presence of the butler's shoeprints.

### 3.2 Avoidance of *ad hoc* elements

Uri Geller is an illusionist who has claimed to have various paranormal abilities, such as psychokinesis and telepathy. In 1973, Johnny Carson invited Geller to display these abilities on the Tonight Show. Renowned magician and skeptic James Randi gave Carson specific instructions about how to prepare the props so that Geller could not cheat by using well-known magician's techniques. Let's call Randi's hypothesis that Geller has no paranormal abilities, "H3":

**H3:** Uri Geller has no paranormal abilities.



When time came for Geller to perform, he was unable to demonstrate any of his supposed paranormal abilities. Let's call this fact "E2":

**E2:** Uri Geller failed to display any paranormal abilities on the Tonight Show.

The problem, claimed Geller, was simply that he was having an off night. We'll call this hypothesis "H4":

**H4:** Uri Geller has paranormal abilities, but he was having an off night.

Now imagine someone watching this episode of the Tonight Show for the first time. Unlike last time, let's stipulate that the rival hypotheses are equally antecedently plausible to our agent. Because I intend for this case to isolate the theoretical virtue of avoiding *ad hoc* elements, I'll also assume that H3 and H4 are tied with respect to the remaining theoretical virtues. Given these assumptions, in light of E2, should our subject be more confident in H3 or H4?

As before, we'll first answer by applying the theoretical conception of IBE. While it is notoriously difficult to give an adequate analysis of what it means for a hypothesis to be *ad hoc*, the general idea is that a hypothesis is *ad hoc* to the degree that it has been modified to accommodate specific and otherwise disconfirming evidence. I take it that whatever else might be said about the notion of an *ad hoc* hypothesis, the hypothesis that Geller has psychokinetic abilities but simply happened to be off on the one night that James Randi controlled his props is a paradigm example of the vice. I conclude that the theoretical conception of IBE favors the

hypothesis that Geller is not telekinetic (H3) over the hypothesis that he is but was having an off night (H4) in light of Geller's failure to perform on the Tonight Show (E2).

As with the previous case, however, I see no reason at all for the subjunctive conception of IBE to recommend either hypothesis over the other. Starting with H3, we apply the subjunctive conception by asking how deeply we would understand Geller's failure (E2) if we were to know that he is not telekinetic (H3). Like our previous hypotheses H1 and H2, H3 leaves some questions unanswered. Here are just three relevant matters H3 is silent on: H3 does not tell us if telekinesis is possible; H3 gives us no indication as to why telekinesis is possible or why it is not; and, H3 does not tell us whether or why it is possible for Geller in particular to have telekinetic abilities. Nevertheless, H3 provides us with a decent understanding of Geller's failure. Geller failed to display any paranormal abilities on the Tonight Show simply because he has no such abilities to display.

However deeply I understand E2 in light of H3, H4's explanation of E2 seems just as deep. Were I to know that Geller really does have paranormal abilities but was simply having an off night (H4), then I think I would just as well understand Geller's failure to display his paranormal abilities on the Tonight Show (E2). H4 leaves the same issues unaddressed as does H3: H4 gives us no indication as to why it's possible for individuals to possess paranormal abilities; H4 is silent about whether people besides Geller have similar abilities; and, H4 has nothing to say about why Geller in particular has paranormal abilities. Still, it seems to me that I have the same depth of understanding of Geller's failure. Were I to learn that Geller really does have paranormal abilities but that he was simply having an off night, I would perfectly well understand his failure to perform. To be sure, I would not understand why his powers worked, or

how he acquire them. But the question is whether I would understand his failure equally well on both hypothetical suppositions, and I submit that I would.

Now we have our second example of the incompatibility between the theoretical and subjunctive conceptions of IBE. The theoretical conception tells us to be more confident in H3 than in H4 in light of E2, but the subjunctive conception tells us to be equally confident in H3 and H4. Once again, I side with the theoretical conception's recommendation. Even open-minded viewers of the Tonight Show should have lost confidence in Geller's claims to have paranormal abilities in light of his failure to perform during Randi's intervention.

I think it should be no surprise that our judgements of which hypotheses are *ad hoc* fail to track our judgements of how well various hypotheses would explain were they true. Once we've settled on what is true, it no longer matters what is *ad hoc*. We find it implausibly convenient that Geller would have an off night exactly when an expert like James Randi has prepared his props. So, we should doubt that H4 is true. However, *were* it in fact true that Geller had the bad luck to be off on the very night Randi tested him, then we would well understand Geller's failure to display his paranormal abilities. That we find it overly convenient that Geller happens to fail on the one night he is tested is a consideration that points in favor of the theoretical conception's advice. That same consideration, however, seems totally irrelevant to any further explanatory considerations once we have granted that Geller actually has paranormal abilities.

### 3.3 Simplicity

Spindrift is a Massachusetts beverage company that distributes carbonated water mixed with small amounts of real fruit and juice. The blackberry flavor is particularly delicious, but it is hard to find. My friend is coming over, and her drive takes her past one of the few local stores

that carries the blackberry flavor. I asked her pick me up an eight pack but I can see now that she is coming up the walk empty-handed.

Let the fact that my friend didn't bring me any Spindrift be our evidence:

**E3:** My friend did not bring me Spindrift.

Suppose that two hypotheses occur to me about why my friend didn't bring Spindrift. It might be that the store was out of Spindrift. We'll call this hypothesis "H5":

**H5:** The store was out of Spindrift and my friend brought her wallet.

But it might also be that the store was out of Spindrift and my friend forgot to bring her wallet, which we'll call "H6":

**H6:** The store was out of Spindrift and my friend forgot to bring her wallet.

In this example, I aim to isolate *simplicity*. Simplicity is yet another theoretical virtue that is difficult to analyze, but I hope it is clear that H6 is less simple than is H5. H5 describes only one reason for E3, but, according to H6, my friend's failure was *overdetermined*: the store was out of Spindrift (just as H5 says), but even if the store had not been out, she would not have brought me Spindrift because she forgot her wallet. I hope it is uncontroversial that, all else being equal, hypotheses that posit overdetermination are more complex than those that do not.

I stipulate that H5 and H6 are both equally antecedently plausible to me. I further assume that neither hypothesis is particularly *ad hoc*. Finally, I assume that the remaining theoretical virtues are equally well-exemplified by both H5 and H6. On the grounds that H5 is simpler than is H6, I conclude that the theoretical conception of IBE recommends H5.

As will now be anticipated, the subjunctive conception's advice conflicts with the theoretical conception's advice. Were I to know that the store was out of Spindrifit (H5), then I would perfectly well understand why my friend did not bring me any (E3). Were I to instead know that there was an additional reason that my friend did not bring me Spindrifit (i.e., that my friend forgot her wallet), I would understand why my friend didn't bring me Spindrifit at least as deeply as I would were I to know H5. A scenario in which my friend brought me no Spindrifit because the store was out is simpler than is a scenario in which my friend brought me no Spindrifit because the store was out *and* she forgot her wallet. In both scenarios, however, my friend's failure to bring me Spindrifit is equally well explained.

And, as will also now be anticipated, it is the theoretical conception that delivers the correct answer. I should be more confident in H5 than in H6 in light of E3, despite the subjunctive conception's recommendation that I regard H5 and H6 as being equally likely.

Finally, I once again think that we should not be surprised; in retrospect, it seems obvious that simplicity considerations are not guides to the relevant sense of "better explanation" or "deeper understanding" that the subjunctive conception invokes. Of course, limited as we are, it is easier *for us* to understand and to explain simple systems than it is for us to understand and to explain more complex systems. But surely there is no impediment, in principle, to an agent coming to have a thoroughly deep understanding, or an ideal explanation, of even very complex situations. Complex situations are explained just as well by complex theories as simple

situations are explained by simple theories. While there may be good reason to be skeptical of overly complex theories, I see no similarly good reason to doubt the potential of complex theories to explain the scenarios in which they are true.

### 3.4 Broad Scope/Unification

In our final case, I aim to isolate the closely related virtues of having broad scope and of being unified. Breadth of scope and unification are like two sides of the same coin. Roughly, to evaluate breadth of scope, we hold fixed a hypothesis (or axiom, or theory, or model, etc.) and evaluate the domain over which the hypothesis yields predictions; the broader a hypothesis, the wider and more diverse is its domain of application. To evaluate unification, we instead hold fixed a particular domain and evaluate how many independent hypotheses (or axioms, or theories, or models, etc.) are required to yield predictions over the entire domain. Roughly, more unified theories require fewer independent hypotheses to accommodate a domain.

Of all the theoretical virtues, unification might seem most plausibly connected to depth of understanding. After all, there are popular theories of scientific explanation on which the role of scientific explainer just is to “unify” some otherwise disparate domain of inquiry. If achieving scientific understanding is a matter of possessing unified theories, then don’t more unified theories always convey deeper understanding?

No. Maybe it is the case (though I think not) that the correct theory of scientific explanation is one on which whatever scientific theory best unifies a domain thereby provides the ideal explanation of that domain. On such a picture, explaining is a kind of systematization, and the best system (of our world) is also an ideal explanation (of our world). (See, e.g., Friedman 1974 and Kitcher 1989) Notice that, also on this picture, being true is a necessary

condition for being part of the best system. In other words, rival systems are not competing accounts of what is true. Instead, rival systems are competing ways of *organizing* the same truths. So, the notion of “unification” at work in unificationist theories of scientific explanation is a notion that presupposes what is true, and then ranks systems of truths with respect to how well they systematize those truths.

Plainly, such a notion cannot be used to make non-trivial modal comparisons of how well various theories would explain where they true, because false theories are immediately disqualified from being “unified” in the relevant sense. Just as plainly, such a notion also cannot correspond to any theoretical virtue that is useful for theory choice. “Unification”, when understood as a theoretical virtue, indicates which of several mutually inconsistent theories is most likely to be true. “Unification”, when understood as a constitutive feature of scientific explanation, holds fixed what’s true and indicates which collection of truths is explanatory. So, the theoretical virtue “unification” *is* the right kind of thing for making transworld comparisons, but *is not* the kind of thing that unificationists appeal to in their theories of scientific explanation. Defenders of the subjunctive conception of IBE thus cannot appeal to the authority of the literature on scientific explanation to claim that more “unified” theories (in the theoretical sense) would, were they true, make for better explanations or deeper understanding .

Of course, one might still insist that hypotheses that are broader or more unified in the theoretical sense also would provide deeper understanding were they true. Suppose that the ideally unified theory of world 1 is *more unified* than is the ideally unified theory of world 2. To defend the idea that the theoretical virtue of “unification” is a guide to otherworldly explanatory potential is to argue that the best explanation of world 1 is *better* or *deeper* than is the best explanation of world 2. We have already seen that there is no pressure from the literature on

scientific explanation to think that, e.g., world 1 is somehow more apt for better or deeper explanation than is world 2. Turning to our final case, I aim to show that our pre-theoretic intuitions likewise fail to motivate the idea that less unified worlds are less well explained or less deeply understood than more unified worlds.

I didn't know much about boxers before I started taking my own dog to the park, but I've noticed the following fact:

**E4:** All the boxers at my dog park have short tails.

I know that Dobermans are born with floppy ears—the erect, pointy-eared look commonly associated with Dobermans is actually the result of ear cropping. Perhaps something similar is true of boxers? Let's call that hypothesis "H7":

**H7:** All short-tailed Boxers are born with long tails, which are then typically cropped.

On the other hand, I also know that Australian shepherds, like boxers, characteristically have short tails. Most Australian shepherds are born with long tails which are then cropped, but approximately 1 in 5 are born with a natural bobtail. Perhaps boxers are like Australian shepherds, i.e., perhaps "H8" is true?

**H8:** Only some short-tailed boxers are born with long tails, which are then typically cropped. The remaining short-tailed boxers are born with natural bobtails.



In light of E4, should I be more confident in H7 or H8?

Both H7 and H8 are equally broad (because each covers exactly the domain of short-tailed boxers), but H7 is more unified than is H8. H7 accommodates my evidence, but posits only one kind of boxer: those born with long tails. H8 also accommodates my evidence, but posits two kinds of boxer: those born with long tails and those born with bobtails. As before, I assume that the remaining theoretical virtues are equally possessed by both hypotheses. So, according to the theoretical conception, that all the boxers at my dog park have short tails (E4) should make me more confident that all short-tailed boxers' tails are cropped (H7) than that only some short-tailed boxers' tails are cropped (H8).

The subjunctive conception, I claim, implies that both hypotheses are equally likely in light of our new evidence.<sup>5</sup> Were I to know that all short-tailed boxers are born with long tails, then I would understand why all the boxers at my dog park have long tails. Were I instead to know that some short-tailed boxers are born with long tails while other short-tailed boxers are born with bobtails, then I would have just as deep an understanding of why all the boxers at my dog park have long tails.<sup>6</sup> A world in which there is one kind of boxer is more unified than a world with two kinds of boxers (all else being equal), but both worlds are equally apt for explanation and understanding. Once again, when compared with the theoretical conception of IBE, the subjunctive conception gives different, and worse, advice.

---

<sup>5</sup> Barnes 1995 draws a similar conclusion.

<sup>6</sup> One might intuit that H7 provides a deeper understanding than H8 because H7 specifies how each dog came to have a short tail. To meet this objection, we can simply reformulate H8 so that it specifies exactly which dogs were born with long tails and which dogs were born with bobtails.

### 3.5 Other theoretical virtues?

One might agree with my descriptions of my cases, and yet still maintain that the subjunctive and theoretical conceptions of IBE are compatible. If successful, I have shown only that antecedent plausibility, avoidance of *ad hoc* elements, simplicity, and breadth of scope/unification are not guides to how well a theory would explain, were it true. So much the worse, perhaps, for the traditional list of theoretical virtues; perhaps the moral of my cases is not that the actual theoretical virtues fail to track otherworldly explanatory potential, but that the virtues I've considered are not the actual theoretical virtues.

The problem with this response is that, in each of the four cases considered above, it is the theoretical conception, and not the subjunctive conception, that yields the intuitively correct recommendation. My cases seem to show that the traditional theoretical virtues are better guides to theory choice than are considerations regarding how well various hypotheses would explain were they true. It will not be enough, then, for my opponent to come up with a list of features that successfully track how well a theory would explain were it true; she must also argue that these new features are better guides to theory choice than are the traditional theoretical virtues. Furthermore, her list of theoretical virtues cannot include even one of the four virtues considered above, since the theoretical and subjunctive conceptions can conflict if even one of the theoretical conception's virtues fails to track our judgements about how well various theories would explain were they true. Perhaps that can all be done, but it has not been done yet.

#### 4. IBE and the New Size Elitism

As I noted in my introduction, the subjunctive conception of IBE does surprisingly little philosophical work. One exception, however, arises in the context of a debate about the nature of scientific explanation in indeterministic contexts.

“Size elitism” is the view that likely occurrences have better explanations than do unlikely occurrences. (Strevens 2000 introduces the term.) The original size elitism is Carl Hempel’s (1965). Hempel argued that scientific theories explain by providing nomic grounds for predicting the fact to be explained. When an unlikely event occurs, there are no nomic grounds for predicting that the event would occur, and so Hempel’s view implies that unlikely occurrences have no explanations at all. The view that likely occurrences are explicable but unlikely occurrences are inexplicable is an example of size elitism.

One of the most famous counterexamples to Hempel’s view of scientific explanation is also a counterexample to his size elitism. Paresis is a neurological disorder that is found only in patients who have had syphilis. However, one’s chance of developing paresis given that one has latent, untreated syphilis is only around 25%. Michael Scriven (1962) famously objected to Hempel’s account of indeterministic explanation on the grounds that, though one should not expect any particular patient to develop paresis, every patient’s paresis is explained by their syphilis. And since the particular chance of paresis seems to make no difference to whether a patient’s paresis is explicable, Scriven’s case seems to defeat size elitism.

##### 4.1 The New Size Elitism

Michael Strevens and Bradford Skow have advocated for a new version of size elitism that can avoid Scriven’s counterexample. Strevens writes,

The paresis case shows that high probabilities are not necessary for explanation. But the nature of probabilistic explanation in statistical mechanics shows that high probabilities explain *better* than low probabilities. The correct account of explanation will underwrite both these conclusions. (Strevens 2000, p. 32)

Strevens's version of size elitism differs from Hempel's in that Strevens does not hold that unlikely occurrences are inexplicable. Instead, Strevens's view is that there is some sense of "better" on which unlikely occurrences are explicable, but likely occurrences are nevertheless better explained than are unlikely occurrences. I'll call this view "the new size elitism".

As we saw in section 1, there are many different dimensions along which explanations might be better or worse. Which dimensions are invoked by the new size elitism? Neither Strevens 2000 nor Skow 2013 contains a precise answer, but each asserts that their brand of size elitism, whatever it amounts to, is denied by Peter Railton's theory of indeterministic explanation. Briefly reviewing Railton's view, then, will help to clarify the new size elitism.

Railton advances a mechanistic account of scientific explanation. He writes,

The goal of understanding the world is a theoretical goal, and if the world is a machine—a vast arrangement of nomic connections—then our theory ought to give us some insight into the structure and workings of the mechanism, above and beyond the capability of predicting and controlling its outcomes. (Railton 1979, p. 208)

Explaining the future in light of the past, on Railton's picture, involves displaying the nomic connections that hold between the future and the past. When an event occurs by chance, these nomic connections take the form of causal relations, probabilistic laws, and chances. The ideal explanatory text for a chance occurrence contains a complete description of the nomic facts relevant to its occurrence, independent of what those nomic facts are. The ideal explanatory text for a likely occurrence, then, includes a description of the event's high chance of occurring. The ideal explanatory text for an unlikely occurrence, similarly, includes a description of the event's low chance of occurring. These texts convey understanding solely in virtue of accurately describing "a vast arrangement of nomic connections"; that some of these texts describe high chances while others describe low chances makes no difference to the quality of explanation or depth of understanding that each conveys.

However, consistent with Railton's view (and outlined in Railton 1981) are a number of ways in which explanations that cite high chances might be "better" than explanations that cite low chances.

First, Railton accepts that acts of explaining that cite high chances might be better, in a sense, than acts of explaining that cite low chances. To see why, note that giving a false answer to a why-question can be an adequate explanatory act if the false answer conveys accurate information about the content of the ideal explanatory text. For example, suppose that the ideal explanatory text for my getting paresis contains the fact that my chance of getting paresis is 25.056%. If we explain my paresis by reporting that my chance of getting paresis is 25%, we say something false. Nevertheless, our act of explaining might be perfectly appropriate for the context because we have sufficiently approximated the ideal explanatory text.

Because acts of explaining can be better or worse the closer they come to approximating the contents of the ideal explanatory text, acts of explaining that cite high chances are often better than acts of explaining that cite low chances—especially in deterministic contexts. If our world is deterministic, then the ideal text for any occurrence includes the fact that its chance of occurring is 1. If the ideal explanatory text ascribes chance 1 to every event, then, *ceteris paribus*, acts of explaining that cite high chances are better than acts of explaining that cite low chances, because higher chances always more closely approximate chance 1. Size elitism is meant to contradict Railton’s view, but Railton’s view admits that acts of explaining might be better when they cite higher chances. The new size elitism, then, must not be about context-sensitive acts of explaining.

Next, notice that Railton’s mechanistic view of scientific explanation is not about which explanatory conjectures one should believe, but rather about what *truths* appear in ideal scientific explanations. Accordingly, Railton need not deny that hypotheses that attribute a higher chance to an event’s occurrence are “better” in the sense that they are, *ceteris paribus*, more likely to be true. Furthermore, Railton need not deny that explanatory conjectures that attribute high chances have theoretical virtues not had by explanatory conjectures that attribute low chances. Thus, the new size elitism cannot merely be the view that explanatory hypotheses that cite high chances are more likely to be true or better exemplify the theoretical virtues.

Here is what Railton’s view seems not to allow: that the ideal explanatory text for a likely event conveys a better explanation or deeper understanding than does the ideal explanatory text for an unlikely event. Since ideal explanatory texts for both likely and unlikely events contain complete descriptions of the nomic connections relevant to those occurrences, neither conveys a better explanation or deeper understanding than the other. Suppose, for illustration, that some

particular atom decays during some particular time interval. Suppose further that, according to the ideal explanation of the atom's decay, its chance of decaying during that time interval is 10%. Railton's view implies that, were the atom's chance of decaying 90% instead of 10%, then the ideal explanation of the atom's decay would have been different, but no better or worse, than is the actual ideal explanation of the atom's decay.

Now we are in a better position to state the new size elitism:

**The New Size Elitism:** Had an event's chance of occurring been higher than it actually is, the event's ideal explanation would have been better, or would have conveyed a deeper understanding, than it actually does. Had an event's chance of occurring been lower than it actually is, the event still would be explicable but its ideal explanation would have been worse, or would have conveyed a more shallow understanding, than it actually does.

While I cannot be sure that the above gloss is precisely what Strevens and Skow had in mind, it fits well with their respective discussions in two ways: it is a view denied by both Hempel and Railton, and it is a view naturally described as the claim that likely occurrences are "better explained" than are unlikely occurrences. It is also a view motivated by the subjunctive conception of IBE.

#### 4.2 From the Subjunctive Conception to New Size Elitism

Some of the arguments Strevens and Skow give for the new size elitism are independent of the subjunctive conception of IBE, but these are easily defanged. For example, Strevens uses

the term “size egalitarian” to refer to views like Railton’s, and argues that size egalitarians cannot endorse the historically important claim that molecular theory combined with statistical mechanics better explains thermodynamic phenomena than does molecular theory on its own.

He writes,

What was novel in statistical mechanics was the introduction of the statistical part of the explanation, so as to show that the kinds of processes posited by the molecular theory were highly likely to give rise to the observed behavior of gases and heat. Thus what gave the old molecular theory a degree of explanatory power that matched, and eventually exceeded, that of the caloric theory, was probability—the high probability that SM ascribed to the behavior to be explained. *If egalitarianism is correct, however, showing that a process that generates some event might possibly occur explains the event just as well as showing that the process has a high probability of occurring...* Thus, egalitarianism completely fails to make sense of this aspect of the history of physics. (Strevens 2000, p. 374, my emphasis)

Strevens argument fails because its description of size egalitarianism is not accurate. As we have just discussed, Railton’s egalitarian view of scientific explanation is that an ideal explanation of an event contains a complete description of the nomic connections relevant to the event’s occurrence. In cases in which chances are among those nomic connections, an ideal explanation contains a complete description of those chances. If, say, our atom’s chance of decaying really is 10%, then an explanation that merely reports that decay is possible—or that the chance of decay is greater than 5%, or between 9%-11%, or...— does not fully reveal the



“structure and workings” of the decay event. In other words, explanations that merely report that an event is possible do not, on Railton’s view, explain as well or convey as much understanding as do explanations that report an event’s precise chance of occurring. Thus, *contra* Strevens, egalitarianism need not imply that reporting that an outcome is possible explains just as well as does reporting an outcome’s precise chance of occurring.

Skow endorses the Strevens argument, and offers an additional consideration in favor of the new size elitism. He writes,

Suppose we do not know the law governing the radioactive decay of this atom. We have two theories. Theory 1 assigns the decay a high chance per unit time; theory 2 assigns it a low chance per unit time. Which should we believe? Certainly theory 1... We should believe theory 1 because it is a better explanation of the quick decay than theory 2. But then egalitarianism is false. (Skow 2013, p. 122-123)

Skow is right that we should be more confident in theory 1 than theory 2 in light of the atom’s decay. However, the egalitarian has the resources to grant that we should be more confident in theory 1 without thereby granting that IBE favors theory 1. Bayesianism, for example, implies that evidence favors whichever hypothesis assigns that evidence the highest probability (all else being equal). Thus, the egalitarian can appeal to Bayesianism to explain why the atom’s decay favors theory 1 over theory 2.

That said, responding to Skow by appealing to Bayesianism admits to a tension between size egalitarianism and IBE. If size elitism is consistent with both Bayesian and IBE-style

accounts of theory choice, while size egalitarianism is consistent only with Bayesianism, then there seems to be a theoretical cost to denying size elitism. Is this cost real?

In light of our distinction between the theoretical and subjunctive conceptions of IBE, the answer depends on which conception of IBE we have in mind. The theoretical conception of IBE is insufficient to motivate size elitism, as the following version of Skow's argument illustrates:

- P1. In light of the atom decaying, IBE favors theory 1 over theory 2. (Grant)
- P2. If IBE favors theory 1 over theory 2, then theory 1 better exemplifies the theoretical virtues than does theory 2. (theoretical conception of IBE)

---

C. Theory 1 better exemplifies the theoretical virtues than does theory 2.

As we saw in section 4.1, egalitarians such as Railton need not deny conclusion C. Establishing C, then, is not sufficient to establish the new size elitism.

In contrast, the subjunctive conception of IBE, coupled with the assumption that IBE favors theory 1 over theory 2, does imply (an instance of) the new size elitism:

P1. In light of the atom decaying, IBE favors theory 1 over theory 2. (Grant)

P2\*. If IBE favors theory 1 over theory 2, then theory 1, were theory 1 true, would provide deeper understanding of the atom decaying than would theory 2, were theory 2 true. (Subjunctive conception of IBE).

---

C\*. Theory 1, had theory 1 been true, would provide a deeper understanding of the atom decaying than would theory 2, were theory 2 true.

Unlike C, C\* is denied by Railton and implied by the new size elitism. Were the subjunctive conception of IBE correct, then, Skow's argument would successfully show that there is a serious theoretical cost to denying the new size elitism. But the subjunctive conception of IBE is not correct, so rejecting it while endorsing the theoretical conception of IBE involves no cost at all.

## 5. Conclusion

I have argued that the theoretical conception of IBE conflicts with the subjunctive conception of IBE. When the two conceptions conflict, I have argued that the theoretical conception's recommendations are more plausible. For these reasons, we should reject the subjunctive conception while retaining the theoretical conception. Perhaps rebranding the theoretical conception of IBE as "inference to the best theory" would help avoid further confusion.

## Bibliography

- Barnes, Eric (1995). Inference to the loveliest explanation. *Synthese* 103 (2):251 - 277.
- Bromberger, S. (1966). 'Why Questions', in *Mind and Cosmos: Essays in Contemporary Science and Philosophy*, R. Colodny, (ed), Pittsburgh: University of Pittsburgh Press.
- Callebaut, Werner (ed.) (1993). *Taking the Naturalistic Turn, or How Real Philosophy of Science is Done*. University of Chicago Press.
- Dellsén, Finnur (2018). The heuristic conception of inference to the best explanation. *Philosophical Studies* 175 (7):1745-1766.
- Elliott, Katrina. Where are the chances? *Synthese*. Forthcoming.
- Friedman, Michael (1974). Explanation and scientific understanding. *Journal of Philosophy* 71 (1):5-19.
- Harman, Gilbert H. (1965). The inference to the best explanation. *Philosophical Review* 74 (1):88-95.
- Hempel, Carl (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. The Free Press.
- Henderson, Leah (2014). Bayesianism and Inference to the Best Explanation. *British Journal for the Philosophy of Science* 65 (4):687-715.
- Kitcher, Philip (1989). Explanatory unification and the causal structure of the world. In Philip . Kitcher & Wesley Salmon (eds.), *Scientific Explanation*. Minneapolis: University of . Minnesota Press. pp. 410-505.
- Kuhn, Thomas S. (1977). Objectivity, value judgment, and theory choice. In *The Essential*

- Tension: Selected Studies in Scientific Tradition and Change*. University of Chicago Press. pp. 320--39.
- Lewis, David K. (1986). Causal explanation. In David Lewis (ed.), *Philosophical Papers Vol. Ii*. Oxford University Press. pp. 214-240.
- (1994). Humean Supervenience Debugged. *Mind* 103 (412):473--490.
- Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). London: Routledge.
- Longino, Helen E. (1995). Gender, politics, and the theoretical virtues. *Synthese* 104 (3):383 - 397.
- Okasha, Samir (2000). Van Fraassen's Critique Of Inference To The Best Explanation. *Studies in History and Philosophy of Science Part A* 31 (4):691-710.
- Railton, Peter (1978). A deductive-nomological model of probabilistic explanation. *Philosophy of Science* 45 (2):206-226.
- (1981). Probability, explanation, and information. *Synthese* 48 (2):233 - 256.
- Scriven, M., 1962, 'Explanations, Predictions, and Laws', in *Scientific Explanation, Space, and Time* (Minnesota Studies in the Philosophy of Science: Vol. 3), H. Feigl and G. Maxwell (eds), 170–230. Minneapolis: University of Minnesota Press.
- Setterfield, Barry. (2013). *Cosmology and the Zero Point Energy*. Natural Philosophy Alliance, INC.
- Skow, Bradford (2013). The role of chance in explanation. *Australasian Journal of Philosophy* (1):1-21.
- Sober, Elliott (2015). *Ockham's Razors: A User's Manual*. Cambridge University Press.
- Strevens, Michael (2000). Do large probabilities explain better? *Philosophy of Science* 67 (3):366-390.

Thagard, Paul R. (1978). The best explanation: Criteria for theory choice. *Journal of Philosophy* 75 (2):76-92.

van Fraassen, Bas C. (1983). The Scientific Image. *Mind* 92 (366):291-293.

----- (1989). *Laws and Symmetry*. Oxford University Press.

Weisberg, Jonathan (2009). Locating IBE in the Bayesian Framework. *Synthese* 167 (1):125-143.

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