Explanatory Inference Un-Locke-d

Abstract

What is explanatory inference and how does it work? Some argue that explanatory inference should be analysed within a Bayesian framework. Others contend that explanatory inference and Bayesianism are incompatible. In this paper, I propose a middle ground between these two views: a new account of explanatory inference that goes beyond Bayesianism while remaining compatible with it. One way to do this is by denying the Lockean Thesis in favour of the Humean Thesis.

Keywords: belief, credence, inference to the best explanation, explanatory inference.

1. Introduction

Inference to the Best Explanation (IBE) is an argument form with the following structure:

1) E is a body of evidence.
2) H explains E better than any other hypothesis.
3) Therefore, H.

Philosophers who defend IBE often claim that it is a rational and ubiquitous form of inference that is essential to scientific reasoning (Thagard 1978; Psillos 2007; Douven 2013). Others defend closely related forms of inference, such as inference to the only explanation (Woodward 2025), inference to the implications of the n-best explanations (Dellsén 2017), inference to the best out of a good lot of explanations, and so on. The common underlying principle behind these inferences is that a good explanation is, all else being equal, more belief-worthy than a poor

explanation. I'll use the term 'explanationism' to refer to the view that IBE, or a closely related form of explanatory inference, is a rational form of inference.[1]

There are different interpretations of how explanatory inference works. Some explanationists claim that explanatory inference should not be interpreted probabilistically. Others claim that explanatory inference should be interpreted within a Bayesian framework. Others argue that explanatory inference is incompatible with Bayesianism and involves adopting diachronically incoherent credences.

In this paper, I sketch a different interpretation of explanatory inference, construing explanatory considerations as a factor that affects the relationship between justified belief and credence. The explanatory virtues of a hypothesis can make a hypothesis more belief-worthy without making it more probable.

2. Existing Accounts of Explanatory Inference

There are several different interpretations of how explanatory inference works, each with its own advantages and disadvantages. As what follows are typically presented as accounts of IBE specifically, rather than accounts of explanatory inferences in general, I will label them as interpretations of IBE in this section.

First, Full-Belief IBE is the view where IBE governs inferences to full beliefs.[2] Someone who believes that E, that H is a good explanation of E, and that H is the best explanation of E, ought to also believe that H is true.

Climenhaga (2017) argues that sometimes, the hypothesis that best explains the evidence at one level is inconsistent with the hypothesis that best explains the evidence at a different level

[1] My use of the term is different from McCain's (2015), where explanationism is a theory of justification.
[2] See (Harman 1965; Vogel 1998) for discussions of Full-Belief explanationism.

(I will discuss this in more detail in section 7). Thus, if we follow Full-Belief IBE and believe the best explanatory hypotheses at both of these levels, we believe an inconsistent set of statements. Dellsén (2017) argues that, sometimes, the best explanation is only slightly better than the second-best explanation and, in these cases, we ought not to believe the best explanation over the second-best explanation. Instead, we ought to believe those propositions which are entailed by the n-best explanations.

The other existing accounts of IBE concern how IBE affects credences. Darwin observes two species with similar features. This is a surprising piece of evidence. He notes, however, that if the two species share a common ancestor, that would explain why they have similar features. Consequently, the fact that they have similar features would not be surprising at all. Therefore, Darwin infers that the two species share a common ancestor.

This line of reasoning can be represented as an instance of Bayesian reasoning. Darwin increases his confidence in the common ancestry hypothesis because it increases the likelihood of a surprising piece of evidence, i.e., the similarity. But this raises the question: how would the explanationist behave any differently from the Bayesian?

According to Heuristic explanationists, IBE is a useful heuristic for non-ideal epistemic agents to approximate Bayesian reasoning. 'H is a good explanation of E' roughly means 'P(H) is high' or 'P(E|H) is high' or 'P(E|H) is higher than P(E).'[3] If we have premises like these, it follows from Bayes' Theorem that P(H|E) is high (or higher than P(H)). According to Heuristic explanationism, IBE concludes that our credence in H, upon learning E, should be $c \pm \theta$, which is

[3] Philosophers who discuss Bayesian interpretations of IBE disagree on whether H's status as a good explanation of E is reflected in both the priors and likelihoods or only in likelihoods. For instance, Samir Okasha (2000) argues that good explanations have high priors. In contrast, many other authors who write on IBE only discuss its effect on likelihood (Schupbach & Sprenger 2011; Climenhaga 2017; Trpin 2024).

high and our best approximation of P(H|E).[4] In short, the Heuristic explanationist would not behave any differently from the Bayesian.

Heuristic explanationism avoids Climenhaga's (2017) objection. But it is often charged with making IBE philosophically uninteresting (Psillos 2007; Weisberg 2009; Cabrera 2017). If there is nothing more to IBE than useful heuristics that non-ideal epistemic agents use to approximate Bayesian reasoning, it seems to be dispensable to our epistemic lives.

Next, we have Objective Bayesian explanationism, on which explanatory considerations provide normative constraints on the priors we use in Bayesian conditionalisation. All else being equal, the fact that common ancestry is a good explanation of similarity provides normative constraints for assigning high values to P(Common ancestry) or P(Similarity|Common ancestry). Thus, by Bayes' Theorem, someone who assigns probabilities according to these normative constraints will be more confident that two species share a common ancestor upon learning that they share similar features.[5]

Objective Bayesian explanationism avoids the charge of making IBE philosophically uninteresting. But it has its own problems. For instance, Cabrera (2017) argues that some of the features we consider as virtues of an explanation are irrelevant for confirmation. So, it's not clear that the goodness of an explanation provides normative constraints for assigning high values to Bayesian priors. Additionally, I've argued in a previous work (2024) that Heuristic explanationism and Objective Bayesian explanationism are incompatible with influential models of scientific explanation.

[4] See (Lipton 2004; McGrew 2003; Dellsén 2018) for discussions and defences of the heuristic account.
[5] See (Weisberg 2009; Poston 2014; Climenhaga 2017) for discussions and defences of the objective Bayesian account.

Another influential objection in the literature comes from Roche & Sober (2013). Suppose we have a dataset on the frequency of lung cancer among heavy smokers, casual smokers, and non-smokers, but don't know the causal relation between smoking and lung cancer. In this scenario, we can estimate the likelihood that Smith has cancer based on the frequency data. We can estimate that P(Smith has cancer|Frequency data & Smith is a heavy smoker) is 0.6, for instance. However, Roche & Sober (2013) argue that the hypothesis that smoking causes cancer makes no difference to this probability. That is, P(Smith has cancer|Frequency data & Smith is a heavy smoker & Smoking causes lung cancer) = P(Smith has cancer|Frequency data & Smith is a heavy smoker) = 0.6. From this equality, it follows that [Smoking causes cancer] is neither confirmed nor disconfirmed by the frequency data, even though it would explain the frequency data. Roche & Sober (2013) conclude, from this case, that explanatoriness is evidentially irrelevant. This objection is problematic for Objective Bayesian explanationism because it purports to show that an explanatory claim does not provide normative constraints for assigning some high value to the likelihood that Smith has cancer.

The fourth version is Ampliative explanationism. On this version, explanatory considerations add probabilistic bonuses to a hypothesis after Bayesian conditionalisation. If Darwin had been an Ampliative explanationist, he would first perform Bayesian conditionalisation and then further increase his confidence in the common ancestry hypothesis upon noting that it is the best explanation for similarity. Thus, Darwin's final credence in the common ancestry hypothesis would be greater than P(Common ancestry|Similarity).[6]

Ampliative explanationism makes IBE a distinctive rule of inference that is not reducible to Bayesianism, it makes IBE philosophically interesting, and it is also compatible with extant

[6] See (Douven 2013; Douven and Schupbach 2015) for defences of the ampliative account.

models of scientific explanation, insofar as no extant model prohibits giving bonus probabilistic points to the best explanation. However, Ampliative explanationism is incompatible with Bayesianism because it leads to diachronic incoherence. Many consider this to be a sufficient reason for rejecting the ampliative account (Okasha 2000; McGrew 2003; Huemer 2009; Weisberg 2009; Poston 2014).

There is room, therefore, for a new version of explanationism. I offer a novel account of explanatory inference, which I call Belief-Justifying explanationism. I will argue that Belief-Justifying explanationism is consistent with Bayesianism, consistent with extant models of scientific explanation, and it gives explanatory inference an epistemic role that cannot be reduced to Bayesian confirmation. The way to achieve this is by denying the Lockean Thesis and accepting the Humean Thesis instead. However, I want to emphasise that my primary aim is not to settle the debate on whether the Lockean Thesis or the Humean Thesis is true. Rather, it is to show explanationists that the grass is greener on the Humean side of the fence.

3. Belief, Credence, and the Lockean Thesis

Let's briefly revisit Roche & Sober's (2013) argument that explanatoriness is evidentially irrelevant. I will grant their claim that P(Smith has cancer|Frequency data & Smith is a heavy smoker & Smoking causes cancer) = P(Smith has cancer|Frequency data & Smith is a heavy smoker). From this, it follows that [Smoking causes cancer] is not confirmed by the frequency data, even though it explains the frequency data, and that explanatoriness is evidentially irrelevant (at least in the Bayesian conception of evidence). However, it does not follow that explanatoriness is irrelevant to the formation of justified beliefs. At least, not if beliefs and credences are distinct.

Belief-Justifying explanationism requires distinguishing between beliefs and credences. There are different epistemological views on the relationship between belief and credence. One of them is the following:

> Lockean Thesis: S is justified in believing that P iff S has a rational credence in P at or above a certain threshold.

Other authors use the term differently. There are versions where the Lockean Thesis only posits a necessary condition for S to have justified belief, and versions where the Lockean Thesis posits a sufficient but not necessary condition. Another variation states that S would have an epistemic duty to believe if/iff their rational credence in P is at or above the threshold. Additionally, there are variations where the Lockean Thesis concerns rational belief rather than justified belief. For my project, I prefer to talk in terms of justified belief, but I think the arguments here can apply, *mutatis mutandis*, if the Lockean Thesis were framed in terms of rational belief.

The Lockean Thesis is a normative principle about when an agent is justified in believing a proposition. Epistemologists who discuss the Lockean Thesis disagree on the nature of this threshold. Some philosophers argue that belief requires having credence 1 (Clarke 2013; Greco 2015; Dodd 2017). But many prefer lower thresholds. Some argue that there is a universal threshold required for rational belief (Lee 2017; Shear and Fitelson 2019). Others argue that the threshold may shift depending on contextual factors (Weatherson 2005; Bach 2008; Ganson 2008; Fantl and McGrath 2010; Pace 2011; Ross & Schroeder 2014).

The Lockean Thesis enjoys some initial plausibility. It seems that someone who has rational high credence in P is justified in believing that P. But like any other thesis in philosophy, there is plenty of room for disagreement. Belief-Justifying explanationism starts by denying the

Lockean Thesis, more specifically, by denying that rational high credence is sufficient for justified belief.[7] Below, I outline some of the arguments often raised against the Lockean Thesis.

First, we have the lottery paradox (Douven 2012; Leitgeb 2014a). Let *r* be the Lockean threshold, and *n* the number of tickets for a lottery. You buy a lottery ticket and your credence that you will lose is (*n*-1)/*n*. If this value is greater than *r*, you ought to believe that you will lose the lottery. Thus, as long as *r* is less than 1, there are possible lotteries such that you are justified in believing that your ticket will lose. But for any ticket in the lottery, your credence that that ticket will lose is (*n*-1)/*n*. So, you are justified in believing, of each lottery ticket, that it will lose. But you also believe that exactly one ticket will win. Thus, you have an inconsistent set of beliefs. Some argue that the lottery paradox shows that the Lockean Thesis is false (Douven 2002; Jackson 2020a).

Other epistemologists offer counterexamples against the Lockean Thesis. Here is an example from Buchak (2014). Suppose your iPhone is missing. There are two possible explanations: either Jake stole your iPhone or Barbara did. You have no evidence about who stole your iPhone and don't know either suspect well, but you know that men are 10 times more likely to steal iPhones than women. In this case, you ought to have high credence that Jake stole your iPhone (roughly 0.91). However, Buchak argues that you should not believe that Jake stole your iPhone. Furthermore, we would not convict Jake merely based on the statistic that men are 10 times more likely to steal iPhones than women (Enoch et al. 2012; Gardiner 2019). To make matters worse, we can push the case by making the probability higher. If Buchak is right about this kind of case, then we ought to deny the Lockean Thesis: rational high credence is not sufficient for justified belief. This is an instance of the problem of naked statistical evidence.

[7] For a helpful overview on the Lockean Thesis and the relationship between belief and credence, see (Jackson 2020b).

Buchak further argues that certain changes to the case that keep your rational credence constant may lead to different judgments. For instance, suppose men and women are equally likely to steal iPhones, but an eyewitness who is reliable 90% of the time claims that they saw Jake steal your iPhone. In the eyewitness case, you would be justified in believing that Jake stole your iPhone, even though your credence is slightly lower than in the original case.

There are, of course, ways to defend the Lockean Thesis in light of the lottery paradox and Buchak's counterexample.[8] However, I cannot settle the current debate on the Lockean Thesis, nor is that my primary goal. Thus, I will now explore the consequences of denying the Lockean Thesis with regard to explanationism.

4. Belief-Justifying Explanationism and the Lockean Thesis

We start by defining Belief-Justifying explanationism.

> Belief-Justifying explanationism: Explanatory inference works, in part, through certain explanatory virtues that can contribute to justification for belief without increasing rational credence.

We use explanatory inferences by noting how well (or how poorly) a hypothesis explains a set of evidence. Traditionally, the appraisal of an explanation is done using the explanatory virtues (Thagard 1978). These include virtues such as simplicity, fruitfulness, unification, mechanism, precision, informativeness, and so on. Good explanations have these virtues and bad explanations lack these virtues.

Belief-Justifying explanationism is an interpretation of how explanatory inference works, meant to be a competitor to Heuristic explanationism, Objective Bayesian explanationism, and

[8] Backes (2019) argues against closure, which would provide a Lockean-friendly solution to the lottery paradox. Freitag & Zinke (2020) argue that statistical evidence does not license high credence, and therefore Buchak's case does not threaten the Lockean thesis.

Ampliative explanationism. According to this interpretation, explanatory inference may be partially—but not wholly—situated within the Bayesian framework. That is, while some of the explanatory virtues can be analysed in terms of how they increase credence via Bayesian confirmation, others cannot. However, those virtues that cannot be analysed in terms of Bayesian priors and likelihoods may still play some positive role in justifying beliefs.

As an illustrative example, we can slightly alter the iPhone case. Suppose that men and women are equally likely to steal phones. Your iPhone is missing, and there are two possible explanations: either Jake stole it or Barbara did. You learn that an associate sent you an email containing pictures which prove that Jake was involved in some criminal activity. Thus, Jake has a motive for stealing your iPhone, whereas Barbara does not. With this new information, you perform Bayesian conditionalisation and increase your credence in [Jake stole the iPhone] to 0.91. In this case, it seems that you are rational to have high credence in [Jake stole the iPhone] and you are justified in believing that Jake stole your iPhone.[9]

Let's look more closely at the original case and the modified case. In the original case, the statistical truth provides some justification for [Jake stole the iPhone], just not enough justification (or not the right kind of justification) for justified belief. So, you are not ultimately justified in believing that Jake stole the iPhone. Compare this to the modified case. Your knowledge that Jake has a motive to steal your iPhone rightly increases your credence in [Jake stole the iPhone] to 0.91, and it justifies the belief that Jake stole the iPhone. Somehow, in this case, you attained more justification than in the original case (or the right kind of justification) such that you are now justified in believing that Jake stole the iPhone. If the analysis, so far, is

[9] This may still not be enough to convict Jake of theft, but not much hangs on this. Buchak defends similar judgments. For instance, if Jake looks guilty and looking guilty is a fairly but not perfectly reliable indicator of guilt, it seems that you are justified in believing that Jake is guilty.

accurate, then there's more to justification than increase in credence or the resulting posterior credence. When determining whether we have enough (or the right kind of) justification for belief, we need to consider factors beyond the increase in credence and the posterior credence attained from the evidence.

What other factors should be considered? The explanationist answers: explanatory considerations! Consider the original iPhone case. Absent some extra information, such as that Barbara is an incompetent thief, P(The iPhone is missing|Jake stole the iPhone) should be equal to P(The iPhone is missing|Barbara stole the iPhone). Furthermore, even though it's true (by stipulation) that men are 10 times more likely to steal iPhones than women, we have no idea what the connection is between being a man and stealing iPhones. Nor do we have any idea why Jake, in particular, might want to steal your iPhone. Thus, [Jake stole the iPhone] does not confer greater understanding of [The iPhone is missing] than [Barbara stole the iPhone].[10]

Thus, it seems that the only advantage that [Jake stole the iPhone] has over [Barbara stole the iPhone] as an explanation of [The iPhone is missing] is that P(Jake stole the iPhone) is greater than P(Barbara stole the iPhone), given the stipulation that men are more likely to steal iPhones than women. Current literature on IBE pulls in different directions regarding this issue. Some argue that higher prior probability is a good-making feature of an explanation (Okasha 2000; Good 1968; Glass 2023). Others argue that explanatory goodness only affects likelihoods (Lipton 2004; Climenhaga 2017). Fortunately, we don't need to settle this issue for the current argument. I will grant that, in the original case, if [Jake stole the iPhone] explains the missing

[10] On Lipton's (2004) account of IBE, an explanation is good to the extent that it confers understanding of the evidence. This is further supported by other authors who argue that there is a close connection between explanation and understanding (Friedman 1974; Strevens 2013).

iPhone better than [Barbara stole the iPhone], it is only in virtue of the fact that [Jake stole the iPhone] has the higher prior.

Now consider the same hypotheses in the modified case. [Jake stole the iPhone] has a higher probability than [Barbara stole the iPhone], given that Jake might have discovered that you have incriminating photos of him. After conditionalising, your credence in [Jake stole the iPhone] is 0.91. But apart from the higher prior, [Jake stole the iPhone] enjoys certain explanatory virtues that it lacks in the original case. In particular, it provides greater understanding. If the hypothesis were true, then you would understand that your iPhone is missing because Jake stole it to destroy evidence of his crime, that you would still have your iPhone if it weren't for those pictures, and so on. Thus, this explanation answers certain important what-if-things-had-been-different questions ('w-questions', for short), which is a mark of a good explanation (Woodward, 2004). In contrast, [Barbara stole your iPhone] confers less understanding. If it were true, you would understand that your iPhone is missing because Barbara stole it, but you remain in the dark regarding why Barbara stole it, and you have no clue whether it is the contents of your iPhone that led Barbara to steal it.

Bayesianism does not have the resources to distinguish between the original case and the modified case. Bayesianism only dictates what your credences should be, and in the two scenarios, your credences are comparable. However, the Belief-Justifying explanationist can say more because the two cases differ in explanatory considerations. In the original case, [Jake stole the iPhone] has the advantage of having the higher prior probability. In the modified case, [Jake stole the iPhone] has the higher probability and it provides greater understanding of the explanandum than its rival hypothesis. The lesson to learn is that different explanatory considerations can provide different amounts (or different types) of justification, and this

difference in justification is not reducible to a difference in rational credence. In the original case, [Jake stole the iPhone] may be a better explanation because it has a higher prior probability, but that is not enough to justify belief. In the modified case, [Jake stole the iPhone] has a higher probability and it confers greater understanding, thus it attains more justification than in the original case (or the right kind of justification) and it is more belief-worthy.

Like the Ampliative explanationist, the Belief-Justifying explanationist claims that explanatory inference involves a process that goes beyond Bayesian conditionalisation. However, the Ampliative explanationist states that explanatory inference consists in boosting one's credence beyond conditionalisation. In contrast, the Belief-Justifying explanationist claims that explanatory inference allows an agent to attain non-credal justification for belief.

A few qualifications should be noted here. First, I envision Belief-Justifying explanatory inference working alongside Bayesian conditionalisation. Thus, explanatory inference may not always yield outright belief. For instance, if one's credence in a hypothesis were too low after conditionalisation, the extra justification from explanatory inference may not be enough to justify outright belief. Another possibility is that sometimes explanatory virtues pull in opposite directions. Cabrera (2017) argues that, in general, the more informative a hypothesis is, the less probable it is. Similarly, Ylikoski and Kuorikoski (2010) argue that there are five explanatory virtues and that some of them are 'systematically in conflict' (p. 208). In this type of case, the explanatory hypothesis attains some *pro tanto* non-Bayesian justification, appropriate to the extent that it has or lacks the explanatory virtues. For this reason, I am not suggesting that we should never believe a hypothesis when the only advantage it has over its rivals is that it enjoys a higher prior probability.

5. The Humean Thesis—An Alternative to the Lockean Thesis

Thus far, I have suggested that some explanatory virtues can justify belief without increasing credence. I've also argued that this interpretation of explanatory inference has significant advantages over extant interpretations of explanatory inference. However, this view is incompatible with the Lockean Thesis, specifically, with the claim that rational high credence is sufficient for justified belief.

But merely denying that rational high credence is sufficient for justified belief might not be enough. We don't yet know what extra ingredient, besides high credence, is required for belief and whether explanatory virtues could provide this extra ingredient. Here, I will argue that the Humean Thesis as defended by Leitgeb (2014; 2017)—one of the main competitors to the Lockean Thesis—can lend plausibility to Belief-Justifying explanationism.[11] That is, the Humean Thesis can provide the details on how some explanatory virtues contribute to justification without increasing credence. Two things should be noted. First, there are other views, apart from the Humean Thesis, that purport to tell us what else, besides high credence, is required for belief. I cannot investigate all such views and explore how each affects Belief-Justifying explanationism, so I will restrict my discussion to the Humean Thesis. Second, as with denying the Lockean Thesis, I will not provide a defence of Leitgeb's view against objections that have been raised against it. I merely want to show that Belief-Justifying explanationism can fit quite well within the framework of the Humean Thesis.

The Humean Thesis is defined as follows.

> Humean Thesis: S is justified in believing P just in case S is justified in assigning P a stably high credence.

[11] I thank an anonymous referee for pointing out that Leitgeb's stability theory of belief would be relevant to my discussion.

On Leitgeb's view, a rational agent believes a proposition X iff they assign X a credence that is stably high under conditionalisation. More precisely, the agent's credence in X conditional on every Y that is doxastically possible to the agent is greater than some threshold. In turn, Y is doxastically possible for an agent just in case they do not believe Y's negation (2017, pp. 79–81).

In a way, Leitgeb's view is orthogonal to Belief-Justifying explanationism. After all, Leitgeb's view expresses a set of necessary and sufficient conditions for justifiably (or rationally) believing a proposition. However, it is silent on the sources of justification—whether, for instance, having testimonial evidence that X justifies an agent in assigning a stably high credence in X. So, it is possible that an agent is justified in believing X iff they assign X a stably high credence, and that one way to be justified in assigning X a stably high credence is by noting that X is a good explanation of some body of evidence.

However, we can go a step further. Belief-Justifying explanationism is the view that explanatory inference works, in part, through explanatory virtues that can increase justification for belief without increasing rational credence. If the Humean Thesis were true, then explanatory virtues can justify belief in a hypothesis H without increasing rational credence in H by bringing one's credence in H closer to stability.

There is some discussion of explanation and stability in the literature on IBE. As a response to Roche & Sober's (2013) claim that explanatoriness is evidentially irrelevant, McCain & Poston (2014) argue that explanatory hypotheses often make our probability functions more resilient.

> In general, a theory about the [causal] relationship between two (or more) events need not change the observed frequencies of those events. Consequently, once the theory is added to one's evidence one shouldn't change one's beliefs about any specific case—e.g., the chance a heavy smoker will get cancer. Yet learning the theory does change the resiliency of those probabilities. Whereas prior to possessing an explanation one only had observation to settle the frequencies, after gaining the explanation one has observation and theory (p. 147).

Thus, McCain & Poston (2014) grant Roche & Sober's (2013) claim that if P(Smith has cancer|Smith is a heavy smoker) = 0.6, learning that smoking causes cancer doesn't change that. The explanatory hypothesis does not change the likelihood that Smith has cancer, given that Smith is a heavy smoker. Instead, what [Smoking causes cancer] does is make that probability assignment more resilient. Specifically, the causal statement allows the probability assignment of 0.6 to 'survive runs of misleading data' (p. 149).

McCain & Poston (2014) do not provide a formal definition of resilience, nor do they provide a measure of resilience. Instead, they provide an example to illustrate resilience. Suppose Sally and Tom have the same credence in H, namely 0.9. Both Sally and Tom learn a new piece of evidence, E. The new piece of evidence does not affect Sally's credence, but it lowers Tom's credence to 0.5. In this case, Sally's probability function is more resilient while Tom's is more volatile (McCain & Poston 2014, p. 149).

Let's apply this to the cancer example. Suppose that P(Smith has cancer|Frequency data & Smith is a heavy smoker) = 0.6. Likewise, P(Smith has cancer|Frequency data & Smith is a heavy smoker & Smoking causes cancer) = 0.6. However, suppose now that we obtain some new (misleading) data: 100 heavy smokers who do not have lung cancer. McCain & Poston's claim is that the new data changes P(Smith has cancer|Frequency data & Smith is a heavy smoker) more significantly than it changes P(Smith has cancer|Frequency data & Smith is a heavy smoker & Smoking causes cancer). In fact, unless we have some principled reason for thinking otherwise,

the new data should be treated in the same way as the old frequency data. If P(Smith has cancer|Frequency data & Smith is a heavy smoker) was estimated to be 0.6 because there are 1000 heavy smokers in the dataset and 600 of them have cancer, then adding the new data yields 1100 heavy smokers, 600 of which has cancer. This means P(Smith has cancer|Frequency data & Smith is a heavy smoker & New data) should be close to 0.55. McCain & Poston claim, however, that P(Smith has cancer|Frequency data & Smith is a heavy smoker & Smoking causes cancer & New data) is higher than 0.55 and closer to 0.6, because the causal statement makes the probability assignment more resilient.

McCain & Poston's resilience and Leitgeb's stability are distinct.[12] For one, Leitgeb does not provide a measure of stability. On his definition, stability is binary. In contrast, McCain & Poston's example suggests that resilience comes in degrees. Furthermore, in McCain & Poston's example, an agent's credence is resilient relative to a particular proposition (Sally's credence in H is more resilient than Tom's credence in H relative to a piece of evidence E), whereas for Leitgeb, an agent's credence is stably high when it is high conditional on each doxastic possibility. However, the two concepts are closely related. An agent's credence is stably high only if their credence is sufficiently resilient conditional on each doxastically possible statement. Furthermore, even though stability is binary, we can meaningfully speak of one credence being closer to stability than another. For instance, the following is a plausible condition. Let D denote the set of all statements doxastically possible to an agent, and A be a proper subset of D, such that P(H1|E) is high for all $E \in A$ and P(H1|E) is low for all $E \in D \setminus A$. In other words, H1 is stably high conditional on each member of A, and no other doxastically possible statement. Now suppose H2 is also stably high conditional on each member of A, and on some other doxastically

[12] It should also be noted that Leitgeb sometimes uses the term 'resilience' and 'stability' interchangeably. However, I will use 'resilience' for McCain & Poston's view and 'stability' for Leitgeb's view.

possible statements. In other words, P(H2|E) is high for all E ∈ *B*, where $A \subset B \subset D$. In this case, we can say that H2 is closer to stability than H1. Note that this only specifies a sufficient condition for being closer to stability. However, it is plausible and provides a bridge from resilience to stability. Namely, if H2 and H1 are high, but H2 is more resilient than H1 relative to a proposition E, then, all else being equal, H2 is closer to stability than H1. To illustrate, if McCain & Poston are right, then since P(Smith has cancer|Frequency data & Smith is a heavy smoker & Smoking causes cancer & New data) is closer to 0.6 than P(Smith has cancer|Frequency data & Smith is a heavy smoker & New data), it follows that, all else being equal, [Smoking causes cancer] brings [Smith has cancer] closer to stability.

Finally, McCain & Poston (2014) claim that an explanatory claim makes a probability function more resilient. But they don't claim that the explanatory claim itself is resilient or more resilient than its less explanatory rivals. That is, [Smoking causes cancer] can make P(Smith has cancer|Smith is a heavy smoker) more resilient. But the result we need to support explanatory inference is that P(Smoking causes cancer) is resilient.

So, given the Humean Thesis, the task of the Belief-Justifying explanationist is to show that some of a hypothesis's explanatory virtues bring credence in that hypothesis closer to stability. One way to do this is to appeal to Quine's (1951) view that when we are forced to revise our belief system, we have a natural tendency to prefer changes that cause minimal disturbance to our total belief system. If we accept both Quine's claim that small revisions to our belief systems are preferable to large revisions and McCain & Poston's claim that explanations, generally, make our credences more resilient, we have a direct path towards a Humean Belief-Justifying explanationism. Explanatory hypotheses, when they are assigned high credences, are, all else being equal, closer to stability than their less explanatory rivals. We do not easily lower

our credences in explanatory hypotheses because they allow us to make small (rather than large) revisions to our belief systems in light of new information.

6. Classifying the Explanatory Virtues

Another way to defend the credence-stabilising effects of explanations is to investigate particular explanatory virtues. In the contemporary literature on IBE, there is a puzzle on the so-called explanatory virtues. Some explanatory virtues seem to be confirmationally relevant, others do not (Salmon 1990; Cabrera 2017). Thus, it isn't clear what role the non-confirmational virtues play in inference.

Consider unification. A unifying explanation accounts for a wide range of phenomena that, apart from the purported explanation, seem independent. The motion of the planets and the tides seem independent. However, Newton's theory of universal gravitation can account for both of these phenomena and, for this reason, it is a unifying explanation.

A probabilistic analysis can show that unification is confirmationally relevant. Let N be Newton's theory of universal gravitation, E1 be a conjunction of facts about the motion of the tides, and E2 be a conjunction of facts about the motion of the planets. N unifies E1 and E2 iff $P(E1\&E2|N) > P(E1|N)P(E2|N)$.[13] According to Bayes' Theorem, $P(N|E1\&E2) = P(N)P(E1\&E2|N)/P(E1\&E2)$. If N unifies E1&E2, then $P(E1\&E2|N) > P(E1|N)P(E2|N)$. If N does not unify E1&E2, then $P(E1\&E2|N) \leq P(E1|N)P(E2|N)$. Thus, all else being equal, the Bayes' factor, $P(E1\&E2|N)/P(E1\&E2)$ would have a greater value if N were unifying than if it weren't. For this reason, unification is a confirmational virtue.

[13] McGrew (2003) provides an extensive discussion of a probabilistic analysis of unification.

But the epistemic import of unification may go beyond probabilistic confirmation. Plausibly, a unifying explanation is the kind of hypothesis that we justifiably assign stable high credence if it is assigned a high credence at all. There are familiar examples from the history of science. Astronomers learned that Uranus's orbit was anomalous, given Newton's theory of gravitation. In Leitgeb's terms, Uranus' anomalous orbit is doxastically possible for astronomers (by virtue of being doxastically actual—i.e. believed). However, they continued to assign Newton's theory of gravitation high credence. Instead of lowering their confidence in Newtonian physics, they opted to lower their credence in the hypothesis that there are seven planets in the solar system. This eventually led to the discovery of Neptune, which accounts for Uranus' initially anomalous orbit.

In this case, then, astronomers believed Newtonian physics (i.e. assigned it stable high credence). But why? Why search for an eighth planet instead of rejecting Newtonian physics? One plausible answer is that Newtonian physics is unifying—i.e. it accounts for a wide range of phenomena. Thus, once it is assigned high credence, we prefer not to revise it to avoid having to revise beliefs about the wide range of phenomena that fall under its domain.

On the other hand, some explanatory virtues don't seem to be confirmationally relevant. Precision is one of these virtues. Schupbach & Sprenger (2011) argue that an adequate probabilistic measure of explanatory power would be a 'monotonically increasing function of the posterior ratio [P(H|E)/P(H|~E)]' (p. 111).[14] Given such a function, a precise hypothesis can have more explanatory power than its less precise rivals. Suppose E is the statement that a pole's shadow is between 5 m and 5.5 m long and the angle of incidence of light is between 42.5° and 47.5°. Under normal conditions, it follows that the height of the pole is between 4.58 m and 6

[14] I thank an anonymous referee for pointing out the relevance of (Schupbach & Sprenger 2011) to my argument here. For objections against probabilistic measures of explanatory power, see (Roche & Sober 2023).

m.[15] Now, let H1 be the hypothesis that the height of the pole is between 4 m and 6 m, and let H2 be the hypothesis that the pole is between 5 m and 6 m tall. Since P(H2|~E) is 0, the explanatory power of H2, with regard to E, is maximal.[16] On the other hand, under normal assumptions, P(H1|~E) is greater than 0. Therefore, P(H2|E)/P(H2|~E) is higher than P(H1|E)/P(H1|~E), i.e. H2 is more precise than H1 and has more explanatory power than H1. In addition, we may also note that H2 enjoys a greater degree of confirmation than H1, insofar as the Bayes factor, P(E|H2)/P(E), is greater than P(E|H1)/P(E).[17] However, since H2 entails H1, P(H2) cannot be higher than P(H1). Likewise, P(H2|E) cannot be higher than P(H1|E). Despite having greater explanatory power, H2 is less probable than H1.

In this type of case, explanatory power only increases with precision up to a certain extent. H3, the hypothesis that the height of the pole is between 4.58 m and 6 m, is less precise than H2. However, it has just as much explanatory power as H2, since P(H3|~E) is 0. Still, the lesson here is that precision may increase explanatory power, but it lowers posterior probability. For this reason, precision is classified as an informational virtue (Cabrera 2017). If we prefer H2 over H1, it is not because H2 is more probable than H1, but because it is more informative.

Thus, in the contemporary literature, there are two classes of explanatory virtues: confirmational virtues and informational virtues. However, a marriage between the Humean Thesis and Belief-Justifying explanation can give rise to a third class: stabilising virtues. As the name suggests, a stabilising virtue is a feature of a hypothesis that, all else being equal, brings that hypothesis closer to stability. The three classes of explanatory virtues may overlap. Indeed,

[15] Tan(42.5º) = 0.916; Tan(47.5º) = 1.091. Thus, the pole can be as short as 4.58 m (5/1.091 m) and as tall as 6.00 m (5.5/0.916 m).
[16] This is a corollary of Schupbach & Sprenger's (2011) measure. The explanatory power of H with regard to E is maximall iff P(H|~E) is 0 (p. 120).
[17] Since the denominators are the same, the only expressions that matter are P(E|H2) and P(E|H1). Since H2 entails E, P(E|H2) is 1. However, H1 does not entail E, therefore P(E|H1) is less than 1.

as outlined above, unification is plausibly both a confirmational and a stabilising virtue. Given the Humean Thesis, a stabilising virtue is an epistemic virtue because it makes a hypothesis more belief-worthy.

Apart from unification, there may be other stabilising virtues. More significantly, some virtues that were previously thought to be merely informational—and therefore lacking in epistemic import, under the Lockean Thesis—could turn out to be epistemic virtues given the Humean Thesis and Belief-Justifying explanationism. Mechanism is such a virtue. According to Lipton (2004), a mechanistic hypothesis provides a causal-nomological explanation of the evidence, and a mechanistic hypothesis, all else being equal, is a better explanation than a non-mechanistic hypothesis. However, a mechanistic hypothesis is often less probable than a non-mechanistic hypothesis. Again, let's return to the smoking example. Suppose H4 is the hypothesis that smoking causes cancer and H5 is a hypothesis describing the carcinogens in cigarettes and the processes by which they damage DNA in lung cells, leading to mutation and uncontrolled cell growth. H5 is a mechanistic hypothesis and H4 is not. But H5 entails H4, so H5 cannot be more probable than H4. Thus, mechanism is not a confirmational virtue.

Can mechanism, however, be a stabilising virtue? This is more challenging to demonstrate. Since H5 entails H4, but H4 does not entail H5, there are statements that would falsify H5 but not H4. Thus, relative to some statements—that Polycyclic aromatic hydrocarbons (PAHs) are not harmful, for instance—P(H4) is more resilient than P(H5).

However, if we are primarily interested in a surface-level state space concerning whether a given smoker has cancer, then H5 seems to be more stable than H4. Arguably, H5 is significantly more effective than H4 at stabilising P(Smith has cancer|Smith is a heavy smoker). H4 merely tells us that smoking causes cancer, but it does not tell us how potent smoking is as a

cause of cancer. In contrast, details about the carcinogenic substances in cigarettes and how they interact with DNA in lung cells can tell us how potent smoking is as a cause of cancer. Thus, H5 can provide an estimate for the likelihood that Smith has cancer, given that Smith is a heavy smoker, independently of the frequency data. Since H5 plays a credence-stabilising role and, according to Quine, we like to avoid making large changes to our belief system, we are justified in maintaining high credence in H5.

Thus far, I've argued that we can defend an interpretation of explanatory inference by denying the Lockean Thesis, that is, by claiming that there's more to justified belief than rational high credence. I've argued that the Humean Thesis can lend plausibility to this interpretation of explanatory inference—if justified belief consists in stably high credence, then some explanatory virtues provide justification for a hypothesis by allowing us to have a more resilient credence in that hypothesis, or a credence that is closer to stability. In the next few sections, I outline several consequences of Belief-Justifying explanationism.

7. Advantages over Other Conceptions of Explanatory Inference

Belief-Justifying explanationism has certain advantages over extant accounts of IBE. In particular, it can avoid the objections raised against the rival accounts of IBE.

First, let's consider Full-Belief IBE. Climenhaga (2017) argues that Full-Belief IBE can lead to inconsistent beliefs. Suppose I have four urns with these contents:

Urn 1: 4 white balls

Urn 2: 2 black balls, 2 white balls

Urn 3: 3 black balls, 1 white ball

Urn 4: 4 black balls.

I will flip a coin twice. If the first flip lands heads, I will flip it again to choose between Urn 1 and Urn 4. If the first flip lands tails, I will flip it again to choose between Urn 2 and Urn 3. Suppose that after the second flip, I draw a black ball from the selected urn.

In this scenario, there are two levels of explanation wherein we may perform Full-Belief IBE. On one level, Urn 4 is the best explanation for the black ball because Urn 4 has the highest ratio of black balls to white balls. On another level, Tails is the best explanation for the black ball because P(Black ball|Tails) is 0.625 while P(Black ball|Heads) is 0.5. However, Tails, in this scenario, is equivalent to Urn 2 ∨ Urn 3. Thus, if we perform Full-Belief IBE at both levels, we have the inconsistent beliefs that Urn 4 is the selected urn and that the first flip landed tails (Climenhaga 2017).

I don't accept that Urn 4 is a better explanation of the black ball than Urn 2 and Urn 3, nor that Tails is a better explanation than Heads. There are substantive accounts of scientific explanation on which making the explanandum probable is not a good-making feature of explanations, and I find these accounts plausible. But this is merely a complaint against this particular example involving urns and balls. We can conjure up similar cases where it is clear that the best explanation at one level is incompatible with the best explanation at another level. What I want to focus on, instead, is how Belief-Justifying IBE can escape Climenhaga-style counterexamples.

There are several ways to handle this kind of counterexample. The most direct way is to insist that meeting some minimal threshold above 0.5 is a necessary condition for justified belief. Some authors endorse this view (Chandler 2010; Pettigrew 2015; Worsnip 2015). Stipulating this as a necessary condition protects the Belief-Justifying explanationist from inconsistent beliefs, so long as they have coherent credences. In the urn example, P(Urn 4|black ball) is 0.444. On the

other hand, P(Tails|black ball) is 0.889. Therefore, at most, the Belief-Justifying explanationist considers the agent justified in believing that the first flip landed tails, but not that Urn 4 was selected.

However, this solution is unattractive to those who sympathise with the view defended by Hawthorne, Rothschild, & Spectre (2016) that having justified belief is compatible with having credence lower than 0.5.[18] In this case, one may look for other methods to handle Climenhaga's objection. For instance, Hawthorne, Rothschild, & Spectre discuss—but do not endorse—the view that belief is partition-relative, so that an agent may justifiably believe that Urn 4 was selected (given the Urn-level partition) and believe that the coin landed tails (given the coin-level partition) (p. 1400).

However, Hawthorne, Rothschild, & Spectre's paper opens up other interesting possibilities for the explanationist. The primary thesis of Hawthorne, Rothschild, & Spectre's paper is that the evidential norms for belief are weaker than the evidential norms for assertion. Thus, while justified belief does not require having credence over 0.5, justified assertion may require having credence over 0.5. Thus, in place of Belief-Justifying explanationism, an explanationist may endorse Assertion-Justifying explanationism. Like the name suggests, this would be the view that the explanatory goodness of a hypothesis can provide justification for asserting that hypothesis without increasing an agent's credence in it. Assertion-Justifying explanationism enjoys similar advantages to Belief-Justifying explanationism, such as being compatible with Bayesianism, granting explanatory inference an important role in our epistemic lives, and being compatible with different models of scientific explanation. For the rest of this

[18] I thank an anonymous reviewer for pointing out the relevance of Hawthorne, Rothschild, and Spectre's (2016) paper to my argument.

paper, however, I will focus my discussion on Belief-Justifying explanationism and leave discussions on Assertion-Justifying explanationism for another time.

Next, consider the apparent dilemma faced by the Heuristic explanationist and the Ampliative explanationist. On the one hand, Heuristic explanationism seems to make explanatory inference dispensable and philosophically uninteresting. On the other hand, Ampliative explanationism is charged with being inconsistent with Bayesianism. Belief-Justifying explanationism avoids both of these problems. It allows us to say that explanatory inference can play an important role in belief formation without denying Bayesianism.

Finally, consider Objective-Bayesian explanationism. There are two important objections raised against it in the current literature. First, the Objective Bayesian explanationist must show that a hypothesis's status as the best explanation of a body of evidence provides normative constraints for setting high priors and likelihoods associated with that hypothesis. More precisely, an agent who recognises that H is the best explanation of E ought to assign P(H) and P(E|H) high values. However, as discussed above, some of the explanatory virtues used in positive appraisal of a potential explanation are not confirmational virtues. That is, though they are good-making features of explanation, they are not probable-making features of a hypothesis (Cabrera 2017).

The second objection concerns substantive accounts of scientific explanation. As I have argued elsewhere (2023, 2024), explanationists should engage some literature on models of scientific explanation when defending explanatory inference. More specifically, I've argued that several influential accounts of scientific explanation are incompatible with Objective Bayesian IBE (2024). For instance, Railton's (1978) Deductive-Nomological-Probabilistic (DNP) model and Salmon's (1971) Statistical-Relevance (SR) model fall in this category. The primary reason

for incompatibility is that these are accounts where making the explanandum probable is not a good-making feature of an explanation. Therefore, H can be a perfectly good SR explanation for E, while P(E|H) is low.

Belief-Justifying explanationism avoids both of these objections. As discussed above, the Belief-Justifying explanationist maintains that some virtues can provide justification for belief in the best explanation without increasing one's credence in the best explanation, by stabilising credences, for instance. Likewise, the Belief-Justifying explanationist can maintain that a hypothesis's status as a good SR or DNP explanation can provide justification for belief in that hypothesis without increasing one's credence in the SR or DNP explanation. Of course, the philosophers who defend accounts of scientific explanation typically do not claim that we attain justification for belief in a hypothesis when it meets their requirements for being a scientific explanation. But neither do they deny this. Thus, Belief-Justifying explanationism is not entailed by existing accounts of scientific explanation, but it is compatible with them.

8. Belief-Credence Dualism

The Lockean Thesis, as defined above, is neutral on the metaphysical relationship between belief and credence. It is compatible with dualism, the view that neither beliefs nor credences reduces to the other.[19] It is also compatible with reductionist views—that beliefs reduce to credences or that credences reduce to beliefs.[20]

Belief-Justifying explanationism starts by denying the Lockean Thesis. Having rational credence over some threshold is not sufficient for justified belief. Does Belief-Justifying

[19] For defences of dualism, see (Hawthorne 2009; Jackson 2019)
[20] Pettigrew (2015, 2016) defends the view that belief reduces to credence. Weisberg (2013, section 3.2) defends the view that credence reduces to belief, but defends a dualist view in (2020).

explanationism also require taking a stand on the metaphysical relation between beliefs and credences? At least, it seems that Belief-Justifying explanationism requires denying that justified belief reduces to rational high credence. Part of the motivation for endorsing Belief-Justifying explanationism is to give explanatory considerations an important epistemic role. But if justified belief reduces to rational high credence, then there is no significant difference between justified belief and rational high credence. If so, it is hard to see how we have truly given explanatory considerations an important role in our epistemic lives.

For this reason, the Belief-Justifying explanationist should, at least, deny that justified belief reduces to rational high credence. Epistemologists who defend this position usually argue that belief can play certain roles that high credence cannot. Several examples are worth citing. Buchak (2014) argues that belief is not reducible to high credence because our norms of blaming someone involve belief, but not high credence. That is, having high credence that S did something morally wrong is not sufficient to warrant blame, but having the belief that S did something morally wrong is sufficient to warrant blame. Ross & Schroeder (2014) argue that belief is required for our epistemic attitudes to be correct or incorrect. If I assign some proposition P a high credence less than 1 (say, 0.95) and P turns out to be false, you cannot say that I was wrong. After all, I assigned not-P some positive credence, too, and it so happens that P is false. By parity of reasoning, I also cannot be right about a proposition simply by assigning that proposition a high credence. Finally, Jackson (2019) defends dualism, in part, by arguing that belief simplifies our reasoning.

9. Belief-Justifying Explanationism and Evidentialism

Evidentialism is the view that whether an agent is justified in believing a proposition is entirely determined by the agent's evidence. Though it's a popular view in epistemology, it has recently been challenged in the literature. One view that conflicts with evidentialism is pragmatic encroachment (Ganson 2008; Fantl and McGrath 2012). Defenders of pragmatic encroachment argue that, apart from evidence, stakes also play a role in determining whether an agent is justified in believing a proposition. Suppose I made an almond butter sandwich and a peanut butter sandwich, and I put them both in my fridge. In normal circumstances, I am justified in believing, on the basis of memory, that the sandwich on the left is the peanut butter sandwich. However, if I have a hungry guest who is fatally allergic to peanuts, I am no longer justified in believing this based on memory. Here, we can see that a change in stakes, unaccompanied by a change in evidence, results in a change in the justification of beliefs.

Does Belief-Justifying explanationism likewise require denying evidentialism? Consider the original iPhone case and the modified iPhone case. In both cases, your final credence in [Jake stole the iPhone] is 0.91, but you do not have the same evidence in the two cases. Your knowledge of possible motives constitutes evidence, but that is only available in the modified case, not in the original case. In short, while your total evidence in the two cases yields the same probability for [Jake stole the iPhone], you don't have the same evidence in the two cases. Thus, these two cases don't tell us whether Belief-Justifying explanationism is compatible with evidentialism.

We can investigate the issue from a different angle. A piece of evidence E and a hypothesis H can have different types of features and relations. E can confirm, weakly confirm, or disconfirm H, E can have high or low initial probability—these are instances of probabilistic features and relations. But there are also non-probabilistic features and relations. E can be

explained by H, for instance. The evidentialist claims that whether an agent is justified in believing H is entirely determined by the agent's evidence. But in claiming this, the evidentialist is not advocating the view that the probabilistic features and relations are not relevant. Whether and to what extent E confirms H is part of the mechanism by which E justifies (or fails to justify) the belief that H. In determining how well E justifies H, the evidentialist accounts for the probabilistic features and relations of E and H.

Stakes, on the other hand, seem external to the mechanisms of how E justifies the belief that H. Indeed, many who defend pragmatic encroachment do not argue that low stakes contribute to the justification for the belief that there is an almond butter sandwich in the fridge. Rather, the stakes lower the standards for justified belief by, for instance, lowering the Lockean threshold at which credence turns into belief (Bach 2008; Ganson 2008). This is why pragmatic encroachment is considered incompatible with evidentialism.

What about the explanatory virtues? Are they, like probabilistic features and relations, internal to the mechanisms of how E justifies H? Or do they, like stakes, lead to justified belief in some roundabout way, such as lowering the standards required for justified belief? My preference is to treat explanatory virtues like probabilistic features and relations rather than stakes. They contribute to justification rather than affect the standards required for justified belief. However, a full answer may require defending a substantive account of explanation, which is beyond the scope of this paper. On pragmatic accounts of explanation, a mere change in context—the interests of the agents involved in giving and receiving explanations—can change the status of a hypothesis from being explanatory to non-explanatory. Van Fraassen (1980) and Achinstein (1983; 1984) are proponents of such views. On non-pragmatic accounts, a hypothesis's status as explanatory depends on the content of the hypothesis or its logical relation

to the explanandum—whether it reports laws of nature (Hempel 1965), whether it reports causal facts (Salmon 1984; Woodward 2004), how well it unifies evidence (Kitcher 1981), and so on.

Suppose Hempel's Deductive-Nomological model of explanation is correct. In this case, I'm inclined to think that explanatory features and relations work in the same way that probabilistic features and relations do. They are internal to the mechanisms of how E justifies H. Thus, in determining how well E justifies H, the evidentialist can account for the probabilistic features of E and H, and the explanatory virtues of H. On the other hand, if a pragmatic account of explanation is correct, then this may lead to a Threshold-lowering explanatory inference, where explanatory virtues play the same role that stakes do for pragmatic encroachment. A change in the interests of agents may lead to a change in which hypotheses are explanatory, which in turn affects the standards for justified belief.

10. Conclusion

Much of current literature on IBE focuses on the question of whether IBE is compatible with Bayesianism. Some philosophers attempt to develop some account of IBE that is compatible with Bayesianism without trivialising it. With Belief-Justifying explanationism, we have a way out of the dilemma. Instead of denying Bayesianism, the explanationist can deny the Lockean Thesis and claim that explanatory considerations help an agent attain justification for a belief without needing to boost their credence beyond conditionalisation.

Acknowledgements

I thank the participants of the 2024 *Probabilistic Reasoning in the Sciences Conference* held at Marche Polytechnic University for helpful comments and discussions on an earlier draft of my paper. I also thank three anonymous reviewers for their comments on this paper.

References

Achinstein, Peter. 1983. *The Nature of Explanation*. Oxford University Press.

Achinstein, Peter. 1984. 'The Pragmatic Character of Explanation'. *Proceedings of the Biennial Meeting of the Philosophy of Science Association* 2: 275–92. https://doi.org/10.1086/psaprocbienmeetp.1984.2.192509

Bach, Kent. 2008. 'Applying Pragmatics to Epistemology'. *Philosophical Issues* 18 (1): 68–88. https://doi.org/10.1111/j.1533-6077.2008.00138.x.

Backes, Marvin. 2019. 'A Bitter Pill for Closure'. *Synthese* 196: 3773–87. https://doi.org/10.1007/s11229-017-1620-8.

Buchak, Lara. 2014. 'Belief, Credence, and Norms'. *Philosophical Studies* 169 (2): 285–311. https://doi.org/10.1007/s11098-013-0182-y.

Cabrera, Frank. 2017. 'Can There Be a Bayesian Explanationism? On the Prospects of a Productive Partnership'. *Synthese* 194 (4): 1245–72. https://doi.org/10.1007/s11229-015-0990-z.

Chandler, Jake. 2010. 'The Lottery Paradox Generalised?' *The British Journal for the Philosophy of Science* 61 (3): 667–79. https://doi.org/10.1093/bjps/axp056.

Clarke, Roger. 2013. 'Belief Is Credence One (in Context)'. *Philosopher's Imprint* 13: 1–18.

Climenhaga, Nevin. 2017. 'Inference to the Best Explanation Made Incoherent'. *The Journal of Philosophy* 114 (5): 251–73. https://doi.org/10.5840/jphil2017114519

Dellsén, Finnur. 2017. 'Abductively Robust Inference'. *Analysis* 77 (1): 20–29. https://doi.org/10.1093/analys/anx049

Dellsén, Finnur. 2018. 'The Heuristic Conception of Inference to the Best Explanation'. *Philosophical Studies* 175 (7): 1745–66. https://doi.org/10.1007/s11098-017-0933-2.

Dodd, Dylan. 2017. 'Belief and Certainty'. *Synthese* 194 (11): 4597–621. https://doi.org/10.1007/s11229-016-1163-4.

Douven, Igor. 2002. 'A New Solution to the Paradox of Rational Acceptability'. *The British Journal for the Philosophy of Science* 53 (3): 391–410. https://doi.org/10.1093/bjps/53.3.391.

Douven, Igor. 2012. 'The Lottery Paradox and the Pragmatics of Belief'. *Dialectica* 66 (3): 351–73. https://doi.org/10.1111/j.1746-8361.2012.01304.x.

Douven, Igor. 2013. 'Inference to the Best Explanation, Dutch Books, and Inaccuracy Minimisation'. *The Philosophical Quarterly* 63 (252): 428–44. https://doi.org/10.1111/1467-9213.12032.

Douven, Igor, and Jonah N. Schupbach. 2015. 'Probabilistic Alternatives to Bayesianism: The Case of Explanationism'. *Frontiers in Psychology* 6 (459). https://doi.org/10.3389/fpsyg.2015.00459.

Enoch, David, Levi Spectre, and Talia Fisher. 2012. 'Statistical Evidence, Sensitivity, and the Legal Value of Knowledge'. *Philosophy & Public Affairs* 40 (3): 197–224. https://doi.org/10.1111/papa.12000.

Fantl, Jeremy, and Matthew McGrath. 2010. *Knowledge in an Uncertain World*. Oxford University Press.

Fantl, Jeremy, and Matthew McGrath. 2012. 'Pragmatic Encroachment: It's Not Just about Knowledge'. *Episteme* 9 (1): 27–42. https://doi.org/10.1017/epi.2011.3.

Fraassen, Bas van. 1980. *The Scientific Image*. Oxford University Press.

Freitag, Wolfgang, and Alexandra Zinke. 2020. ‘Statistics and Suspension’. *Philosophical Studies* 177: 2877–80. https://doi.org/10.1007/s11098-019-01344-7.
Friedman, Michael. 1974. ‘Explanation and Scientific Understanding’. *The Journal of Philosophy* 71 (1): 5–19. https://doi.org/10.2307/2024924.
Ganson, Dorit. 2008. ‘Evidentialism and Pragmatic Constraints on Outright Belief’. *Philosophical Studies* 139 (3): 441–58. https://doi.org/10.1007/s11098-007-9133-9
Gardiner, Georgi. 2019. ‘The Reasonable and the Relevant: Legal Standards of Proof’. *Philosophy & Public Affairs* 47 (3): 288–318. https://doi.org/10.1111/papa.12149
Glass, David. 2023. ‘How Good Is an Explanation?’ *Synthese* 201: 53. https://doi.org/10.1007/s11229-022-04025-x.
Good, Irving John. 1968. ‘Corroboration, Explanation, Evolving Probability, Simplicity, and a Sharpened Razor’. *The British Journal for the Philosophy of Science* 19: 123–43. https://doi.org/10.1093/bjps/19.2.123
Greco, Daniel. 2015. ‘How I Learned to Stop Worrying and Love Probability 1’. *Philosophical Perspectives* 29 (1): 179–201. https://doi.org/10.1111/phpe.12059.
Harman, Gilbert. 1965. ‘The Inference to the Best Explanation’. *The Philosophical Review* 74 (1): 88–95. https://doi.org/10.2307/2183532.
Hawthorne, John. 2009. ‘The Lockean Thesis and the Logic of Belief’. In *Degrees of Belief*, edited by Franz Huber and Christoph Schmidt-Petri. Springer.
Hawthorne, John, Daniel Rothschild, and Levi Spectre. 2016. ‘Belief Is Weak’. *Philosophical Studies* 173: 1393–404. https://doi.org/10.1007/s11098-015-0553-7.
Hempel, Carl G. 1965. ‘Aspects of Scientific Explanation’. In *Aspects of Scientific Explanation and Other Essays*. The Free Press.
Huemer, Michael. 2009. ‘Explanationist Aid for the Theory of Inductive Logic’. *The British Journal for the Philosophy of Science* 60 (2): 345–75. https://doi.org/10.1093/bjps/axp008
Jackson, Elizabeth. 2019. ‘How Belief-Credence Dualism Explains Away Pragmatic Encroachment’. *The Philosophical Quarterly* 69 (276): 511–33. https://doi.org/10.1093/pq/pqz006.
Jackson, Elizabeth. 2020a. ‘Belief, Credence, and Evidence’. *Synthese* 197: 5073–92. https://doi.org/10.1007/s11229-018-01965-1
Jackson, Elizabeth. 2020b. ‘The Relationship Between Belief and Credence’. *Philosophy Compass* 15 (6): e12668. https://doi.org/10.1111/phc3.12668.
Kitcher, Philip. 1981. ‘Explanatory Unification’. *Philosophy of Science* 48 (4): 507–31. https://doi.org/10.1086/289019.
Lee, Matthew. 2017. ‘On the Arbitrariness Objection to the Threshold View’. *Dialogue* 56 (1): 143–58. https://doi.org/10.1017/s0012217317000154
Leitgeb, Hannes. 2014a. ‘The Review Paradox: On the Diachronic Costs of Not Closing Rational Belief under Conjunction’. *Noûs* 48 (4): 781–93. https://doi.org/10.1111/nous.12020
Leitgeb, Hannes. 2014b. ‘The Stability Theory of Belief’. *Philosophical Review* 123 (2): 131–71. https://doi.org/10.1215/00318108-2400575
Leitgeb, Hannes. 2017. *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford University Press.
Lipton, Peter. 2004. *Inference to the Best Explanation*. 2nd edn. Routledge.
McCain, Kevin. 2015. ‘Explanationism: Defended on All Sides’. *Logos and Episteme* 6 (3): 333–49. https://doi.org/10.5840/logos-episteme20156325.

McCain, Kevin, and Ted Poston. 2014. ‘Why Explanatoriness Is Evidentially Relevant’. *Thought* 3: 145–53. https://doi.org/10.1002/tht3.122.
McGrew, Timothy. 2003. ‘Confirmation, Heuristics, and Explanatory Reasoning’. *The British Journal for the Philosophy of Science* 54 (3): 553–67. https://doi.org/10.1093/bjps/54.4.553.
Okasha, Samir. 2000. ‘Van Fraassen’s Critique of Inference to the Best Explanation’. *Studies in History and Philosophy of Science* 31 (4): 691–710. https://doi.org/10.1016/S0039-3681(00)00016-9.
Pace, Michael. 2011. ‘The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James’ "Will to Believe"’. *Noûs* 45 (2): 239–68. https://doi.org/10.1111/j.1468-0068.2010.00768.x.
Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. 2nd edn. Cambridge University Press.
Pettigrew, Richard. 2015. ‘Accuracy and the Credence-Belief Connection’. *Philosopher’s Imprint* 15 (16): 1–20.
Pettigrew, Richard. 2016. *Accuracy and the Laws of Credence*. Oxford University Press.
Poston, Ted. 2014. *Reason and Explanation: A Defense of Explanatory Coherentism*. Palgrave Macmillan.
Prasetya, Yunus. 2023. ‘Towards a Synthesis of Two Research Programmes: Inference to the Best Explanation and Models of Scientific Explanation’ *Australasian Journal of Philosophy* 101 (3): 750-64. https://doi.org/10.1080/00048402.2022.2041058.
Prasetya, Yunus. 2024. ‘Which Models of Scientific Explanation Are (In)Compatible with Inference to the Best Explanation?’ *The British Journal for the Philosophy of Science* 75 (1): 209–32. https://doi.org/10.1086/715203.
Psillos, Stathis. 2007. ‘The Fine Structure of Inference to the Best Explanation’. *Philosophy and Phenomenological Research* 74 (2): 441–48. https://doi.org/10.1111/j.1933-1592.2007.00030.x.
Quine, Willard Van Orman. 1951. ‘Two Dogmas of Empiricism’. *The Philosophical Review* 60: 20–43. https://doi.org/10.2307/2266637.
Railton, Peter. 1978. ‘A Deductive-Nomological Model of Probabilistic Explanation’. *Philosophy of Science* 45: 206–26. https://doi.org/10.1086/288797.
Roche, William, and Elliott Sober. 2013. ‘Explanatoriness Is Evidentially Irrelevant, or Inference to the Best Explanation Meets Bayesian Confirmation Theory’. *Analysis* 73 (4): 659–68. https://doi.org/10.1093/analys/ant079.
Roche, William, and Elliott Sober. 2023. ‘Purely Probabilistic Measures of Explanatory Power–A Critique’. *Philosophy of Science* 90 (1): 129–49. https://doi.org/10.1017/psa.2021.39.
Ross, Jacob, and Mark Schroeder. 2014. ‘Belief, Credence, and Pragmatic Encroachment’. *Philosophy and Phenomenological Research* 88 (2): 259–88. https://doi.org/10.1111/j.1933-1592.2011.00552.x
Salmon, Wesley. 1990. ‘Rationality and Objectivity in Science or Tom Kuhn Meets Tom Bayes’. In *Scientific Theories*, edited by C. Wade Savage. University of Minnesota Press.
Salmon, Wesley. 1971. ‘Statistical Explanation’. In *Statistical Explanation and Statistical Relevance*, by Wesley C. Salmon, Richard C. Jeffrey, and James G. Greeno. University of Pittsburgh Press.
Salmon, Wesley. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton University Press.

Schupbach, Jonah N., and Jan Sprenger. 2011. ‘The Logic of Explanatory Power’. *Philosophy of Science* 78 (1): 105–27. https://doi.org/10.1086/658111.
Shear, Ted, and Branden Fitelson. 2019. ‘Two Approaches to Belief Revision’. *Erkenntnis* 84 (3): 487–518. https://doi.org/10.1007/s10670-017-9968-1.
Strevens, Michael. 2013. ‘No Understanding Without Explanation’. *Studies in History and Philosophy of Science* 44 (3): 510–15. https://doi.org/10.1016/j.shpsa.2012.12.005.
Thagard, Paul. 1978. ‘The Best Explanation: Criteria for Theory Choice’. *The Journal of Philosophy* 75 (2): 76–92. https://doi.org/10.2307/2025686.
Trpin, Borut. 2024. ‘Affirming the Explanandum’. *Analysis* 84 (4): 788–96. https://doi.org/10.1093/analys/anae003.
Vogel, Jonathan. 1998. ‘Inference to the Best Explanation’. In *Routledge Encyclopedia of Philosophy*, edited by Edward Craig. Routledge.
Weatherson, Brian. 2005. ‘Can We Do Without Pragmatic Encroachment?’ *Philosophical Perspectives* 19 (1): 417–43. https://doi.org/10.1111/j.1520-8583.2005.00068.x.
Weisberg, Jonathan. 2009. ‘Locating IBE in the Bayesian Framework’. *Synthese* 167: 125–43. https://doi.org/10.1007/s11229-008-9305-y.
Weisberg, Jonathan. 2013. ‘Knowledge in Action’. *Philosopher’s Imprint* 13 (22): 1–23.
Weisberg, Jonathan. 2020. ‘Belief in Psyontology’. *Philosopher’s Imprint* 20 (11): 1–27.
Woodward, James. 2004. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.
Woodward, James. 2025. ‘The Place of Explanation in Scientific Inquiry: Inference to the Best Explanation vs Inference to the Only Explanation’. *European Journal for Philosophy of Science* 15 (21). https://doi.org/10.1007/s13194-025-00649-1.
Worsnip, Alex. 2015. ‘Belief, Credence, and the Preface Paradox’. *Australasian Journal of Philosophy* 94 (3): 549–62. https://doi.org/10.1080/00048402.2015.1084343.
Ylikoski, Petri, and Jaakko Kuorikoski. 2010. ‘Dissecting Explanatory Power’. *Philosophical Studies* 148: 201–19. https://doi.org/10.1007/s11098-008-9324-z.