There Are No Bad Lots, Only Bad Formulations of Inference to the Best Explanation.

FORTHCOMING IN THE BRITISH JOURNAL FOR THE PHILOSOPHY OF SCIENCE.

Kevin Davey

Abstract

In this paper I argue that van Fraassen's famous bad lot objection against Inference to the Best Explanation only manages to refute poorly formulated versions of this rule of inference. I formulate a version of Inference to the Best Explanation for which I claim bad lot cases do not arise. I do this by requiring that one have certain sorts of justified beliefs about the set of all easily available explanations for a phenomenon before one can apply Inference to the Best Explanation. This forces Inference to the Best Explanation to occur only relatively late in the process of inquiry when one already has fairly sophisticated knowledge of the subject matter in question, thereby stopping bad lot cases from arising.

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 main 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 and discuss 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 §5-6 I discuss hard bad lots, and in §7 I refine IBE further to yield a version that is unrefuted by them. This final version of IBE therefore avoids bad lots cases quite generally. In §8 I make some concluding remarks, suggesting what we should regard as the most serious threat to IBE now that bad lots have been disposed of.

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:

$$\frac{\text{Of all the explanations of } P \text{ available to me, } H \text{ is the best.}}{H \text{ is true.}} \qquad \text{IBE}_1$$

This is essentially ABD1 of (Douven [2022], p. 45), or the version of IBE presented in (Schupbach [2014], 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 IBE_1 does not guarantee the truth of the conclusion. This of course need not worry us, as IBE_1 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 IBE_1 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, the set of explanations is so poor that even the best of them is not particularly likely to be true (I will give an example of this in §3.)

Might we then claim that IBE_1 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 [2014]) seems to suggest something like this.[1] If this were true, it would follow that the sets of potential explanations available to us generally include the truth. Van Fraassen claims in (van Fraassen [1989], 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 IBE_1 want to use this principle.

I think arguments like these are compelling objections to these particular ways of trying to understand the goodness of IBE_1.[2] However, this does not mean that any way of trying to understand the goodness of IBE_1 is doomed to fail. An alternative way of understanding the goodness of IBE_1 is that whenever belief in its premise is justified, then so is belief in its conclusion. That is to say, consider the following principle

$$\frac{\begin{array}{l}\text{I am justified in believing that of all the}\\ \text{explanations available to me, } H \text{ is the best.}\end{array}}{\text{I am justified in believing that } H \text{ is true.}} \qquad \text{IBE}_2$$

If the truth of the premise of this argument entailed the truth of its conclusion – that is, if IBE_2 were deductively valid – then this fact would give an entirely satisfactory sense in which IBE_1 is a good principle of reasoning.[3] To say that IBE_2 is deductively valid is, however, not to say that whenever the premise of IBE_1 is true then its conclusion is probably true, nor is it even to say that this conditional generally holds.

I think in fact that a case can be made that a slightly tweaked version of IBE_2 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 (Dellsen [2017]), Dellsen claims that it is possible to have justified belief in something like the premise of IBE_2 but not its conclusion, and thus I assume would claim that IBE_2 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 IBE_2.

I suspect that many proponents of IBE have believed in the deductive validity of something like IBE_2 all along. However, I also suspect that many proponents of IBE

[1] Schupach's actual claim is that IBE '*reliably preserves good material content*' (Schupbach [2014], p. 59). What exactly this means is never explained. I interpret Schupbach here as claiming that IBE reliably takes us from a true premise to a true conclusion.

[2] There is of course much ongoing debate on this topic – see (van Fraassen [1980, 1989]), (Kinkaid and Day [1994]), (Ladyman et. al. [1997]), (Lipton [2004]), (Psillos [1996]), (Dellsen [2017]), and (Schupbach [2014]) for example. I will not try to survey this vast literature here.

[3] I think IBE_2 actually 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.

have thought that the deductive validity of something like IBE_2 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 IBE_1 in order to explain it. Their hope, I think, was that the reliability or probability of IBE_1 would yield IBE_2 as some sort of 'corollary'. For an analogy, consider the following principle

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

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 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 and some variants of it.

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 attempted 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 (van Fraassen [1989]), amounts to the worry that there will be scenarios in which even though some H is the 'best' of all presently available explanations, there might nevertheless be insufficient evidence to justify belief in 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 am able to come up with might all be somewhat implausible, far-fetched, or lame in one way or another. 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 works such as (Lipton [2004]) 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. 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 (though we will revisit this requirement later in the paper.) 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 (we will revisit this requirement too later in the paper.) 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 me concluding that α caused the mess. More generally, IBE_2 always allows me to infer the first minimally plausible explanation that occurs to me in any situation. Such a fact quickly leads to counterexamples to the deductive validity of IBE_2 such as the one just given.

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 in believing that α caused the mess, and thus that IBE_2 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 IBE_2, perhaps it could somehow be required that a certain minimum number of hypotheses be entertained before applying IBE_2, 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 hypothesis that I really ought to be considering manages 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 in generating a set of potential explanations which leads us to think that they are not justified in inferring that the best explanation available to them must be true. Call bad lot cases that arise from this sort of failure to exercise appropriate epistemic due diligence in generating a set of potential explanations *easy bad lot cases*. (We will later contrast these with so-called *hard bad lot cases*.)

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 intuitive requirement into IBE. One natural way which I will follow involves broadening the conception of precisely 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 my anger stops it from explicitly coming to mind.

For purposes of spelling out one of the additional requirements on IBE that I think we need, I shall 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 could easily come up with given a little time are easily available to me, while rival explanations requiring extraordinary ingenuity that no-one has yet imagined 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 the best 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 someone need not 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 perfectly 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 a hypothesis

has been *severely tested*.[4] We shall return to this point later, and I will argue that Mayo is in fact making an even stronger claim than this.) Perhaps this sometimes involves some sort of induction from the fact that so many ingenious attempts over such a long period of time to come up with a rival to some explanatory hypothesis H have failed, to the conclusion that we are justified in believing that there is no easily available rival to H superior to it. The details of how such an induction might work are not essential for this paper. For our purposes, it will simply 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. My suggestion is that this is what we should regard as the starting point of an IBE.

Armed with this, I propose the following initial refinement of IBE:

$$\frac{\text{I am justified in believing that } H \text{ is the best of all easily available explanations for } P.}{\text{I am justified in believing that } H \text{ is true.}} \qquad \text{IBE}_3$$

(Again, this will not be our final version of IBE, but it nevertheless represents an important improvement over IBE_2.) Easy bad lot cases are not counterexamples to IBE_3, because in easy bad lot cases there is an easily available, highly plausible explanation being ignored. In such cases, we have not performed due diligence in generating a set of rival explanations to H, and thus cannot be said to be justified in believing that H is the best of all easily available explanations. The premise of IBE_3 is false in such cases. Such cases are therefore not counterexamples to the deductive validity of IBE_3.

5 Hard Bad Lots

Eliminating the threat of easy bad lots does not completely get rid of the bad lot problem. For there are cases in which through meticulous investigation, I am quite aware of all (or enough) of the easily available explanations that I can form a justified belief that some H is the best of all easily available explanations, and yet I am still not justified in believing H itself. In such a case the premise of IBE_3 would be true (unlike easy bad lot cases), but its conclusion false. Call such a case a *hard bad lot*. In hard bad lot cases, I am simply not justified in believing that the truth lies among the hypotheses easily available to me, even though I have a firm grip on the set of easily available explanations.

Hard bad lot cases often occur 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, (Dellsen [2021]) points out that around the 1820s and 1830s, even though it was conceded that Fresnel's transverse wave theory of light provided the best available explanation

[4]See in particular the discussion of Perrin and Brownian motion in (Mayo [1996], chapter 7).

for various optical phenomena, physicists did not take themselves to be justified in believing it.[5] One problem was that there were results Fresnel's theory was unable to explain. For example, (Brewster [1883]) raised challenging questions about the nature of the 'ether' through which light undulated, as well as questions about how Fresnel's theory was able to account for the specific way in which chromium, potassium, gold and various other substances interacted with light. Such worries left scientists unable to fully endorse Fresnel's theory, even while endorsing its extraordinary success. For a different example, (Novick and Scholl [2022]) and (Novick [2017]) argue 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. For a further example, string theory might currently be the best easily available explanation of various phenomena of quantum gravity (as well as the best easily available explanation of the diversity of particles), though it is hard to imagine arguing that we are justified in believing it solely in virtue of this fact given the extraordinary claims string theory makes about the dimensionality of space, as well as it resistance to empirical testing.

In addition to these examples, it will also be helpful to have a simpler and cleaner example of a 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 of 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 would presumably be premature to judge that H_0 is the best of all easily available explanations of my failure to spot Bigfoot. 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

[5]In particular, (Dellsen [2021], p. 163) 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.

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 of my failure to spot Bigfoot, and test them all accordingly. Assuming that all these searches also fail to yield a sighting, it will at some point become reasonable for me to believe that H_0 is the best easily available explanation of my failure to spot Bigfoot. (Perhaps I do not even have to test literally every rival explanation to reach this conclusion.)

In more detail, consider something like the following rough list of potential minimally plausible 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 skillfully able to 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.
$\vdots$

At this point in my investigation as I have described it, it looks very reasonable for me to think of H_0 as better than any other explanation 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 calls for explanation – exactly *how*, after all, does Bigfoot continue to avoid contact with me, if he is anything like the other creatures in the forest? The fact that H_1 calls for 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 at least somewhat implausible, even if it still meets our standards for minimally plausibility. And so I am justified in believing that H_0 is the best easily available explanation of my failure to spot Bigfoot.

However, the key thing to note is that all this is compatible with my acknowledging that there is still an enormous amount I do not know about the forest and its wildlife. 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 overall 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. I simply know too little about the forest to have that sort of

justified belief. For example, I might in the future discover that there are animals in the forest with an extraordinary ability 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. We can assume that I do not presently have sufficient evidential grounds on which to dismiss such possibilities. As of now, H_0 is the best easily available explanation I have of my failure to spot Bigfoot. But given how little I know about the forest, I am certainly *not* justified in thinking that H_0 will remain the best easily available explanation of my failure to spot Bigfoot as more evidence rolls in. In sum, we have another example of a hard bad lot: I am certainly justified in believing that H_0 is the best of all currently easily available explanations of my inability to spot Bigfoot, but I am nevertheless not yet justified in believing H_0 itself.

Before looking at such examples more closely, it is worth pointing out that there are some people who outright deny the existence of bad lots altogether. For example, (Shaffer [2019, 2021]) argues that we are right '*to select the best explanation of a phenomenon from among known hypothesis on the basis of known evidence*' (Shaffer [2019], p. 427). Shaffer emphasizes, however, that our acceptance of the best available explanation should be understood as defeasible, and that the scientist is under an obligation to continue probing the relevant phenomena to see if there might be even better explanations. This is the best we can do, and to demand more of theory acceptance is to '*[conflate] ideal standards of rationality with epistemic standards of rationality*' (Shaffer [2021], p. 268). In this way, Shaffer takes himself to have refuted the bad lot argument. Shaffer would say that whenever I am justified in believing that H is the best of all easily available explanations, then I am justified in (defeasibly) believing H. To use the terminology of this paper, Shaffer would therefore certainly deny the existence of what I have called hard bad lots.[6]

Shaffer is of course right to point out that the best of all currently easily available explanations occupies a special status in the process of scientific inquiry. Paying particular attention to the best of all currently easily available explanations and probing it as deeply as possible is a good strategy for arriving at even better explanations, as the history of science clearly attests. Perhaps there is some sense in which we are then right to 'accept' the best of all currently easily available explanations. For example, even Brewster seemed willing to take something like this attitude towards Fresnel's wave theory of light.

The problem however is that this notion of theory acceptance does not amount to *belief*, which is the thing with which the bad lot argument is concerned. The question being probed by van Frassen in the bad lot argument is whether something being the best explanation available to us justifies our believing it. Brewster might have been right to pay special heed to Fresnel's theory, and we might be right to pay similar heed

[6]Perhaps Shaffer would also say something similar about easy bad lots, but I will not pursue this point here.

to string theory, and we might also be right in the scenario described to pay special heed to the hypothesis that Bigfoot does not exist. Shaffer is undeniably right about all this. But this does not mean that Brewster would have been justified in *believing* the wave theory, or that we would be justified in believing string theory, or that in the scenario described we would be justified in believing in Bigfoot's non-existence. In each of these situations there is simply not enough evidence to justify such belief. And so the kinds of considerations Shaffer raises, while telling us something useful about good methods for scientific inquiry, do not seem to me to entail the non-existence of hard bad lots.

6 More on Hard Bad Lots

Let us examine our examples of hard bad lots more closely, and try to discern their most general common features. We find something noteworthy by comparing the Bigfoot example with that of Fresnel's wave theory of light. In the case of Fresnel's wave theory, scientists like Brewster took themselves to have positive reasons for thinking that there were better explanations of the optical phenomena purportedly explained by the wave theory than those offered by the wave theory, even if such better explanations were not yet easily available. In particular, Brewster seemed to think that there must exist explanations that got right the things that Fresnel's theory got right, while also getting right the things that Fresnel's theory appeared to get wrong. In this way, it was reasonable for Brewster to think that the easily available explanations of the time did not include the correct explanation, and that scientific inquiry would eventually uncover an explanation even better than Fresnel's wave theory.

By contrast, in the case of the Bigfoot example we have no particularly strong positive reason for thinking that there are better explanations for our observations than the hypothesis that Bigfoot does not exist. The problem is rather that our general lack of knowledge means that we have no good reason one way or another for thinking that the hypothesis that Bigfoot does not exist will remain the best explanation as our inquiry proceeds. The Bigfoot example therefore shows that it is not just when we have good reason to think that scientific inquiry will eventually uncover an even better explanation that we are not justified in inferring the truth of the currently best easily available explanation. Rather, merely lacking grounds for thinking that the currently best easily available explanation for some phenomenon will remain the best easily available explanation as inquiry proceeds means that we are not justified in inferring the truth of the currently best easily available explanation. The lack of justified belief that our currently best easily available explanation will remain the currently best explanation as inquiry proceeds is in fact the most general commonality of all hard bad lot examples I have considered, and I know of no example of a hard bad lot that lacks this feature.

Let us turn now to cases in which we *are* justified in inferring the best easily

available explanation. For example, consider the explanation of the orbits of the planets in terms of the inverse square law of gravitation, or the explanation of the way in which species change over time in terms of Darwin's theory of natural selection, or the explanation of the flow of heat in terms of the second law of thermodynamics. To focus on the first case, the extraordinary success of the inverse square law of gravitation in explaining the orbits of the planets, where this includes very detailed and risky predictions (such as the existence of Neptune) that turn out to be true, convinces us that the existence of a rival explanation that does as well is extremely unlikely. It is important to emphasize that what such successes convince us of is not merely that it is very unlikely that there is no *easily available* rival explanation that does as well, but rather that that it is very unlikely that there is *any* rival explanation – easily available or otherwise – that does as well.

It is useful to connect this with Mayo's notion of *severe testing*. According to Mayo, a severe test T of a hypothesis H is one such that 'there is a very high probability that the test procedure T would not yield ... a passing result, if H is false.' (Mayo [1996], p. 180.) A severe test thus renders rival hypotheses or explanations improbable. Mayo is clear that is it not just *known* rival hypotheses that are rendered improbable, but hypotheses such that 'we may not even know what they are' (Mayo [1996], p. 195.) Mayo goes on to argue

> I can rule out the killer's being over six-feet tall without scrutinizing all six-footers. A single test may allow ruling out all six-footers. Using a similar strategy Jean Perrin was able to rule out, as causes of Brownian motion, all factors outside a certain liquid medium. He did so by arguing that if the observed Brownian motion were due to such external factors – whatever they might be – the motion of Brownian particles would follow a specified coordinated pattern. His experimental tests, Perrin argued, would almost surely have detected such a pattern of coordination, were it to exist; but only uncoordinated motion was found. (Mayo [1996], p. 196.)

Severe tests thus render not just known rivals unlikely, but *all* rivals unlikely. Perrin's experiments rendered unlikely not just known ways of trying to account for Brownian motion in terms of external factors, but *all* such ways, 'whatever they might be'. In (Mayo [1996]), Mayo has much to say about what constitutes a severe test, and much to say in response to those who doubt the cogency of the very idea of a severe test, but the details of all this are not our main concern. The important point is that there is precedent in the philosophy of science for the claim that we can sometimes come to believe with justification that an explanation is better than not just any known (or easily accessible) rival, but in fact better than even rivals that are not known, or not easily accessible. I think we find precisely this in the explanation of the orbits of the planets in terms of the inverse square law of gravitation, the explanation of the way in which species change over time in terms of Darwin's theory of natural selection,

the explanation of the flow of heat in terms of the second law of thermodynamics, or the explanation of Perrin's experimental results in terms of Brownian motion. In each case, we are justified in thinking not just that the explanation in question is the best easily available explanation, but rather that the explanation in question will continue to be the best easily available explanation, even as our scientific investigations continue, more data is gathered, and hitherto unimagined putative explanations are brought to our attention.

Reflecting on examples in which we are justified in inferring the best easily available explanation as well as examples in which we are not, we naturally arrive at the conclusion that it is precisely when we are justified in thinking that the best easily available explanation H would remain the best easily available explanation even if scientific inquiry were to freely continue that we are justified in inferring H. In such a case, we are entitled to move from a merely comparative judgment about H (that H is superior to its rivals) to a non-comparative judgment about H (that the evidence warrants belief in H.) That in this sort of situation this sort of inference is possible should not be surprising. Justified belief, after all, 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) grounds for thinking that further evidence will likely not change this verdict. Without this, it is difficult to see what would entitle us to render a verdict, and thus difficult to see how we could have justified belief at all. In different language, justified belief only arises after the activation of some sort of appropriately chosen 'stopping rule'. Once we have grounds for thinking that the best easily available explanation H would remain the best easily available explanation even if scientific inquiry were to freely continue, there is then some sense in which further evidence gathering is not necessary, our inquiry is complete for the purposes of justified belief formation, and we are right to make the non-comparative judgment that the evidence warrants belief in H.

Of course, in such cases the belief that the given explanation would remain the best even if scientific inquiry were to freely continue is highly defeasible. Further evidence could erode our confidence in this belief. This sort of defeasibility is of course present in all non-deductively acquired justified belief. This, however, does not render unjustified the belief that the explanation in question would remain the best even if scientific inquiry were to freely continue, and so does not deprive us of the possibility of performing an IBE. In some sense then, even though some aspect of inquiry closes upon formation of justified belief, we must maintain readiness to re-open inquiry should further evidence compel us to.

7 Formulating IBE correctly

With the argument of the previous section in mind, I propose the following final version of IBE:

I am justified in believing that both (i) H is the best of all easily available explanations for P, and (ii) H would remain the best of all easily available explanations for P were scientific inquiry to freely continue.	IBE_4
I am justified in believing that H is true.	

In easy bad lot cases, we are not justified in believing that the hypothesis H we would like to infer is the best of all easily available explanations, and in hard bad lot cases, we are not justified in believing that the best of all currently easily available explanations would remain the best of all easily available explanations were scientific inquiry to freely continue. Thus, in all bad lot cases the premise of IBE_4 is false. Bad lot cases are therefore not counterexamples to the claim that IBE_4 is deductively valid. In this way, IBE_4 is immune to bad lots. Three sets of remarks will help to clarify the view here and ward off misunderstandings.

(1) First of all, it must be re-emphasized that the premise of IBE_4 involves not just the judgment that as scientific inquiry freely proceeds, some explanatory hypothesis H will remain the best of our set S of *currently* easily available explanatory hypotheses, but rather the stronger judgment that H will be the best of our *future* sets of easily available explanatory hypotheses, whatever those might be.

Now, one might worry that some sort of pessimistic meta-induction tells us that we can never have confidence in any explanatory hypothesis remaining the best of all easily available alternatives in this way. But pessimism of this sort does not seem to me to be warranted by an examination of the history of science. Even as our best scientific theories evolve, it is not unusual for the correctness of old explanations to persist. For example, the Newtonian explanation of the motions of the planets is still the correct explanation, even if now we understand its correctness a little differently – namely, as a consequence of the fact that general relativity yields classical mechanics in certain limits. Arguably, there is reason to think that the physics of the future will look very different from the physics of the present, but that does not mean that all our current explanations will be abandoned (though surely some of them will.) I do not see grounds – historical or otherwise – for holding a kind of skeptical position according to which we can *never* have confidence that an explanatory hypothesis will remain the best even as scientific inquiry freely proceeds. In fact, if Mayo is right, then surviving severe testing gives us good (though defeasible) reason to think that a hypothesis *will* remain the best as scientific inquiry freely proceeds. And so even though explanatory hypotheses can sometimes be supplanted by new explanatory hypotheses when further data is acquired, it does not follow that we can never be (defeasibly) justified in believing that this will not happen. The skeptical worry that

we can never have the kind of justified belief required by the premise of IBE_4 therefore strikes me as unfounded.

(2) Secondly, note the presence of the requirement that 'scientific inquiry freely continue' in clause (ii) of IBE_4, and the way in which this turns (ii) into a counterfactual conditional. For our purposes, we work with a semantics of counterfactuals according to which 'if X were true, then Y would be true' holds in the actual world just in case the closest $X\&Y$ worlds are closer to the actual world than the closest $X\&\neg Y$ worlds. This basic idea goes back to (Lewis [1973]) and can be fleshed out in a multitude of ways. There is now an enormous literature devoted to figuring out the details of this as well as rival views on the semantics of counterfactuals, but as this will not really matter for anything I discuss it will suffice to proceed somewhat naively.

To see the benefit of formulating (ii) in this way, consider a case in which an anti-science dictator arises who wants to strongly restrict scientific activity. We may well then have grounds for believing that our currently best available explanation for some phenomenon will remain our currently best available explanation for that phenomenon in the future, as perhaps we have reason to think that the creation of new plausible hypotheses will require the use of experimental resources that the dictator has prohibited. However, if the closest worlds in which inquiry is allowed to freely continue are worlds in which the dictator did not seize power, then clause (ii) may well turn out *not* to be true. This is as it should be, as artificially stifling future inquiry should not be enough to render belief in a hypothesis justified.

Of course, there are cases in which inquiry is stifled not by dictators, but by natural phenomena. In our Bigfoot example, let us suppose that due to natural causes our forest suddenly becomes completely inaccessible to us. (An extreme scenario could be one in which the forest is completely destroyed by an asteroid.) Unlike the dictator case, we are still as free as ever to continue scientific inquiry. There are however now certain sorts of data that we simply cannot collect. In virtue of this, perhaps it is now easier for (ii) to be true. This is because whereas before we might have been unsure whether further investigation of the forest might reveal surprising phenomena supporting rival or new explanations of our inability to spot Bigfoot, now that the forest is inaccessible, we might well be very confident that even in the future the non-existence of Bigfoot will remain our best explanation of our former inability to spot Bigfoot. Thus, according to IBE_4 the sudden inaccessibility of the forest has made it *easier* to be justified in believing in Bigfoot's non-existence.

This should not be surprising and in my view is not a cause for concern. The picture of inquiry with which I am working is one according to which we sometimes find ourselves in a situation in which we have evidence favoring a proposition X over its negation, and good grounds for thinking that further inquiry will likely not change this verdict. In such a case inquiry reaches what we may think of as a natural stopping point, and we are justified in believing X.[7] In the case in which certain methods of

[7]This is obviously not intended to be a reductive analysis of justification. Nor is it supposed to

inquiry become impossible, it is then easier to reach such a stopping point. Conversely, in cases in which new methods or tools of inquiry become available to us (for example, if the forest or records of the forest unexpectedly become accessible to us again) what was once a natural stopping point in inquiry may no longer be a natural stopping point, and further investigation using these new methods or tools may be needed before we can come to be justified in believing a hypothesis.

Putting together these last few points, we then see that human restriction of inquiry does not in general lower the standard for justified belief, though restrictions of inquiry by natural causes can. By contrast, the creation of new methods of inquiry (from either human or natural causes) can have the effect of raising the standard for justified belief. Given the conception of justified belief with which we work here, these results strike me as intuitive.

It should be noted that there will also be cases in which it is simply not clear – and perhaps even indeterminate – whether (ii) is true. For example, suppose the only way we can advance inquiry in some area of physics is to build a prohibitively expensive piece of equipment that would deplete a non-trivial amount of the earth's resources. Would the political decision to not build this piece of equipment count as a restriction of free scientific inquiry? If the machine only required a trivial amount of the earth's resources, then there are at least cases in which a political decision to not build it *would* count as a restriction on free scientific inquiry. If however the machine required a massive proportion of the earth's resources, then the decision to not build it should arguably *not* be viewed as restriction on free scientific inquiry, but rather as something closer to a physical or natural constraint on inquiry. The case I want to consider lies in the zone between these two extremes where it is simply unclear or even indeterminate whether we want to say that not building the machine counts as a restriction of free scientific inquiry. Whether there is any fact of the matter as to the truth value of (ii) in such cases is far from clear. Is strikes me as unsurprising that there are awkward situations like this where it is unclear and perhaps even indeterminate whether inquiry has reached a stopping point, and thus unclear or even indeterminate whether we are justified in believing the best of our currently easily available explanations for some phenomenon. The existence of such borderline cases should not be counted as an argument against IBE_4 as a formulation of IBE. In fact, if we view justification as a vague predicate we should instead be entirely unsurprised by the existence of borderline cases in which there is perhaps not even a fact of the matter as to whether belief in a claim is justified.

(3) Finally, we return to the question of the necessity of additional requirements on IBE discussed in §3 - in particular, the requirement that the best explanation be minimally plausible, and the requirement that it be significantly better than any competitor. These additional requirements seem to me to be no longer necessary

exclude the possibility that we decide to continue inquiry for whatever reason – we reach a stopping point only in the sense that our search for evidence that would warrant belief one way or another has been successful.

when IBE is presented in the form IBE_4. The requirement of minimal plausibility is supposed to make sure that we do not rashly infer something implausible in cases in which more plausible alternatives are unavailable. This is something we must worry about when we are at the very beginning of the process of inquiry and both our evidence and our set of candidate explanations are lacking. However, by the time we have reached the point of inquiry at which we can form justified beliefs about which explanations will persist as the best explanations, we are in a very different situation. Perhaps, for example, we must have a hypothesis that has passed one of Mayo's 'severe tests' to reach this point. In such a case, the worry that our currently best easily available explanation might nevertheless fail to be minimally plausible seems to me to be unfounded. There is therefore no reason to add a 'minimal plausibility' condition to IBE_4.

In the case of the requirement that our best easily available explanation be significantly better than any competitor, the worry that this requirement is supposed to address is that if a hypothesis is only marginally better than a rival, we cannot be confident that further evidence might not tip the scales in favor of what is now the second-best explanation. But the requirement that we have justified belief that the currently best easily available explanation will remain the best easily available explanation eliminates this worry. When the currently second-best easily available explanation is only marginally worse than the best easily available explanation, the requirement that we be justified in thinking that the currently best easily available explanation will maintain its status as the best explanation as scientific inquiry proceeds is of course a very demanding requirement. But as long as it is met, as IBE_4 requires, there is no need for an additional requirement that the best easily available explanation be significantly better than any competitor.

As a result, I claim that IBE_4 stands in no further need of additional requirements, and that unlike other versions of IBE may be taken as a complete statement of IBE.

8 Concluding Remarks

The intuition that there are bad lot cases refuting the validity of IBE arises from a bad 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 at some moment, assign each 'points' for simplicity, explanatory power, prior plausibility and so on, tally up these points in some way or other, and then see which scores highest. (In Table 1 for example, hypothesis H_2 scores highest with 9.4 points.) In this way we find out which hypothesis H is the best explanation. According to this image of IBE, if we perform this process sufficiently carefully, we are then in a position to infer H.

What this description of IBE conceals is the fact that to perform an IBE, we must not just be justified in thinking of some H that it is the best of all currently easily available explanations, but also that this would continue to be the case were

	Simplicity	**Explanatory Power**	...	**Total**
H_1	1.3	0.5		4.3
H_2	7.1	0.1		9.4 ✓
H_3	3.2	2.2		7.2
...				

Table 1: an example of the naive model of IBE

scientific inquiry to continue (as it almost certainly will), and new evidence acquired and new explanations imagined. Only then are we justified in inferring H. At any moment of time we can certainly take a 'snapshot' of how H fares relative to its rivals at that point, but this on its own does not allow us to perform an IBE until we are capable of forming a justified belief that H's occupying the first position will persist in a sufficiently robust counterfactual sense. As a result, we can only perform an IBE at a relatively late stage in the process of inquiry, when we have a good grasp of the general contours of the set of possible explanations for a phenomenon, and have substantive justified beliefs about this set. This suffices for blocking the possibility of bad lots, as I have argued.

To show that IBE is not refuted by bad lot examples is however not to have shown that IBE is free from all problems. Perhaps IBE faces other challenges. For example, one concern is whether it is possible to have enough evidence to believe the premise of IBE_4 without *already* having grounds for believing H independently of IBE-type considerations. One might think for instance that H surviving a severe test already gives us direct grounds for believing H, without needing to argue indirectly that H surviving a severe test gives us grounds for believing the premise of IBE_4, which in turn gives us grounds for believing H. More generally, might one not be able to infer H directly from whatever evidence it is that warrants belief in the premise of IBE_4, without passing through IBE? In (Davey [2023]), I have recently argued that in many cases, explanations may be directly inferred from evidence without the need for principles such as IBE. If anything in any of these lines of argument turn out to be right, then IBE is in some sense *redundant*, insofar as it only allows us to infer things we could already have inferred without it. Call this the *redundancy objection.*

The redundancy objection might not seem particularly compelling if one sets a very low bar for what is required for an IBE. For example, suppose that as described above, merely taking a 'snapshot' of how H fares relative to its rivals at one moment of time and calculating that it is the highest scoring explanation is enough to allow us to infer H. Then the redundancy objection does not have much force, as it is difficult to see how the fact that H scores better than its rivals at one moment of time gives us enough evidence to infer H without IBE. But to set the bar for IBE so low in this way is to work with a poor formulation of IBE. I have argued that we require much

more evidence to perform an IBE - in particular, we need sufficient grounds to have the justified beliefs described in the premise of IBE_4. But given the larger body of evidence that is now required for performing an IBE, it is no longer obvious that we do not already have the means for inferring H without IBE. Thus the redundancy worry arises.

It will not be the job of this paper to assess whether this objection is fatal to IBE or not. I simply note that when IBE is properly understood, the set of objections to it have a very different character than has generally been supposed. Whether IBE can withstand these sorts of objections is perhaps an open question. However, if IBE fails it will not be because of anything like van Fraassen's bad lot worry.

9 Acknowledgements

Thanks to the anonymous referees, whose comments greatly improved this paper.

References

Brewster, D. [1833]: 'Observations on the absorption of specific rays, in reference to the undulatory theory of light', *Philosophical Magazine*, **2**, pp. 360-363.

Davey, K., [2023]: 'On Inferring Explanations and Inference to the Best Explanation', *Episteme* (2023), pp. 1-18.

Dellsen, F., [2017]: 'Reactionary responses to the Bad Lot Objection', *Studies in History and the Philosophy of Science*, **61**, pp. 32-40.

Dellsen, F., [2021]: 'Explanatory Consolidation: From Best to Good Enough', *Philosophy and Phenomenological Research*, **103**, pp.157-177.

Douven, I., [2022]: *The Art of Abduction*, MIT press.

Kincaid, H. and Day, T. [1994]: 'Putting inference to the best explanation in its place', *Synthese* **98(2)**, pp. 271-295.

Ladyman, J., Douven, I., Horsten, L., and van Fraassen, B., [1997] 'A Defence of van Fraassen's Critique of Abductive Inference: Reply to Psillos', *The Philosophical Quarterly*, **47(188)**, pp. 305-321.

Lipton, P. [2004]: *Inference to the Best Explanation*, 2nd edition, Routledge.

Lewis, D. [1973]: *Counterfactuals*, Harvard University Press.

Mayo, D. [1996]: *Error and the Growth of Experimental Knowledge*, University of Chicago Press.

Novick, A. [2017]: 'Metaphysics and the Vera Causa Ideal: The Nun's Priest's Tale', *Erkenntnis*, **82**, pp. 1161-1176.

Novick, A. and Scholl, R. [2022]: 'Presume It Not: True Causes in the Search for the Basis of Heredity', *Brit. J. Phil. Sci*, **71**, pp. 59–86.

Psillos, S., [1996]: 'On Van Fraassen's Critique of Abductive Reasoning', *The Philosophical Quarterly*, **46(182)**, pp. 31-47.

Schupbach, J. [2014]: 'Is the Bad Lot Objection Just Misguided?', *Erkenntnis*, **79 (1)**, pp. 55-64.

Shaffer, M., [2019]: 'The Availability Heuristic and Inference to the Best Explanation', *Logos and Episteme*, **X(4)**, pp. 409-432.

Shaffer, M., [2021]: 'Van Fraassen's Best of a Bad Lot Objection, IBE and Rationality', *Logique and Analyse*, **255**, pp. 267-273.

van Fraassen, B. [1980]: *The Scientific Image*, Oxford University Press.

van Fraassen, B. [1989]: *Laws and Symmetry*, Oxford University Press.