

There Are No Bad Lots.*

Kevin Davey

Abstract

In this paper I argue that the bad lot objection only manages to refute poorly formulated versions of Inference to the Best Explanation. A version of Inference to the Best Explanation is then formulated that is unrefuted by bad lot cases. A consequence that is pointed out is that this provides a way of detaching debates about realism and antirealism in science from questions about whether we should or should not accept Inference to the Best Explanation.

1 Introduction.

Perhaps the most common criticism of Inference to the Best Explanation (henceforth IBE) is that it is refuted by ‘bad lot’ cases. These are cases in which although one explanation of some phenomenon strikes someone as better than any other they know, they are still not justified in believing the truth of that explanation. The goal of this paper is to argue that the kinds of examples of bad lots generally proposed in the existing literature only manage to refute poorly formulated versions of IBE, and that a properly formulated version of IBE is able to resist them.

In §2 I introduce some different formulations of IBE, discussing their motivations. In §3 I introduce a basic bad lot example, and show how certain simple maneuvers to avoid it are not effective. In the remainder of the paper I distinguish two types of bad lots - *easy* bad lots and *hard* bad lots. In §4 I discuss easy bad lots, and introduce a version of IBE that is not refuted by them. In §4 I argue that this very same version of IBE is also unrefuted by hard bad lots. The version of IBE I propose is therefore able to avoid bad lots cases quite generally. In §6 I make some concluding remarks, including some comments on the role IBE plays in the realism / anti-realism debate.

*An earlier version of this paper when submitted for presentation to the 2022 PSA conference had the title ‘*Bad News for Inference to the Best Explanation, but Good News for the Epistemology of Science.*’ The paper was subsequently reworked in such a way that this title no longer made sense, and so it has now been renamed.

2 Some formulations of IBE.

The rule of IBE has been formulated in the literature in many different ways. To state the simple version of IBE with which we begin, let us suppose some phenomenon P calls out for explanation. We then have the following principle with which we might reason:

Of all the explanations H_1, \dots, H_n of P available to me, H is the best. IBE₁
 H is true.

This is essentially ABD1 of Douven [3] (p. 45), or Schubach's version of IBE in [11] (p. 58).

The question of what makes an explanation the 'best' of those presently available to me is a thorny one, and I will not try to define this term with any precision. Suffice it to say that judging that an explanation scores high relative to its rivals on measures such as simplicity, explanatory power, prior plausibility and so on is generally regarded as sufficient for concluding that an explanation is the best of those presently available to me. This rough characterization will suffice for the purposes of this paper.

It is clear that the truth of the premise in IBE₁ does not guarantee the truth of the conclusion. This of course need not worry us, as IBE₁ is a principle of non-deductive reasoning, and thus should not be expected to meet deductive standards of validity. But in what sense then is IBE₁ a good principle of reasoning? Might it be the case that the truth of the premise merely guarantees that the conclusion is *probably* true? This proposal does not seem correct. There are surely cases in which although I can correctly judge which of the explanations available to me is best, even the best explanation I have relies on some extraordinary co-incidence which I know to have objectively low probability, even given the (admissible) evidence I have.

Might we then claim that IBE₁ is a good principle of reasoning in that it *reliably* takes us from a true premise to a true conclusion (or even just a probably true conclusion), without always doing so? (Schupach [11] suggests something like this.¹) This would mean that the sets of potential explanations available to us generally include the truth. Van Fraassen argues in [14] (p. 143) that this would amount to the claim that human inquirers enjoy a type of epistemic *privilege*, and goes on to argue that there is little reason to suppose this is true. Separate from van Fraassen's objection, also note that the claim that the sets of potential explanations available to us tend to include the truth can only be useful to us in the context of an inquiry I if we have reason to think that I is typical of all inquiries, and that we can therefore use facts about typical inquiries when pursuing I . But it is far from clear that we are entitled to do this in anything like the variety of cases in which proponents of IBE₁ want to use this principle.

¹Schupach's actual claim is that IBE '*reliably preserves good material content*' ([11], p. 59). What exactly this means is never explained. I interpret Schupach here as claiming that IBE reliably takes us from a true premise to a true conclusion.

I think arguments like these are compelling objections to these particular ways of trying to understand the goodness of principle IBE₁.² However, this does not mean that any way of trying to understand the goodness of principle IBE₁ is doomed to fail. An alternative way of understanding the goodness of the principle IBE₁ is that whenever belief in its premise is justified, then so is belief in its conclusion. That is to say, consider the following principle

$$\frac{\text{I am justified in believing that of all the explanations } H_1, \dots, H_n \text{ of } P \text{ available to me, } H \text{ is the best.}}{\text{I am justified in believing that } H \text{ is true.}} \quad \text{IBE}_2$$

If the truth of the premise of this argument entailed the truth of its conclusion – that is, if IBE₂ were deductively valid – then this fact would give an entirely satisfactory sense in which IBE₁ is a good principle of reasoning. To say that IBE₂ is deductively valid is, however, not to say that whenever the premise of IBE₁ is true then its conclusion is probably true, nor is it even to say that this conditional generally holds. It is completely compatible with IBE₂ that there are many situations in which of the explanations available to me, H is the best, and yet H is not even probably true. Such situations would simply have to be ones in which either (i) I am *not* justified in believing that H is the best explanation available to me (even though it is in fact the best such explanation), or (ii) I *am* justified in believing H , in spite of it turning out not to be probably true in any substantive sense.

I think in fact that a case can be made that a slightly tweaked version of IBE₂ is deductively valid, and not refuted by bad lot cases (of any sort), worries about privilege, or any other problems that are sometimes raised against the very idea of IBE. In [1], Dellsen claims that it is possible to have justified belief in something like the premise of IBE₂ but not its conclusion, and thus I assume would claim that IBE₂ is deductively invalid. I think that his sort of worry can also be adequately addressed. Much of the rest of the paper will be spent defending the deductive validity of a slightly tweaked version of IBE₂.³

I suspect that many proponents of IBE have believed in the deductive validity of something like IBE₂ all along. However, I also suspect that many proponents of IBE have thought that the deductive validity of something like IBE₂ cannot be some sort of brute fact, but must rather be somehow explainable. They have thus been tempted to posit the reliability or probability of something like IBE₁ in order to explain it. Their hope, I think, was that the reliability or probability of IBE₁ would yield IBE₂ as some sort of ‘corollary’. For an analogy, consider the following principle

$$\frac{\text{I am justified in believing } A \& B.}{\text{I am justified in believing } A.} \quad (\star)$$

²There is of course much ongoing debate on this topic – see [13, 14, 4, 5, 6, 10, 1, 12], for example – but I will not try to survey this vast literature here.

³In fact, I think IBE₂ gives a better way of fleshing out Schupbach’s idea that IBE reliably preserves good material content, rather than his more reliabilist articulation of this idea.

This is a good principle of reasoning, but it is a good principle of reasoning in virtue of the fact that the following is a good argument

$$\frac{A \& B}{A} \quad (\star\star)$$

In just the way that the validity of the epistemic rule (\star) is explained by the validity of the underlying logical principle ($\star\star$), I think some have been tempted to hope that the validity of the epistemic rule IBE_2 could be explained by the reliability or probability of some underlying ‘logical’ principle such as IBE_1 .

But this way of seeing things is not obligatory. I would like to urge that that IBE_2 (or the slightly tweaked version of it I will later introduce) is a deductively valid principle about justified belief that does not need to be grounded in some fact about the logical realm, entirely outside of epistemology. Of course, there would be nothing bad about it turning out that the validity of this variant of IBE_2 could be grounded in some sort of purely logical, non-epistemic principle, but nor would there be anything bad about it turning out that it could not. The point of departure for this paper will be the idea that IBE_2 itself is worthy of investigation, and that the defense of such a principle need not involve some sort of reduction of it to some other very different sort of principle. In what follows we therefore put IBE_1 to the side, and focus on IBE_2 .

3 Best Explanations

In what follows, I will use the term IBE to refer to inference to the best explanation informally and broadly construed, as opposed to any precisification of it such as IBE_1 or IBE_2 .

Perhaps the most famous objection to IBE is the ‘bad lot’ objection, and much of the rest of the paper will revolve around it. This objection, originally introduced by van Fraassen in [14], revolves around the fact that IBE gives only a *comparative* judgment that some explanation H is the best of the set of presently available explanations. It is not clear however what could justify moving from a comparative judgment of this form to the non-comparative judgment that belief in H is in fact justified by the evidence. The set of presently available explanations might have the property that even though some H is the ‘best’ of all its elements, there is nevertheless insufficient evidence to justify believing H . The best explanation drawn from a bad lot is not necessarily worthy of belief, and so IBE_2 is surely not deductively valid.

There is something compelling about this objection. If something calls for explanation, and I am lacking in imagination, fatigued, or intoxicated, the only explanations I might be able to come up with might all be somewhat implausible or far-fetched. Surely the evidence does not justify me believing whatever happens to be the least worst element of this set. The question then is whether adding a few reasonable conditions to IBE can protect it from easy counterexamples like this. For example, some authors (e.g., Lipton in

[6]) add the requirement that the best presently available explanation meet some minimal standard of plausibility in order for us to infer it via an IBE. Any plausible statement of IBE must surely take some requirement of this sort on board. Of course, this minimal standard of plausibility should not be so high that a hypothesis H meeting this standard already renders belief in H justified, for then IBE would be redundant. But it should be high enough to eliminate worries about crudely put together sets of far-fetched explanations. Exactly what the right standard to impose here is not clear, but let us be charitable and suppose that some such standard of ‘minimal plausibility’ can be formulated, and that meeting this standard is a precondition for inferring an explanation via IBE. To avoid carrying around too much jargon, we will henceforth assume that for a hypothesis to be the ‘best’ of a set of candidates it must meet this standard of minimal plausibility. So in a case in which no explanation in some set of possible explanations meets this minimal standard, there is no best element of this set.

It is also sometimes supposed that for us to infer an explanation via IBE, that explanation must be *significantly* better than any competitor. If two potential explanatory hypotheses H_1 and H_2 are both highly plausible, but H_1 is only very slightly superior to H_2 (though both H_1 and H_2 are significantly better than any other competitor explanation), a good case can be made that we should be reluctant to infer H_1 on this basis alone. Instead, we should seek further evidence that reveals more decisively which of H_1 and H_2 is genuinely superior before making any judgments. Thus, it seems reasonable to require that an explanation be *significantly* better than any competitor in order for us to infer it via an IBE. Again, to avoid carrying around too much jargon, we will henceforth assume that for a hypothesis to be the ‘best’ of a set of candidates, it must be significantly better than its rivals in this way. So in the case in which there is only a marginal difference between the strongest explanations in some set, there is no best element of this set.

Adding these sorts of provisos to IBE nevertheless only helps so much in avoiding bad lots. Consider the following example. Suppose one day I wake up to find a mess in the kitchen of the apartment that I share with my roommates α and β . I am quickly filled with anger - it is a horrible mess that I must clean up if I am to prepare my breakfast. Currently not being on good terms with α , I immediately find myself thinking that α created the mess. Suppose that for whatever reason, in my anger the thought that the mess might have been created by β never even crosses my mind. Let us also suppose that the hypothesis that α caused the mess meets our standard of minimal plausibility – any roommate can sometimes create a mess, and so long as we suppose that α is an ordinary person who sometimes creates messes, our minimal plausibility condition will easily be met. However, let us also suppose that objectively speaking, the evidence before me doesn’t really justify my believing that α caused the mess. Let us suppose that had I reflected further on the situation in a calm state, then without acquiring new evidence I would actually have come to regard the hypothesis that β caused the mess as just as likely as the

hypothesis that α caused the mess.

In this example, we may suppose that at the moment the hypothesis that α created the mess occurs to me, it is the only minimally plausible explanation I am aware of. It meets the standard of minimal plausibility, and is significantly better than any minimally plausible rival I know of, because there are no minimally plausible rivals that I know of at that moment. IBE therefore allows me to infer that α caused the mess. But this is surely an unwelcome conclusion, as by assumption the evidence simply *does not* justify concluding that α caused the mess. More generally, IBE₂ always allows me to infer the first minimally plausible explanation that occurs to me in any situation. Such a fact quickly lead to counterexamples to the deductive validity of IBE₂.

Perhaps someone could try to bite the bullet and say that in my moment of anger I *really was* justified in believing that α caused the mess, but that later on when the hypothesis that β caused the mess entered my consciousness and I realized that it was just as likely to be true, I was no longer justified in believing that α caused the mess. But this seems to me to be a very unintuitive way of thinking of justification. Surely we want to say, for example, that anger sometimes makes us believe things we are not justified in believing; or that anger can even make us believe we are justified in believing things that in fact we are not justified in believing at all. Not only anger, but delusion, confusion, blindspots or love can also do the same thing. Giving up this conception of anger, delusion, confusion, blindspots and love is surely too much. And so I think that the bullet in question is simply too unpleasant to bite. I think it must be conceded that in my anger, I simply *wasn't* justified in believing that α caused the mess, and thus that IBE₂ is not deductively valid. Insofar as the example I have described is that of a 'bad lot' case, it looks like the bad lot objection is good.

To defend IBE₂, perhaps it could somehow be required that a certain minimum number of hypotheses be entertained before applying IBE₂, thereby blocking the consequence that I may infer the first minimally plausible explanation that occurs to me in any situation. But this approach is not promising. For even the idea that I am justified in inferring the best of the first 7 explanations that occur to me is no less subject to unpleasant counterexamples. In my anger I might after all first come up with 7 explanatory hypotheses (at least one of which is minimally plausible and significantly better than the other 6) before the important rival explanatory hypotheses that I really ought to be considering manage to enter my consciousness.

4 Easy Bad Lots

Counterexamples to IBE of this sort are so easy to generate that one cannot help but wonder whether they are simply a consequence of working with a poor formulation of IBE. In the kind of counterexamples I have been considering, an agent is working with an artificially restricted set of possible explanations for

some phenomenon – that is to say, for some reason (anger, delusion, confusion, blindspots, love etc.), there is some highly plausible explanation that they are not considering that they really ought to be considering. It is this failure to exercise appropriate epistemic due diligence which leads us to think that they cannot be justified in inferring that, because something seems to them to be the best explanation, it must be true. Call bad lot cases that arise from this sort of failure to exercise appropriate epistemic due diligence *easy bad lot cases*. (We will contrast this with *hard bad lot cases* in the next section.)

Note that in the sorts of scientific cases where applications of IBE look most compelling, we are not working with small sets of explanatory possibilities haphazardly put together in a moment of anger, but rather with larger sets of explanatory possibilities very carefully generated and refined over a long period of time. In such cases, due diligence has been meticulously exercised in generating the set of rival explanatory hypotheses that must be carefully weighed against each other in order to determine the ‘best’. The obvious (and I think, entirely correct) response to easy bad lot cases should therefore be to add the requirement to IBE that the set of candidate explanatory possibilities be the result of a process performed with sufficient due diligence, in such a way that no candidate hypotheses that really ought to have been considered has been omitted from this set.

There are different ways to build this sort of requirement into IBE. The way I will choose involves broadening the conception of which explanations are ‘presently available’ to me. It is tempting to think that what makes an explanation presently available to me is its psychologically occurring to me in some explicit sort of way. That is certainly one conception of present availability, but it is not the only one. The word ‘available’ has a modal character – it refers to the things of which it is *possible* for me to avail myself. Like any modal, the corresponding notion of possibility can be interpreted broadly or narrowly. It is certainly true that I may avail myself of hypotheses that have explicitly psychologically occurred to me. But there are other things of which I could surely be said to easily avail myself. Suppose I would like an explanation for why a physics experiment returned the result it did. Suppose further that while no good explanation comes to my mind, my best friend who is a physicist is standing right beside me, and happens to know the correct explanation. There is surely a sense in which the explanation the physicist has is easily available to me – I could just ask her, and she would tell me. Likewise for any explanation in a textbook at hand. For a different case, return to the example of the mess in the kitchen. If I were to contemplate the situation just a little more carefully in a calmer state of mind, the hypothesis that β caused the mess would no doubt arise in my mind. This hypothesis too is surely easily available to me, even if right now I my anger stops it from explicitly coming to mind.

For purposes of spelling out the additional requirement on IBE that I think we want, I would like then to distinguish explanations that are easily available to me from those that are not. The concept of ‘easy availability’ is of course vague, and there is no sharp line separating the things that are easily available

to me from those that are not. No doubt there is context dependence here. Nevertheless, there is surely a clear sense in which general relativistic explanations for cosmological phenomena were not easily available to the ancient Greeks, while Ptolemaic explanations were. There is surely a clear sense in which the standard explanation of natural selection in terms of genes is easily available to more or less anyone in the 21st century, even if they have not learned about it yet. And there is surely a clear sense in which various commonsensical explanations for everyday phenomena that I would easily come up with given a little time are easily available to me, while rival explanations requiring extraordinary ingenuity are not. A concept can be intuitive and philosophically useful even if it is vague or context dependent, and in this spirit I would like to use the concept of easy availability to get a clearer understanding of IBE.

In cases where IBE seems most compelling, it seems to me that I begin with a judgment that a certain explanation is the best not just of those that have explicitly psychologically occurred to me, but rather of those easily available to me. For example, when in preparation for an IBE an expert scientist carefully generates a set of rival explanatory hypotheses for some phenomenon in consultation with the broader scientific community and goes on to judge that some hypothesis H is the ‘best’ of these, it seems to me that they are typically making *more* than the claim that H is the best of the rival explanatory hypotheses of which they are explicitly psychologically aware. Rather, they are (rightly or wrongly) making the stronger claim that H is the best of *any* hypothesis easily available to them. In the context of good science, the standards for ‘easy availability’ are generally sufficiently liberal that this includes a large set of possible explanations. The scientist’s claim that some hypothesis is the best of those easily available to them is therefore a bold one.

Of course, if the scientist is very confident that having exercised sufficient due diligence they are explicitly psychologically aware of *all* easily available explanations, then the move from the claim that H is the best of all explanations of which they are explicitly aware to the claim that H is the best of all explanations easily available to them is a trivial one. But it need not be the case that someone needs to be explicitly psychologically aware of all easily available explanations in order to have justified belief that some H is the best of all explanations easily available to them. There are presumably cases in which an expert can judge in a justified way that H is the best of all explanations easily available to them merely by considering an appropriately large sample of rivals to H , for example. (Mayo seems to think this is sometimes possible when one has *severely tested* a hypothesis.⁴) The details of how this might be possible are not essential for the argument of this paper however. For our purposes, it will suffice to note that we can and often do make justified claims that some H is the best of all explanations easily available to us of some phenomenon. I shall claim that this is what we should take as the starting point of an IBE.

⁴See in particular the discussion of Perrin and Brownian motion in chapter 7 of [7].

Armed with this, I propose the following version of IBE:

I am justified in believing that H is the best
of all easily available explanations for P . IBE₃
I am justified in believing that H is true.

Easy bad lot cases are not counterexamples to IBE₃, because in easy bad lot cases there is an easily available, highly plausible explanation being ignored. The premise of IBE₃ is therefore false in such cases. Such cases are therefore not counterexamples to the deductive validity of IBE₃.

5 Hard Bad Lots

Eliminating the threat of easy bad lots does not completely get rid of the bad lot problem. For one might wonder whether there are cases in which through meticulous investigation, I am quite aware of all (or enough) of the easily available explanations and can form a justified belief that some H is the best of all easily available explanations, and yet I am not justified in believing H itself. Call such a case (if it exists) a *hard bad lot case*. A hard bad lot case would refute the deductive validity of IBE₃, and so we must ask whether there really are hard bad lot cases.

It has I think seemed obvious to many that there must be hard bad lot cases. Surely there are cases in which I am not justified in believing that the truth lies among the hypotheses easily available to me, even though of the explanations that happen to be easily available to me, one is clearly the strongest. Perhaps this occurs when we feel our grasp of some domain of inquiry is lacking, in spite of even the experts having exercised as much due diligence as possible in investigating it in order to explain some phenomenon. For example, in [2] Dellsen argues that in the 1920s, even though it was conceded that Fresnel's transverse wave theory of light provided the best available explanation for various optical phenomena, physicists did not take themselves to be justified in believing it.⁵ Novick and Scholl argue in [9] and [8] that biological theories that do not identify sufficiently plausible causes are not regarded as objects of justified belief, regardless of how well they explain the phenomena. Arguably, string theory might be the best explanation of various phenomena of quantum gravity, though it is hard to imagine arguing that we are justified in believing

⁵In particular, on p. 163 of [2] Dellsen makes the following claim of Fresnel's transverse wave theory of light:

... Fresnel's theory was already in 1819 considered to be explanatorily superior to its available alternatives, including Newton's corpuscular theory, which had been accepted up to that time. So Fresnel's new theory was considered lovelier than Newton's theory, which must have been considered sufficiently lovely to be accepted. And yet Fresnel's theory was viewed with considerable suspicion by many prominent optical physicists for most of the 1820s.

it solely in virtue of this fact. Perhaps some or all of these provide examples of hard bad lots.

I will argue that in spite of initial appearances, such examples are not counterexamples to the deductive validity of IBE_3 . In particular, I maintain that in these cases the premise of IBE_3 is false, and that we are not in fact justified in believing that the hypothesis H in question is the best of all easily available explanations for the phenomena in question. I will return to some of these specific examples shortly, but it will be helpful to begin with a much more accessible hard bad lot case with which to sharpen certain intuitions.

Suppose I have heard rumors that Bigfoot exists somewhere in some specific large forest, and that I decide to investigate the matter for myself. Let us suppose that I start my investigations with no strong convictions either way as to whether these rumors are true. Let us also suppose that I start with only highly superficial knowledge of the wildlife of the region, and no real sense of what Bigfoot-like creatures might be. Furthermore, suppose that even the best experts I know of are in essentially this position – they all know very little about the wildlife in the forest in question and what Bigfoot-like creatures might be. If because our current extensive knowledge of the wildlife on Earth this sort of ignorance is hard to imagine, the modern reader should imagine themselves in a time in which such extensive knowledge of the Earth’s wildlife did not exist and was not in any way easily available. In what follows, let H_0 be the hypothesis that Bigfoot does *not* exist in this forest.

Suppose I begin my investigation with a quick and relatively superficial search of the forest, failing to catch sight of Bigfoot. At this point, it is of course too soon to say that the best explanation of my failure to find Bigfoot is H_0 . There will still be numerous moderately plausible rival explanations to H_0 as to why I have not spotted Bigfoot that I am not yet in a position to dismiss and that are easily available to me – perhaps Bigfoot only comes out at night and my quick search occurred during the day, or perhaps Bigfoot has very sensitive hearing and was always able to hear me coming and run away, or perhaps Bigfoot hibernates during the winter which is when I conducted my quick search, and so on. Recognizing these moderately plausible rival explanations to H_0 of my failure to spot Bigfoot, I might modify my search procedure in various ways, doing a few random searches during the night, during the summer, and leaving carefully concealed video cameras at a number of points in the forest to try and catch sight of Bigfoot. Let us suppose that none of these further investigations yields a sighting of Bigfoot.

Perhaps in consultation with others I might try to come up with yet further minimally plausible rival explanations to H_0 as to why I have not spotted Bigfoot, and test them all accordingly. Assuming that all these searches also fail to yield a sighting, I will at some point come to think that every minimally plausible, easily available rival explanation to H_0 of my failure to spot Bigfoot is unlikely to be true. (Perhaps I do not even have to test literally every rival explanation of this sort in order to be justified in reaching this conclusion.)

Consider something like the following rough list of potential minimally plau-

sible explanations of why I have not observed Bigfoot:

H_0 : Bigfoot does not exist in the forest.

H_1 : Bigfoot exists, but desires to avoid me and my equipment and is able to skillfully do so.

H_2 : Bigfoot exists and is not deliberately trying to avoid contact with me, but simply lies in a part of the forest I have not checked yet.

...

At this point in my investigation as I have described, it looks very reasonable for me to rank H_0 ahead of anything else on this list. The more I investigate the forest, the more H_1 , if it were true, would be a surprising fact that in turn stands in need of explanation – exactly *how*, after all, does Bigfoot continue to avoid contact with me? The fact that H_1 stands in need of explanation in this way renders it much less explanatorily attractive than H_0 . The hypothesis H_2 also becomes less plausible the more I search, and we can suppose that after the searches described is starting to look decently implausible, even if it still meets our standards for minimally plausibility. And so it looks like H_0 is the best explanation I have of my failure to spot Bigfoot.

But in spite of this, it must be remembered that at this stage I still do not have a clear sense of exactly what wildlife exists in the forest. Perhaps when conducting my investigations in the forest I have been continually surprised by unusual botanical and biological phenomena, and I still lack any sort of grasp of the constitution and ecology of the forest. I still have no right to expect that what I discover as I continue to explore the forest will resemble what I have discovered so far. Given this, would I really be justified in this case in asserting that Bigfoot does not exist in this forest? Is my evidence really sufficient to justify belief in H_0 ? It is very natural to think that it is not, and that I would not yet be justified in believing H_0 . This is in spite of the fact that H_0 seems to be the most reasonable, easily available explanation for my failure to spot Bigfoot. And so it looks like we have a hard bad lot case – that is, a case in which the premise of IBE_3 is true, and its conclusion false.

I think however that this analysis of the situation is in error. First of all, I think that the intuition that belief in H_0 would not be justified in this situation is a good one. Why? To say that I am justified in believing H_0 is to say in part that I am justified in thinking that my future investigations will not yield a sighting of Bigfoot. (How could someone meaningfully say that they have justified belief in X , but not justified belief that subsequent investigations will yield results in accordance with X ?) My past investigations of the forest, however, have been consistently surprising enough that I cannot claim to have evidence that my future investigations will yield results in accordance with H_0 . I may of course *conjecture* that Bigfoot does not exist in this forest, and this conjecture may seem to me to be a very compelling one, but I do yet not have sufficient evidence to justify thinking that this conjecture will *continue* to look compelling as I gather more evidence. That would require a type of knowledge of the forest that I simply lack at present. Once my investigations of the forest

pass a certain point of thoroughness, I may well be justified in thinking that further investigation will not uncover evidence that would render belief in H_0 unjustified. But until I reach that point, I do not have justified belief in H_0 . I think this sort of reasoning lies behind the intuition that I do not have justified belief in H_0 . Moreover, I think this sort of reasoning is basically correct.

I think something like this criterion is being used in declaring that in the early 1820s we would not have been justified in believing Fresnel's transverse wave theory of light, or in declaring that we are not presently justified in believing string theory. In the case of Fresnel's transverse wave theory, Dellsen diagnoses the unwillingness of the scientists of the time to accept it as follows:

Fresnel's presentation of the theory in 1818 acted as a catalyst for speculation about whether yet other theories, currently unconceived, could explain the relevant phenomena as well or better. [2], p. 167.

Presumably, this is a situation in which scientists took themselves not to be justified in thinking that further evidence would continue to confirm Fresnel's theory, as opposed to some (in this case, yet unknown) rival. I agree with Dellsen that in such a situation, one could not be said to have justified belief in Fresnel's theory. In the case of string theory, the case might be made that we simply do not understand the phenomena of quantum gravity well enough to be confident that further evidence will not radically change what we are warranted in believing in this domain. I think that these sorts of considerations correctly show that in these cases we do not have justified belief in these theories.

The problem however is that precisely the same considerations speak against the claim that I have justified belief that H_0 is the best easily available explanation for my inability to spot Bigfoot so far. Certainly, it is not unreasonable for me to conjecture that H_0 is the best easily available explanation for my inability to spot Bigfoot. But of course, as I learn more about the forest, that might change. I might discover that there are animals in the forest with extraordinary abilities to flee more rapidly than any animal of which I presently know, at which point H_1 might actually start to look as plausible as H_0 . Perhaps it could turn out that the habitable part of the forest is much greater than I had previously estimated, so that H_2 becomes a more serious rival than it was before. If I remain sufficiently ignorant of the forest, I ultimately have scant grounds for thinking that my current view as to which easily available explanation is best will remain unchanged as I continue my investigation. But if I have only scant grounds for thinking that my opinion that H_0 is the best of all easily available explanations will continue to look compelling as I gather more evidence, then I am *not* justified in thinking that H_0 is the best of all easily available explanations for my inability to spot Bigfoot. Again, it might be very reasonable for me to conjecture that H_0 is the best easily available explanation for my inability to spot Bigfoot so far, but such a conjecture would not amount to justified belief. The premise of IBE₃ consequently fails, and so we do not have a counterexample to the deductive validity of IBE₃. This case is not a hard bad lot case after all.

Likewise, given everything else he says, I do not think Dellsen is right in taking the scientists of the 1820s to have had justified belief that Fresnel's wave theory was the best explanation for the optical phenomena with which they were concerned. If the scientists at the time felt that they lacked grounds for thinking that further investigation would not uncover evidence that some other theory provided a better explanation for the phenomena in question, then in virtue of that fact scientists were simply *not* justified in thinking that Fresnel's theory was the best easily available explanation. It might of course have been very reasonable of such scientists to *conjecture* that Fresnel's theory was the best easily available explanation, but such a conjecture would not have risen to the level of justified belief. Likewise for the case of string theory.

I think the intuition that there are hard bad lot cases arises from a temptation to set too low a bar for justified belief that some H is the best of all easily available explanations for P , while maintaining a high bar for justified belief in H itself. The temptation for setting too low a bar for justified belief that some H is the best of all easily available explanations often comes from a certain image of how IBE works. According to this image, to perform an IBE, we write down all the easily available potential explanations for some phenomenon, assign each 'points' for simplicity, explanatory power, prior plausibility and so on, tally up these points in some way or other, and see which scores highest. In this way, we find out which H is the best explanation. It is tempting to think that if we perform this process sufficiently accurately, that we are automatically justified in believing – or even automatically *know* – that H is the best easily available explanation. From this, we infer H .

I do not think, however, that merely performing this sort of tallying process guarantees justified belief that H is the best of all easily available explanations for P . As we acquire more evidence, the results of this sort of tallying process may change, as a result of which H may or may not continue to receive the highest score. To be justified in thinking that H is the best of all easily available explanations for P is not just for the tallying procedure to declare victory for H now, but for me to be justified in predicting that it will continue to do so as more evidence is acquired. Unless I have some reason for thinking that H 's victory will persist, I do not have justified belief that H is the best of all easily available explanations for P , even if the tallying process currently gives H the highest score.

It might be thought that I am requiring too much of justified belief here. Justification is relative to a body of evidence, and it might be thought that to have justified belief in something relative to a body of evidence it is not necessary to have *any* views about how one's justified beliefs might change if further evidence were required. One might use this claim to try to argue that in hard bad lot cases the premise of IBE₃ really is true after all, and the conclusion false, and thus that such cases really do give us counterexamples to the deductive validity of IBE₃.

I do not think however that this way of trying to salvage hard bad lot cases works. First of all, insofar as this lowers our standards for justified belief in the

premise of IBE_3 , in the name of consistency it must also lower the standards for justified belief in the conclusion of IBE_3 . The obvious threat is that this will result in the conclusion of IBE_3 being true after all, and that such examples will again fail to be counterexamples to the deductive validity of IBE_3 . The burden lies on the proponent of hard bad lots to argue that there is a plausible conception of justification here that renders the premise of IBE_3 true and its conclusion false. I know of no promising way to try to meet this burden.

More importantly, I do not think that it is right to view justified belief (relative to a body of evidence) as not entailing anything about what might happen if further evidence is acquired. Justified belief is something that comes at the *end* of a process of evidence gathering; it arises at a point where we have sufficient evidence to render a reasonable verdict about something. Presumably, one of the things that makes a body of evidence sufficient for rendering a reasonable verdict about something is that we have some sort of (perhaps inductive) reason for thinking that further evidence will likely not change this verdict. Without this, it is difficult to see what entitles us to render a verdict, and thus difficult to see how we could have justified belief.⁶ The idea that at every stage of the evidence gathering process we have justified belief relative to the evidence collected thus far thus strikes me as incorrect – surely we only have justified belief in cases in which have grounds for thinking that further evidence collection will likely not change things much. In this way, justified belief really does involve belief about how one’s beliefs might change if further evidence were acquired. Prior to reaching the point of justified belief, we merely have a conjecture (or perhaps as one gets closer to justified belief, an *intriguing* or *compelling* conjecture), but we must not make the mistake of mistaking these with justified belief.

If justification is thought of in this way, it seems to me that in traditional examples of hard bad lots the premise of IBE_3 is not true, and thus that such cases are not counterexample to the deductive validity of IBE_3 . Hard bad lots are therefore not a threat to IBE, correctly understood.

6 Conclusion

Appropriately formulated, IBE is not refuted by ‘bad lot’ examples, either easy or hard. This of course does not on its own constitute an argument that IBE is one of our basic principles of reasoning, as it may face other challenges. If IBE fails, however, it will not be because of anything like van Fraassen’s bad lot worry.

Some of the ideas of this paper have interesting consequences for the realism / anti-realism debate in the philosophy of science. Much literature in this debate has revolved around the thought that the scientific realist is someone who accepts IBE, and in virtue of that, has an easy IBE-based argument for

⁶In different language, justified belief only arises after the activation of some sort of appropriately chosen ‘stopping rule’.

realism, while the anti-realist is someone who rejects IBE, and thus is unable to infer the truth of scientific claims from their empirical adequacy. (This way of understanding the division between realists and anti-realists goes back at least to van Fraassen's seminal works [13, 14].) But is this way of viewing what is at stake in the realism / anti-realism debate correct? The anti-realist need not, after all, reject IBE₃ – it would suffice for them to simply reject the premise of IBE₃ in the situations in which the realist wants to apply it. In particular, it would suffice for the anti-realist to make the following claim:

Anti-Realism: We are not justified in believing that the best of all easily available explanations for the success of some scientific theory is the truth of that scientific theory.

In fact, van Fraassen's 'neo-Darwinian' conception of scientific theories, according to which the success of scientific theories is not due to their truth, but rather due to the way in which they are '*born into a life of fierce competition, a jungle red in tooth and claw*' in which '*only the successful theories survive*' ([13] p. 40), seems to be a way of thinking about scientific theories that avoids realism *not* first and foremost by rejecting IBE₃, but rather by making the anti-realist claim above and rejecting the *premise* of IBE₃ in some crucial case. Although anti-realists like van Fraassen have rejected IBE₃, in fact this need not be considered a central part of the anti-realist platform, as their rejection of the premise of IBE₃ in certain crucial cases is enough for them to make their anti-realist argument.

It thus seems to me that the realism / anti-realism debate has *not*, in spite of appearances, really been a debate about the admissibility of IBE, but has rather been a debate about whether the premises of certain putative IBEs are true. Getting clear on the structure and premises of IBE as we have done helps to make this clear. Anti-realists are skeptical about the possibility of justified belief that the best of all easily available explanations for the success of some scientific theory is the truth of that scientific theory, and thus make the anti-realist claim above, while realists claim to the contrary that justified beliefs of this form are in fact possible. This helps to make clear what need and need not be under dispute in the realism / anti-realism debate.

References

- [1] Dellsen, F., *Reactionary responses to the Bad Lot Objection*, Studies in History and the Philosophy of Science 61, 32-40.
- [2] Dellsen, F., (2021) *Explanatory Consolidation: From Best to Good Enough*, Philos Phenomenol Res. 2021;103:157-177.
- [3] Douven, I.,(2022) *The Art of Abduction*, MIT press.
- [4] Kincaid, H. and Day, T. (1994), *Putting inference to the best explanation in its place*, Synthese 98 (2):271-295.

- [5] Ladyman, J., Douven, I., Horsten, L., van Fraassen, B., (1997), *A Defence of van Fraassen's Critique of Abductive Inference: Reply to Psillos*, The Philosophical Quarterly, Volume 47, Issue 188, 305-321.
- [6] Lipton, P. (2004) 'Inference to the Best Explanation,' 2nd edition, Routledge.
- [7] Mayo, D. (1996) 'Error and the Growth of Experimental Knowledge', University of Chicago Press.
- [8] Novick, A. (2017), 'Metaphysics and the Vera Causa Ideal: The Nun's Priest's Tale', *Erkenn* (2017) 82:1161–1176
- [9] Novick, A. and Scholl, R. (2022) *Presume It Not: True Causes in the Search for the Basis of Heredity*, *Brit. J. Phil. Sci.* 71, 59–86.
- [10] Psillos, S., (1996), *On Van Fraassen's Critique of Abductive Reasoning*, The Philosophical Quarterly, Vol. 46, No. 182, pp. 31-47.
- [11] Schupbach, J. (2014), *Is the Bad Lot Objection Just Misguided?*, *Erkenntnis* 79 (1):55-64.
- [12] Shaffer, M., (2021), *Van Fraassen's Best of a Bad Lot Objection, IBE and Rationality*, *Logique and Analyse*, 255, pp. 267-273.
- [13] van Fraassen, B. (1980) 'The Scientific Image', Oxford University Press.
- [14] van Fraassen, B. (1989) 'Laws and Symmetry', Oxford University Press.