Inferring to the Best Explanation from Uncertain Evidence

Finnur Dellsén

Accepted for presentation at PSA2024

Abstract

This paper presents a new problem for the inference rule commonly known as *Inference to the Best Explanation* (IBE). The problem is that uncertainty about parts of one's evidence may undermine the inferrability of a hypothesis that would provide the best explanation of that evidence, especially in cases where there is an alternative hypothesis that would provide a better explanation of only the more certain pieces of evidence. A potential solution to the problem is sketched, in which IBE is generalized to handle uncertain evidence by invoking a notion of *evidential robustness*.

1 Introduction

According to standard accounts of *Inference to the Best Explanation* (IBE), one may infer a hypothesis just in case it would provide a better explanation of one's total evidence than any other available explanatory hypothesis (e.g., Harman, 1965; Thagard, 1978; Lipton, 2004). Some influential accounts of IBE conceive of it as a basic and autonomous form of non-deductive inference warranting belief in, or acceptance of, its conclusions (Lycan, 1985; Harman, 1989; McCain and Moretti, 2022). However, it is now more common for proponents of IBE to argue that it serves to complement some more sophisticated form of ideal reasoning, e.g. Bayesian reasoning, by serving as a useful heuristic for mathematically limited beings who cannot be expected to assign precise credences to all relevant propositions, to do so in a probabilistically coherent manner, or to update these credences strictly by conditionalizing on newly obtained evidence (Okasha, 2000; Lipton, 2001; Dellsén, 2018). On this latter view, IBE is still a rule of non-deductive inference, warranting some form of belief or acceptance in its conclusions, but the point of this is ultimately to approximate the ideal form of (typically Bayesian) reasoning.

Although these accounts thus disagree on whether IBE is a epistemically fundamental or instrumental, they share a core commitment that may roughly be summarized as follows:

> **$\mathbf{IBE_{ST}}$:** If H would provide a better explanation of one's total evidence $\mathbb{E}$ than any other available hypothesis, then one may defeasibly infer H.

I will refer to views that conform to IBE_{ST} as *standard accounts of IBE*. Not quite everything that has been referred to as 'IBE' fits this description, but much of it does.[1] Of course, much will depend on how proponents of such accounts spell out key notions such as 'better explanation', 'available hypothesis' and 'defeasibly infer'. But let us set such details aside. This paper presents a new problem for IBE that applies regardless of how it is fleshed out in these respects.

In short, the problem stems from the fact that some or all of one's total evidence $\mathbb{E}$ may be *uncertain* to various degrees. Since this problem is analogous to a well-known problem for Bayesianism, I begin by briefly discussing the latter, before then turning to presenting the new problem of uncertain evidence that partains to IBE (§2). I then outline a possible solution to the problem, in which the standard account of IBE is generalized

[1] In particular, 'IBE' is also sometimes used to describe accounts of non-deductive reasoning that are really just variations on, or extensions of, standard forms of Bayesianism. For example, Weisberg (2009b) and Huemer (2009) supplement the Bayesian framework with the requirement that higher prior probabilities should be assigned to more explanatory hypotheses *a priori*, while van Fraassen (1989) and Douven (2022) have explored an updating rule in which bonus probability points are awarded to best explaining hypotheses. However, as van Fraassen (1989, 145) himself notes, classifying these accounts as forms of IBE would make 'Inference to the Best Explanation' a misnomer, since these accounts do not really involve any *inference*. In any case, my discussion below largely sidesteps these IBE-inspired variations on Bayesian reasoning because, as we shall see (§2), Bayesianism already faces a more familiar problem of the sort I will be discussing.

so as to handle cases of uncertain evidence in a plausible way (§3). Finally, I briefly illustrate how the account may be applied to scientific inferences by briefly considering how Einstein handled evidential uncertainty about Dayton Miller's experiments on 'ether drift' in the 1920s (§4).

2 IBE's Problem of Uncertain Evidence

Let us start with a familiar problem of uncertain evidence. *Bayesian Conditionalization* (BC) says that when you learn some evidence E, your new subjective probability in any hypothesis H, $P_n(H)$, should equal your previous conditional probability in H *given* E, i.e. $P_o(H|E)$. Here, 'learning' E means becoming absolutely certain that E is true, i.e. setting $P_n(E) = 1$.[2] This implication of BC is generally considered implausible in and of itself, for surely we should distinguish between many evidential propositions and patently self-evident truths such as $1 + 1 = 2$. Moreover, requiring certainty of any empirical proposition is highly problematic within the Bayesian framework, where such extreme probabilities cannot ever be revised through later applications of BC.

A common response to this Bayesian problem of uncertain evidence, influentially suggested by Jeffrey (1965), is to replace BC with a more general updating rule. *Jeffrey Conditionalization* (JC) says that when your subjective probabilities in some partition of evidential propositions $E_1, ..., E_n$ changes in any way, your new probability for H should be a weighted sum of your previous conditional probabilities of H given each E_i, where each weight is your new probability in E_i, i.e. $P_n(H) = \sum_{i \leq n} P_n(E_i) P_o(H|E_i)$. JC nicely allows for each evidential proposition in the partition to be uncertain, since we can have $P_n(E_i) < 1$ for all i; while still giving us BC as a special limiting case in which some evidential proposition E_i in the partition is assigned probability 1. Thus, while JC may itself not be entirely unproblematic (see, e.g., Weisberg, 2009a; Trpin, 2020), it is safe to say that the possibility of obtaining uncertain evidence has motivated a profound revision of the Bayesian

[2]To see why this must be so, note that in the Bayesian framework E is just another hypothesis prior to updating, and since $P_o(E|E) = 1$, it follows from BC that $P_n(E) = 1$.

framework.

By contrast, standard accounts of IBE have not yet accommodated this possibility of obtaining uncertain evidence. The only extant discussion of uncertain evidence with reference to 'IBE' appears in Trpin and Pellert's (2019) account of how JC may be combined with the idea, due originally to van Fraassen (1989) and later developed by Douven (2022), that the hypotheses which provide better explanations of some evidence should be given bonus probability points as one updates on that evidence.[3] Trpin and Pellert show how this can be done in an elegant way by generalizing Douven's IBE-inspired updating rule, EXPL, which is itself a generalization of BC, and argue that this new probabilistic updating rule is preferable to JC in a number of simulations. What the literature does not contain, however, is any hint at how uncertain evidence may be handled in standard accounts of IBE, i.e. those according to which IBE is a form of defeasible non-deductive inference (rather than a mere variation on an probabilistic updating rule such as BC).

This is an important lacuna in the literature, for—as I'll now argue—such standard accounts of IBE face a structurally similar problem concerning uncertain evidence, which in turn calls for a correspondingly profound revision of these accounts. As before, the problem stems from the fact that some or all of one's total evidence $\mathbb{E}$ may be uncertain for various reasons. In many cases, such evidential uncertainty undermines the inferrability of a given hypothesis H that provides the best explanation of $\mathbb{E}$, for H may only provide a better explanation than an alternative H' on the assumption that this uncertain evidence is indeed accurate. However, as we shall see, standard accounts of IBE as involving an inference to H from $\mathbb{E}$ are not even capable of representing such uncertainty about the evidence—let alone its implication for the inferrability of a given hypothesis.

For example, consider an everyday case discussed by van Fraassen (1980, 19-20): "I hear scratching in the wall, the patter of little feet at midnight, my cheese disappears—and I infer that a mouse has come to live with me."

---

[3] See also footnote 1.

Setting aside van Fraassen's (and others') concerns about the cogency of IBE, it seems *prima facie* reasonable to make the inference from (E_s) *s*cratching in the wall, (E_p) *p*attering of little feet, and (E_d) *d*isappearance of cheese, to the hypothesis (H_m) that a *m*ouse has come to live with me. After all, H_m arguably provides the best explanation of E_s, E_p, and E_d taken together. So far so good. But what if you are not so sure about E_d, i.e. that the cheese actually disappeared, e.g. because you don't quite remember whether you ate the cheese yourself the night before? Surely, such an uncertainty about the evidence should be taken into account in considering whether, or the extent to which, H_m can be inferred from E_s, E_p, and E_d in the situation in question.

This is especially apparent if there is another available hypothesis that might provide an equally good, or indeed better, explanation of the remaining evidence, i.e. the evidence of which you are certain. For example, it is not hard to imagine a case in which the scratching and the pattering, i.e. E_s and E_p, would be best explained by (H_s) the presence of a *s*quirrel outside your house (rather than a mouse inside it). In that case, since you are uncertain about the crucial piece of evidence E_d that discriminates between this alternative hypothesis H_s and the original mouse hypothesis H_m, it seems at best dubious to infer H_m—even though there is a clear sense in which it provides the best explanation of the entirety of your evidence (including evidence about which you are uncertain). Indeed, it would seem to be equally dubious to infer the alternative hypothesis H_s in such a situation, even though it provides the best explanation of the evidence of which you

[4] One might think that in situations of this sort the obvious response is to infer the disjunction of the two hypotheses that would provide the best explanations in each case, i.e. $H_m \vee H_s$. This is indeed a plausible inference, and below I'll outline a modification to IBE which validates this thought. But can standard accounts of IBE—on which one is supposed to infer the best explanation of one's evidence, as per IBE_{ST}—deliver this intuitively correct result? It is hard to see how, for two reasons. First, disjunctions of explanatory hypotheses are often notoriously bad explanations—if indeed they count as explanations at all (see Weslake, 2013, and references therein). Thus, it is at best unclear whether $H_m \vee H_s$ would come out at as providing the best (or indeed any) explanation of the evidence at hand. Second, it is also not clear what evidence $H_m \vee H_s$ would be providing the best explanation of. After all, that evidence cannot simply be E_s, E_p, *and* E_d, for that combination of evidence is best explained by H_m. Nor can it simply be E_s and E_p (excluding E_d), for that combination is best explained by H_s.

are certain. After all, you are not certain that E_d is false either—rather, you are uncertain whether E_d is true or false—and if E_d *were* true then H_s would arguably *not* be the best explanation of your evidence.[4]

In sum, then, standard accounts of IBE, on which one may infer a hypothesis H just in case it would provide the best explanation of one's total evidence $\mathbb{E}$, are simply unsuited to handle cases of uncertain evidence. The underlying reason for this should be clear at this point. On these accounts, a given piece of potential evidence E_i is either included in, or excluded from, the set of total evidence $\mathbb{E}$ of which a given hypothesis H must provide the best explanation if H is to be inferrable by IBE. If E_i is uncertain, it occupies a sort of in-between space, neither clearly included nor clearly excluded from $\mathbb{E}$. It is unclear at present how this in-between space could or should be conceptualized in accounts of IBE; nor is it clear what sort of claims IBE should allow one to infer from evidence that partly occupies such an in-between space.

One possible response to this problem of uncertain evidence is to jettison standard accounts of IBE altogether, e.g. by replacing them with a purely probabilistic framework in which hypotheses are never really inferred at all. Before we take such drastic measures, however, let us consider whether it may be possible to modify—or, rather, generalize—standard accounts of IBE so as to handle the problem in a more accommodating manner.

3 Evidentially Robust IBE

The key idea behind the solution I wish to explore is that, in cases of evidential uncertainty, it may still be that the same hypothesis provides the best explanation regardless of whether uncertain pieces of evidence are included in the set of evidence to be explained. To put the point differently, it may be that a given hypothesis's status as the best explanation is *robust* across all the different possibilities that are left open by one's uncertainty about the evidence; in that case, the evidential uncertainty should not matter to the inferrability of H.

To flesh out this idea, let us model an agent's total evidence at a given time as a set of evidential propositions, $\mathbb{E} = \{E_1, ..., E_m\}$. Given some some standard for what makes an evidential proposition *certain* (or, equivalently, *uncertain*), there is a subset $\mathbb{E}_c \subseteq \mathbb{E}$ which contains only evidential propositions of which one is certain.[5] Now, let's say that $\mathbb{E}^o_k$ is an *open evidential combination* if and only if $\mathbb{E}_c \subseteq \mathbb{E}^o_k \subseteq \mathbb{E}$. So an open evidential combination has every certain evidential proposition as a member, and none, some, or all of the uncertain evidential propositions. For example, in our variation on van Fraassen's mouse case, supposing that E_s and E_p are certain but E_d is uncertain, the open evidential combinations are simply $\{E_s, E_p, E_d\}$ and $\{E_s, E_p\}$. (Had E_p also been uncertain, the open evidential combinations would also have included $\{E_s, E_d\}$ and $\{E_s\}$.)

With this notion of open evidential combinations in hand, my suggestion is—to a first approximation—that IBE should be modified to state that one may infer a hypothesis if it provides the best explanation of *any* open evidential combination. More precisely:

> **IBE$_{\text{ER}}$*:** If H would provide a better explanation of any open evidential combination $\mathbb{E}^o_k$ than any other available hypothesis, then one may defeasibly infer H.

The thought here is that regardless of which of the open evidential combinations turns out to be accurate, the inferred hypothesis would provide the best explanation of it. In the limiting case where there is no uncertainty about any evidence, i.e. $\mathbb{E} = \mathbb{E}_c$, the only open evidential combination would be the total evidence itself, i.e. $\mathbb{E}^o_k = \mathbb{E}$. In that special case, IBE$_{\text{ER}}$* reduces to IBE$_{\text{ST}}$. In this sense, IBE$_{\text{ER}}$* is a generalization of standard accounts of IBE.

However, IBE$_{\text{ER}}$* is still too restrictive in some cases of evidential uncertainty. Consider cases in which there is no single hypothesis that pro-

[5] Admittedly, this bifurcation of one's total evidence into certain ($\mathbb{E}_c$) and uncertain ($\mathbb{E} \setminus \mathbb{E}_c$) is somewhat crude, in that (a) it does not capture the way in which evidential uncertainty is arguably a matter of degree, and (b) it requires setting down some—perhaps arbitrary—threshold for an evidential proposition to be included in $\mathbb{E}_c$. However, in so far as our aim is to rescue the idea of IBE as a defeasible inference rule—rather than, say, reducing IBE to a form of Bayesian reasoning—a certain amount of crudeness in this respect seems inevitable.

vides the best explanation of every open evidential combination. Indeed, our variation of van Fraassen's mouse inference is a case in point: H_m provides the best explanation of $\{E_s, E_p, E_d\}$, while H_s provides the best explanation of $\{E_s, E_p\}$. So there would be no hypothesis H that could be inferred via $\text{IBE}_{\text{ER}}{}^*$. However, note that in such cases there may nevertheless be some *other* claim that is robust across these explanatory hypotheses, in that it is *implied* by each one of them. For example, H_m and H_s both imply (H_r) that there is a *r*odent in or near your house. This logically weaker claim, implied by each of the best explanations of the two open evidential combinations, should surely be inferrable by our generalization of IBE_{ST}.

To accommodate this thought, we may modify $\text{IBE}_{\text{ER}}{}^*$ to say that one may infer a claim C if it is implied by each hypothesis that provides the best explanation of some open evidential combination:

> **IBE_{ER}:** If C is implied by each hypothesis H_i that provides a better explanation than any other available hypothesis of some open evidential combination $\mathbb{E}_k^o$, then one may defeasibly infer C.

In the limiting case where there is no uncertainty about any evidence ($\mathbb{E} = \mathbb{E}_c$), IBE_{ER} reduces to the view that one may infer a claim only if it is implied by the hypothesis that provides the best explanation of one's total evidence. Since any hypothesis implies itself, this is itself a generalization of IBE_{ST}, albeit a very natural one (see Lipton 2004, 63-4; Dellsén 2017, 20-21). IBE_{ER} clearly delivers the desired verdict in the variation of van Fraassen's mouse case in which E_d is uncertain, viz. that H_r may be inferred (see Figure 1).[6]

Indeed, it also delivers plausible verdicts in further variations on the case in which, say, E_p is also uncertain. In that case, the open evidential combinations would include, in addition to $\{E_s, E_p, E_d\}$ and $\{E_s, E_p\}$, also $\{E_s, E_d\}$ and $\{E_s\}$. Now suppose—as seems plausible—that H_m and H_s still provide

[6]It also delivers the verdict, promised in footnote 4, that $H_m \vee H_s$ would be inferrable, since that is also an implication of both H_m and H_s. This illustrates a general point about IBE_{ER}, viz. that when different open evidential combinations are best explained by different hypotheses, their disjunction will always be inferrable via IBE_{ER}. Indeed, in such cases, this disjunction will always be the strongest proposition inferrable via IBE_{ER}. However, various weaker but more concrete hypotheses, such as H_r, will *also* be inferrable, and these will often be the more natural inferences to make (though not always; see Figure 3).

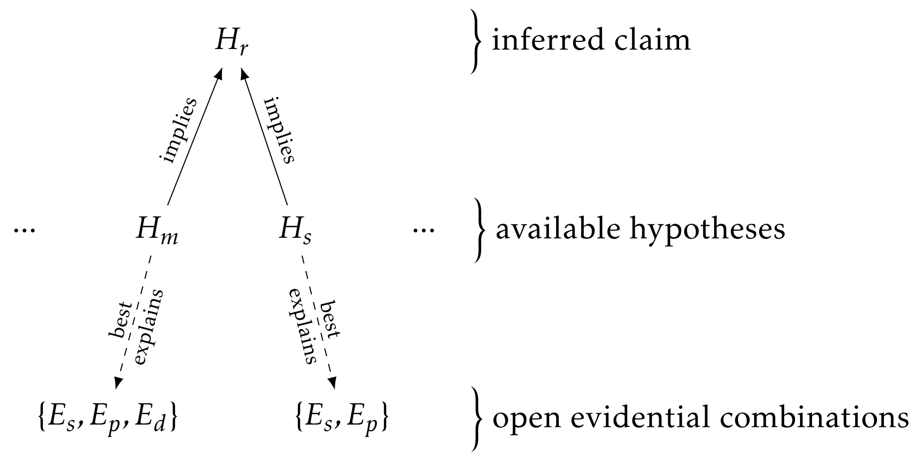

Figure 1: H_r is implied by both H_m and H_s, each of which provides the best available explanation of an open evidential combination.

the best explanations of these open evidential combinations, e.g. because H_m best explains $\{E_s, E_d\}$ and H_s best explains $\{E_s\}$. In that case, H_r would still be inferrable via IBE_{ER} (see Figure 2).

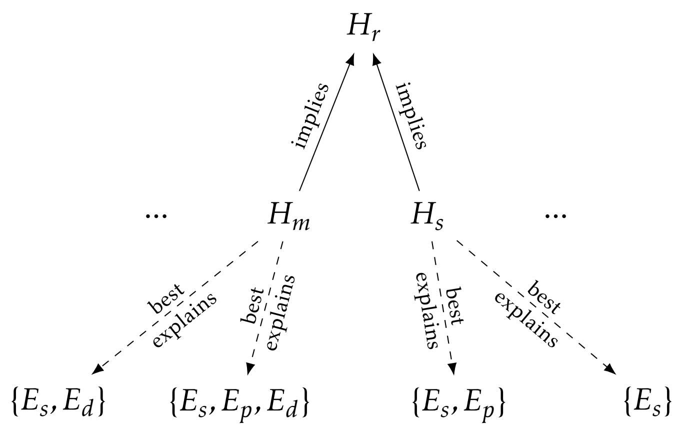

Figure 2: H_r is implied by both H_m and H_s, each of which provides the best available explanation of two open evidential combinations.

With that said, things would be different if some entirely different hypothesis, e.g. the hypothesis (H_p) that you are being *p*ranked by a member of your family, provided the best explanation of one of these evidential combinations, e.g $\{E_s\}$. In that case, IBE_{ER} would not license an inference to either H_r or H_p, although it would allow you to infer their disjunction

$H_r \vee H_p$, which of course is implied by H_m, H_s, and H_p (see Figure 3).

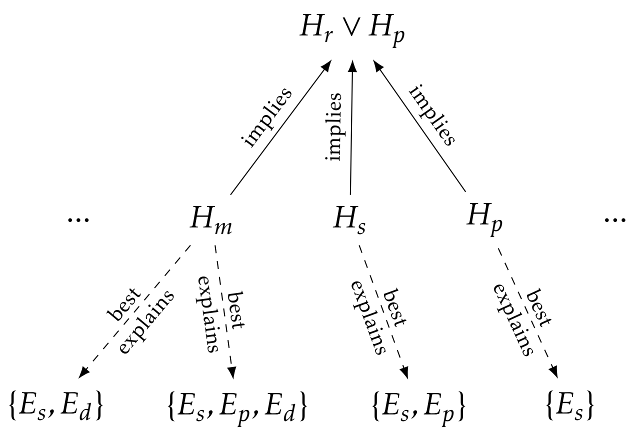

Figure 3: $H_r \vee H_p$ is implied by all of the three hypotheses, H_m, H_s, and H_p, each of which provides the best available explanation of at least one open evidential combination.

As this final variation of the case illustrates, IBE_{ER} will sometimes only license inferences to quite modest conclusions, such as disjunctions of genuinely explanatory hypotheses, especially when one is uncertain about several evidential propositions. Is this a problem for IBE_{ER}? I don't think so. On the contrary, when there is more uncertainty about the evidential propositions from which one is inferring, it seems entirely appropriate to hold off on making ambitious inferences unless and until the evidential situation changes, either by one's becoming more certain about the previously-uncertain evidence, or—which is presumably more common—by one's obtaining other evidence that some of the relevant hypotheses cannot plausibly explain at all. For example, supposing that H_p does indeed provide the best explanation of $\{E_s\}$, and that this is an open evidential combination, it does seem reasonable not to infer H_r unless and until you have gathered further evidence to rule out H_p, e.g. by interrogating your family members. So while IBE_{ER} is a quite cautious inference rule in situations of evidential uncertainty, it is arguably cautious in just the right way.

4 Einstein on the Mt. Wilson Experiments

Thus far, I have motivated IBE_{ER} with variations on a case drawn from everyday life. But does IBE_{ER} also help us understand how hypotheses are inferred from uncertain evidence in scientific practice? In this section, I briefly discuss one case which suggests that it does—although of course further work is needed to evaluate this in a more systematic fashion.

Einstein's special theory of relativity famously discards with the luminiferous ether posited by classical theories of electromagnetism. In so doing, it elegantly explains the famous result, due originally to Michelson and Morley (1887), that no effects of an 'ether drift'—i.e. movement of the earth relative to the ether—were observed in increasingly meticulous experiments. Although special relativity also provided rather compelling explanations of other phenomena, such as the speed of light in water flowing in and against the light's direction of travel (Fizeau experiments), it is safe to say that ether drift experiments were extremely important to establish the plausibility of special relativity in its early years, as Einstein acknowledged (see, e.g., Fölsing, 1998, 219). Moreover, the fact that special relativity elegantly *explains* the lack of an observed ether drift is often taken to show that Einstein used something like IBE in arguing for his theory (e.g. Douven, 2002; Janssen, 2002).

In the 1920s, however, the significance of earlier ether drift experiments were cast into doubt by several new experiments performed at a much greater altitude than before, viz. at the top of Mt. Wilson near Pasadena (Hentschel, 1992). These experimental efforts were led by the highly regarded experimentalist Dayton Miller, who had not only collaborated with Morley on influential replications of the original ether drift experiment in 1905, but was also the president of the American Physical Society. In a series of papers, Miller reported small but positive interference effects of the sort Michelson, Morley, and himself, had failed to observe at lesser altitudes (e.g., Miller, 1925, 1933). These experiments seemed to lend support to versions of classical electrodynamics on which the effects of ether drift could only really be observed at very high altitudes.

Miller's results caused a major upheaval in the scientific community at the time. Einstein himself remained relatively skeptical, however, expressing doubts both privately and publicly about the reliability of Miller's results. But Einstein nevertheless took the apparent results sufficiently seriously to visit Miller specifically to discuss them in person. It seems, then, that Einstein was *uncertain* about Miller's results from Mt. Wilson in a way that he was not regarding various other experiments relevant to his theory, e.g. the Michelson and Morley and Fizeau experiments. Accordingly, we could say—imposing the terminology of §2—that for Einstein there were at least two open evidential combinations relevant to special relativity at the time, viz. one consisting of all the empirical results available to Einstein at the time, and one consisting of all those results except for Miller's results from Mt. Wilson.

Einstein was convinced that it would not be possible to *modify* special relativity so as to provide a plausible explanation of ether drift. The reason for this was Einstein's insistence that ether drift would contradict the constancy of the speed of light, which in turn forms one of the two fundamental principles of his theory. The internal coherence of special relativity would therefore be irredeemably lost by any accommodation of Miller's results within the theory. Accordingly, while special relativity remains the best explanation of the open evidential combination in which Miller's results are excluded, Einstein seems to have ruled out an appropriately modified version of special relativity as the best explanation of the more inclusive open evidential combination in which Miller's results are included. Since there was thus, by Einstein's lights, no weaker version of special relativity implied by the best explanations of *both* evidential combinations, no such version of special relativity could be inferred via IBE_{ER} while these evidential combinations are both open. This accords with Einstein's insistence that "[i]f the results of the Miller experiments were to be confirmed, then relativity theory could not be maintained" (Einstein, 1926; quoted in, and translated by, Hentschel, 1992, 606).

Of course, this situation did not last forever, and IBE_{ER} also sheds light on how the situation was resolved. Subsequent analyses of Miller's results suggested that they were due to temperature changes that caused unequal expansion of certain rods in the experimental setup, leading to an unexpected experimental error. Although Einstein seems to have suspected that something like this might explain Miller's results, there was no way to know this for certain until several years later (Shankland et al., 1955). At that point, however, IBE_{ER} suggests that special relativity could once again be defeasibly inferred, since the evidential uncertainty about Miller's experiments had been resolved in such a way that alternatives to special relativity would no longer provide best explanations of any open evidential combination. Special relativity could, then, reassume its position as the default theory in its domain.

5 Conclusion

I have argued that standard accounts of IBE face a problem of uncertain evidence that is analogous to a well-known problem for Bayesian Conditionalization. If IBE is to survive as a rule of inference, conceived of either as a fundamental rule of inference or as a heuristic for some other form of reasoning, then IBE must be modified so as to accommodate circumstances in which one's evidence is uncertain. I have sketched a potential solution to this problem in the form of a generalization of IBE—IBE_{ER}—which roughly requires an inferred claim to be implied by the best explanations of any combination of one's evidence that the evidential uncertainty does not rule out.

It is worth noting that IBE_{ER} does not require agents to make probabilistic calculations, e.g. by weighing different pieces of evidence by a quantitative estimation of their (un)certainty. Hence, although IBE_{ER} is in some respects rather crude (at least in comparison to highly idealized updating rules that can handle uncertainty, such as Jeffrey Conditionalization), it holds some promise of constituting a workable heuristic inference rule of the sort that many have hoped that IBE might provide us with.

4372 words

References

Dellsén, F. (2017). Abductively Robust Inference. *Analysis*, 77:20–29.

Dellsén, F. (2018). The Heuristic Conception of Inference to the Best Explanation. *Philosophical Studies*, 175:1745–1766.

Douven, I. (2002). Testing inference to the best explanation. *Synthese*, 130:355–377.

Douven, I. (2022). *The Art of Abduction*. Cambridge, MA. MIT Press.

Einstein, A. (1926). Meine Theorie und Miller's Versuche. *Vossische Zeitung*, January 19, 1926.

Fölsing, A. (1998). *Albert Einstein: A Biography*. Penguin, New York.

Harman, G. (1965). The Inference to the Best Explanation. *The Philosophical Review*, 74:88–95.

Harman, G. (1989). *Change In View: Principles of Reasoning*. MIT Press, Cambridge, Mass.

Hentschel, K. (1992). Einstein's attitude towards experiments: Testing relativity theory 1907-1927. *Studies in History and Philosophy of Science Part A*, 23:593–624.

Huemer, M. (2009). Explanationist Aid for the Theory of Inductive Logic. *British Journal for the Philosophy of Science*, 60:345–375.

Janssen, M. (2002). Reconsidering a scientific revolution: The case of Einstein versus Lorentz. *Physics in Perspective*, 4:421–446.

Jeffrey, R. (1965). *The Logic of Decision*. McGraw-Hill, New York.

Lipton, P. (2001). Is Explanation a Guide to Inference? A Reply to Wesley Salmon. In Hon, G. and Rakover, S., editors, *Explanation: Theoretical Approaches and Applications*, pages 93–120. Kluwer Academic Publishers, Dordrecht.

Lipton, P. (2004). *Inference to the Best Explanation*. Routledge, London, 2nd edition.

Lycan, W. G. (1985). Epistemic Value. *Synthese*, 64:137–164.

McCain, K. and Moretti, L. (2022). *Appearance and Explanation*. Oxford University Press, Oxford.

Michelson, A. A. and Morley, E. W. (1887). On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, 34:333–345.

Miller, D. C. (1925). Ether-drift experiments at Mount Wilson. *Proceedings of the National Academy of Sciences of the United States of America*, 11(6):306–314.

Miller, D. C. (1933). The ether-drift experiment and the determination of the absolute motion of the earth. *Rev. Mod. Phys.*, 5:203–242.

Okasha, S. (2000). Van Fraassen's Critique of Inference to the Best Explanation. *Studies in the History and Philosophy of Science*, 31:691–710.

Shankland, R. S., McCluskey, S. W., Leone, F. C., and Kuerti, C. (1955). New analysis of the interferometer observations of Dayton C. Miller. *Reviews of Modern Physics*, 27:167–178.

Thagard, P. R. (1978). The Best Explanation: Criteria for Theory Choice. *The Journal of Philosophy*, 75:76–92.

Trpin, B. (2020). Jeffrey conditionalization: Proceed with caution. *Philosophical Studies*, 177(10):2985–3012.

Trpin, B. and Pellert, M. (2019). Inference to the best explanation in uncertain evidential situations. *British Journal for the Philosophy of Science*, 70:977–1001.

van Fraassen, B. C. (1980). *The Scientific Image*. Clarendon, Oxford.

van Fraassen, B. C. (1989). *Laws and Symmetry*. Clarendon Press, Oxford.

Weisberg, J. (2009a). Commutativity or holism? A dilemma for conditionalizers. *British Journal for the Philosophy of Science*, 60(4):793–812.

Weisberg, J. (2009b). Locating IBE in the Bayesian Framework. *Synthese*, 167:125–143.

Weslake, B. (2013). The problem of disjunctive explanations. Unpublished manuscript, available at <https://philpapers.org/rec/WESTPO-11>.