Episteme (2024), **21**, 1120–1137
doi:10.1017/epi.2023.19

episteme

ARTICLE

On Inferring Explanations and Inference to the Best Explanation

Kevin Davey

Department of Philosophy, University of Chicago, Chicago, IL, USA
Email: kjdavey@uchicago.edu

(Received 5 September 2022; revised 6 February 2023; accepted 20 February 2023; First published online 05 May 2023)

Abstract

Although the inferring of explanations plays an important role in both our everyday lives and in the workings of science, I argue that inference to the best explanation as it is commonly conceived is often not the best way to capture this sort of reasoning. I suggest that a different form of reasoning – so-called *immediate explanatory inference* – is instead often much better suited to this task. This is a form of inference in which we are justified in believing explanations for the evidence before us purely in virtue of this evidence, and not in virtue of the evidence plus some general principle or rule of non-deductive reasoning. I defend the idea of such a notion of inference, and argue that it plays a central role in both ordinary life and science.

Keywords: Explanation; non-deductive inference; inference to the best explanation; Peirce

1. The Problem

One thing we often do in our non-deductive reasoning is infer explanations for known phenomena. Typically, inference to the best explanation (henceforth, IBE) has been regarded as one of the best ways to spell out how this is possible.[1] And yet, attempts to spell out the details of how exactly IBE works immediately befall a large set of intractable problems. For one, the question of what counts as an explanation admits no easy answer.[2] The question of what makes one potential explanation 'the best' of some set of potential explanations is no less thorny.[3]

Perhaps of more serious concern is the following dilemma. If performing an IBE involves judging that some explanation E is better than *any* other potential explanation for some phenomenon, then this means that it involves judging that E is superior to rivals that we have not even thought of yet. How any such judgment could be justified is not clear. (This issue has been carefully discussed in

[1]Perhaps the best case for this is made in Lipton (2004).

[2]A detailed historical account of theories of explanation is provided in Salmon (2006).

[3]For a noteworthy attempt at tacking this question, see Thagard (1978). Here, Thagard tries to account for what makes one explanation better than another in terms of Whewell's notion of *consilience*.

Stanford 2006.) But if on the other hand IBE only involves the judgment that *E* is better than any rival we can currently think of, then it is difficult to see how this gives us good reason to believe *E*. Surely, after all, there will be cases in which none of the explanations we can think of are particularly good, and *E* being the best of such a set of explanations hardly seems to render belief in *E* justified. This is known as the 'bad lot' argument, and was first formulated explicitly in van Fraassen (1989). It has been further developed and defended in papers such as Ladyman *et al.* (1997) and Dellsen (2017).[4]

Dilemmas like this have led some to give up entirely on the idea of IBE.[5] This paper proceeds under the assumption that what is needed is not a complete rejection of IBE, but rather a significant reformulation that somehow avoids the dilemma in the first place. In the terminology of Dellsen (2017), what is needed is a *revisionary* rather than a *reactionary* response to the puzzles surrounding IBE.

To this end, I argue in this paper that there is a way of inferring explanations that, while more modest than IBE in its original form, adequately captures many of our actual non-deductive inferential practices while entirely avoiding the above dilemma. I shall call this reformulation of IBE *immediate explanatory inference*. I do not claim that every case of inferring an explanation is a case of immediate explanatory inference, but I think that a surprisingly large number of cases are. The task of this paper is thus to defend the idea behind this type of inference, and argue for its centrality in both science and our everyday lives.

In §2, I present several examples of the inferring of explanations, and argue that these examples are better understood as cases of immediate explanatory inference rather than IBE. In §3 I then defend the idea of immediate explanatory inference from a number of attacks. In §4, I argue more broadly that many traditional examples of IBE both in and outside of science are nothing but immediate explanatory inferences. I also argue that in at least some cases in which a putative IBE does not seem to be naturally describable as an immediate explanatory inference, what we have is in fact not an inference at all, but something more like a reasonable conjecture. I contrast immediate explanatory inference with Peircean abduction/creative abduction in §5, and make some concluding remarks in §6.

2. Immediate Explanatory Inference

2.1. Basic examples

Consider the following situations.

> **Dog Example**: I come home and see that my favorite houseplant has been torn from its roots, and is lying in pieces in a filthy mess on the floor. I also see my dog's pawprints surrounding it, and my dog cowering in the corner of the room, head downcast and eyes pointing nervously upwards, sadly whimpering. His nose and paws are covered in soil, and there are no other suspicious circumstances. I infer that the dog tore up the plant.

> **Torn Page Example**: I find a small irregularly shaped scrap of paper on my desk. I then notice that a book on my desk has the corner of one of its pages missing. The scrap of paper fits the missing corner perfectly in color, shape and all other

[4]For responses, see Psillos (1996) and Schupbach (2014).
[5]See Norton (2021), for example.

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

> discernible characteristics. I infer that the scrap of paper was somehow torn from this very page.[6]

In both of these examples, we have a state of affairs that calls for explanation. Why is the plant in this state? Why does the scrap of paper fit the book perfectly? In each case, an explanation for the state of affairs is not only offered but inferred.

What is striking in both of these examples is the force and immediacy with which the inferred explanation jumps out at us. Given the evidence, the claim that the dog tore up the plant jumps out at us with such strength that we find it impossible to resist. Likewise for the claim that the scrap of paper was somehow torn from the page in question.

We do not find it necessary in such cases to carefully weigh the epistemic virtues of rival hypotheses to determine where to invest our credences. Nor do we find ourselves worrying about what *really* constitutes an explanation, worrying about what makes one explanation better than another, or worrying about whether we might be choosing from a bad lot. We do not even judge that there is only one reasonable explanation, and then on the basis of this infer that explanation. In these examples we simply look at the evidence, and then find ourselves immediately and irresistibly inferring the obvious explanation more or less without any explicit deliberation whatsoever. Moreover, I claim that we are *right* in such cases to infer the obvious explanation on the basis of this evidence, without deliberating about what the relative virtues of any rival explanations might or might not be. There are presumably many situations in which explicit, sustained deliberation of some sort is needed to infer an explanation, but neither of our above examples describe such situations. In this paper, much of what follows is a consequence of trying to take this aspect of our examples as seriously as possible.

Our day to day lives are filled with examples like these. We infer the recent presence of people on a beach from footprints in the sand. We infer the existence of cars outside from the sounds of traffic. We infer from the smell of tobacco that someone is or has recently been smoking nearby. We infer that a friend sincerely believes something from the fact that they assert it in a serious manner. All these inferences are of course defeasible. For instance, if we sometime later learn that a burglar is running rampant in our neighborhood whose modus operandi involves uprooting house plants and framing dogs, we would need to expand our set of potential explanations and search for further evidence to settle the question of the dog's guilt. But absent good reason to take such possibilities seriously, we are justified in immediately inferring the overwhelmingly compelling explanation that the dog is guilty, without consideration of alternatives, or any process involving the weighing of rival explanations.

In these examples, against the backdrop of our background knowledge and experience, I would like to suggest that the first-order evidence (the houseplant torn from its roots, the dog's pawprints surrounding it, etc.) immediately justifies belief in the explanation that the dog tore up the plant. I shall call the inference that occurs in these sorts of examples *immediate explanatory inference*. Although some effort will be required to specify exactly what is happening in such cases, I take the basic phenomenon to be perspicuous enough that a discussion may begin here.

Perhaps the idea that we perform immediate explanatory inferences might initially seem appealing, or even obvious. Nevertheless, it is very difficult for the philosopher to resist the pull of the idea that if things like the dog and torn page examples are

[6]This is adapted from an example from Peirce (1878: 475).

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

good explanatory inferences, they must be so in virtue of some principle or rule with which they conform and with which we reason. As soon as one tries to spell out what such a principle or rule must look like complexities rapidly appear, and what initially looked like straightforward inferences turn out to hide all sorts of intricacies which must be carefully thought through before we can infer the conclusions we want. This robs our examples of their immediacy. The goal of much of this paper is to return us to our former innocence, and argue that inferences like these really are, in a certain sense, immediate. Much of §3 of the paper is devoted to combatting our impulse to search for 'principles' that license these inferences. In §4 of this paper, I argue that even in the context of careful scientific inquiry, such impulses can be resisted.

Although in this paper I will not need to choose a specific account of explanation from the bewildering set of possibilities that exist today, some very general remarks on the way I view explanation will perhaps help to avoid confusion in what follows. Very loosely speaking, accounts of explanation are divided between those that claim that explanation always involves subsumption under *laws*, and those that think it is enough to cite a set of relevant or salient facts in the causal history of the event to be explained, without the need for (necessarily) citing laws. This basic schism emerges in the earliest contemporary philosophical discussions on explanation (especially in the papers of Hempel 1942; Hempel and Oppenheim 1948; Scriven 1959, 1962), and is still an important schism in the explanation literature today. For the purposes of this paper, I locate myself on the latter side of the schism. Thus, the dog tearing up the plant is the (salient) cause of the mess before me, and the scrap of paper having been torn from the page is the (salient) cause of the 'fit' we see between the page and the scrap, and it is in virtue of this that the facts in question function as explanations, regardless of whether or not we can subsume such events under general laws. There are of course many questions about how the details of such a story work, but addressing such questions will not be the job of this paper.

Understanding explanation in this way also helps us to make sense of the role of background knowledge in our main examples. In our main examples, it is only against the backdrop of our background knowledge that the first-order evidence justifies belief in the explanations in question. It is clear that our background knowledge must be involved in such inferences somehow – without background knowledge about dogs, plants, and what happens when paper is torn, surely none of the inferences in our main examples can be justified. (Such background knowledge may or may not contain Hempelian covering laws; though I think typically it need not.) The reason background knowledge is needed is that it is only against the backdrop of background knowledge that we can come to recognize a known or posited part of an object's causal history as particularly salient in explaining some fact. Recognizing something as a salient cause of something else cannot be done in a vacuum. In this way, the inferring of explanations from evidence always demands background knowledge. Although we may sometimes suppress reference to background knowledge in what follows, it should be understood as always present and playing an indispensable role in our explanatory inferences.

2.2. Misunderstandings to avoid

In order to emphasize what is important in our examples, it will help to warn against four possible misunderstandings. First of all, it must not be thought that in immediate explanatory inference it is some sort of overwhelming feeling of conviction that justifies

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

belief in the inferred explanation. Feelings of conviction can occur even when we have little or no reason to believe something, and the responsible inquirer recognizes the need to not be led astray by such feelings. In our examples, it is the first-order evidence (the toppled houseplant, the dog's pawprints, and so on) that justifies belief in the inferred explanation, and not some mere feeling of conviction inside of us.

Second, what is *not* being suggested is a kind of perceptual story in which we 'see' the truth of an explanation, in such a way that the immediacy of the inferences in our main examples reduces in some way to the immediacy of perception. In both main examples, the explanation in question is a claim about the past, and we simply do not see the past in any literal and relevant sense. Although we might say that we *see that* the dog tore up the plant, in this context, 'see that' surely means 'infer that'. We certainly do not *see* the dog tearing up the plant in any sense.[7]

Third, the claim that in our main examples consideration of rival hypotheses plays no role in justifying belief in the relevant explanations should *not* be confused with a merely psychological claim about what does or does not occur in our minds at the moment of inference. A different example will help bring out the importance of this point. When asked what 6×8 is, most people will immediately and without hesitation answer 48, even though what ultimately justifies this answer is a non-trivial calculation or piece of reasoning. The immediacy and simplicity of what goes on in our minds when we answer 48 should *not* be taken to entail, however, that the *justification* of this answer is in some way immediate or simple. With this sort of phenomenon in mind, someone might object that even though in our main examples we immediately and without hesitation infer that the dog tore up the plant, this is perfectly compatible with a complex piece of argumentation being needed to justify this belief – perhaps a piece of reasoning involving detailed comparisons with rival hypotheses.

My central claim however is that immediate explanatory inference does not have this character. The claim that in immediate explanatory inference considerations of rival hypotheses do not contribute in any way to justifying the conclusion is *not* the claim that at the moment we make the inference, thoughts of rival hypotheses do not enter our mind – they may or may not, though typically they do not. Rather, the claim being made is that what justifies belief in the given explanation does not involve consideration of rival hypotheses at all. In such cases, against the backdrop of our background knowledge, what justifies belief in the given explanation is simply the first-order evidence and nothing more, and the justification is immediate, involving no comparisons of any sort. So for example, given my background knowledge it is purely in virtue of seeing what lies before me that I am justified in inferring that the dog tore up the plant.

As a corollary of this, I will urge that it is not in virtue of the evidence before us plus some sort of general principle (be it IBE or otherwise) that we are justified in inferring that the dog tore up the plant. Rather, against the backdrop of our general knowledge, the evidence before us entirely justifies *on its own* our inferring that the dog tore up the plant. This is a claim about the context of justification. Of course, at this point it may not be clear that such justified inferences are possible. In much of what follows, my goal will be to dispel the kinds of prejudices that tempt us into thinking that they could not be.

Fourth, given what has just been said it should be clear that immediate explanatory inference is *not* just a case of IBE in which an explanation E is inferred on the basis that

[7]For an illuminating discussion of the relationship between seeing and seeing that, see Dretske (1973).

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

it seems utterly implausible that any rival explanation could explain some given state of affairs as well as *E*. In such a case, it would be in virtue of some (perhaps justified) belief about rivals to *E* that *E* is inferred. But no matter how obvious or trivial this process is, the fact that it involves comparisons of *E* with its rivals means that it is no longer a case of immediate explanatory inference. In immediate explanatory inference, it is the first-order evidence, and not facts about rival hypotheses, that justifies belief in some explanation *E*. Thus immediate explanatory inference is not just some sort of trivial or limiting case of IBE in anything like the sense in which IBE is normally understood.

3. Defending Immediate Explanatory Inference

3.1. Worries about IEI

Even if it is granted that the dog and torn page examples are cases in which we can infer an explanation *E* without explicit consideration of rival hypotheses, it does not follow that we have established the possibility of *immediate* explanatory inference. For just because we can infer *E* without explicit consideration of rivals does not mean that we can infer *E* in any sort of immediate way. Perhaps complicated considerations are involved in inferring *E*, even if those considerations do not involve explicit comparisons with rival hypotheses.

In fact, an argument against the possibility of anything like immediate explanatory inference might go as follows. Surely there are conditions that must be met in order that immediate explanatory inference be successfully performed. Whatever these conditions are, part of successfully performing immediate explanatory inference must be verifying that these conditions are met. But insofar as verification that these conditions are met is needed, we cannot describe the inferential process in question as *immediate*. And so even in our main examples, explicit deliberation of some sort is needed.

Here is one reply that might be given to this argument. It is not the reply I will ultimately want to endorse, but it will serve as a useful stepping-stone. Note that if the argument just given were correct, it would seem to speak against the possibility of *any* sort of immediate inference. And insofar as we might want to view reasoning as a process in which a sequence of immediate inferences are sequentially assembled into a larger argumentative structure, the very possibility of reasoning would then come under threat. Something must therefore be wrong with this argument. To see what has gone wrong, consider a rule like &-Elimination, according to which from *A*&*B* we may infer *A*. In typical particular cases, this inference is surely immediate if any is. And yet of course there are things to be checked when one applies such a rule – in particular, one must verify that the sentence *A*&*B* really is a conjunction. In a complicated context this might involve the careful counting and matching of parentheses. But this surely does not challenge the idea that in any particular case the inference from a conjunction to a conjunct is immediate. The argument in question takes the non-immediacy of knowing in certain contexts that the conditions are right for acting in accordance with a certain rule as an argument against the immediacy of acting in accordance with the rule; but that is surely a non-sequitur. The non-immediacy of the knowledge that the conditions for acting in accordance with the rule are met does not mean that subsequent action in accordance with the rule involves any further considerations. Such an argument against the immediacy of the inference in question therefore fails.

But there is something unsatisfying about trying to use this reply to bolster the case for immediate explanatory inference, as there seems to be something fundamentally different between & -Elimination and immediate explanatory inference. In the case of

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

& -Elimination, we can state with precision the conditions under which the rule may be applied (namely, the sentence in question should be a conjunction, and the conclusion one of the conjuncts), even if verification of these conditions is sometimes non-trivial. But what are the conditions under which immediate explanatory inference can be performed? It is impossible to state anything more than vague rules of thumb here. This puts us in an awkward position. For the rule of & -Elimination to yield justified belief, I must be justified in thinking that certain appropriate conditions are met. But how can I be justified in thinking that the conditions for immediate explanatory inference are met, when I cannot even clearly state what these conditions are?

More importantly, even if there are conditions for performing an immediate explanatory inference that we must be justified thinking are met in order to infer some explanation, this sort of picture ultimately robs the inference in question of the *immediacy* that we originally wanted. Suppose it is true, for example, that once we have checked that there is only one overwhelmingly plausible explanation that we are then justified in inferring an explanation without further considerations. Given the amount of work that might go in to checking all this – the careful search for alternatives, and the checking of their plausibility – the original intuition that there is something deeply *immediate* about the inference in question would then be completely lost. Our original intuition was that given my background knowledge, there are circumstances in which after surveying the situation before me I am *without further consideration* justified in inferring that the dog tore up the plant. Surely this idea has been lost if it is only after a potentially complex process of checking that this and that condition for some rule of inference are met that we can then 'immediately' infer the conclusion in question.

The difficulty of stating even vague non-circular conditions for applying some sort of rule of immediate explanatory inference, along with the unsatisfactory consequences of actually succeeding in doing so, helps make clear that it is not in virtue of satisfying the conditions for the application of some *rule* that immediate explanatory inference gives us justified belief in an explanation. Inspired by the character of deductive inference, we might have thought that in order for an inference to be justified it must be an instance of some sort of general rule, which will of course have conditions for its application that must be verified. But I think that we must reject this way of thinking about inference, at least for certain sorts of non-deductive reasoning. It is not in virtue of being subsumable under some sort of rule (a rule that philosophers embarrassingly find themselves unable to articulate) that we find ourselves able to infer that the dog tore up the plant. It is rather in simply seeing the first-order evidence before us that belief in the relevant explanation is justified. There is no rule that is additionally needed to magically confer justification on this belief.

The idea that at least some of our non-deductive inferences are not grounded in universal rules is actually not new, and has appeared in recent epistemology and philosophy of science in several different guises. It has perhaps appeared most prominently in the work of some Bayesian epistemologists. For example, van Fraassen in Churchland and Hooker (1985: 279) denies that there is any 'viable inductive method, inductive logic, or set of canons of induction'. This view is also echoed in Okasha (2001), who thinks that it can be used to avoid at least a certain form of Hume's problem of induction.[8] Such Bayesians have usually defended this view by resorting to a radical subjectivism about our choice of prior

[8]Whether such Bayesians are *right* to think that Bayesianism rids us of canons or principles of induction is a separate question, a full discussion of which goes beyond the scope of this paper.

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

probabilities. (Perhaps Okasha (2001) is most explicit about this.) Subjectivism of this sort however seems highly problematic, for reasons outlined in Earman (1992). Taking a different approach, in Norton (2010, 2021), Norton strenuously argues against the idea of universal rules for non-deductive inference, though for Norton this involves critiquing Bayesianism, rather than finding consolation in it. Norton also ultimately replaces universal rules of induction with 'local' rules of induction. In the view of the present paper, however, local rules are just as unnecessary as universal rules, at least in the examples we have considered thus far. So although the kind of rule-free conception of certain types of non-deductive inference that I am urging here agrees in part with the spirit of these prior approaches, I do not wish to follow any of them in detail. It is not in terms of some fact about our psychological credences (in the style of subjective Bayesianism) or local rules (in the style of Norton) that our belief that the dog tore up the plant is to be justified. It is rather in terms of the first-order evidence alone that belief in such explanations is justified.

But with the idea that the inferences in our main examples are not governed by a rule, a new set of worries emerge. What sense can be made of an inference that does not proceed under some rule or other? A natural concern is that without a rule, there is no obvious way to resolve doubts or disputes about whether an inference of the relevant sort is justified. Imagine our dog example, but this time suppose that the only evidence we have is the upturned plant. Does an immediate explanatory inference allow us to infer the dog's guilt in this case too? Or suppose we see a friend who we know to be in need of money talking with someone we know to be very rich. Does an immediate explanatory inference allow us to infer that the friend is asking for money? It is difficult to know how to answer these questions, and one could imagine doubts or disputes about what immediate explanatory inference tells us in these cases. If we have clear rules telling us when and how immediate explanatory inference can be performed, we have a chance of resolving such disputes. But without such rules, it is difficult to see how these disputes could *ever* be resolved. And if such disputes can never be resolved, one might wonder whether immediate explanatory inference is a justified inferential practice at all. Presumably we will also want our non-deductive forms of inference to play a fundamental role in our scientific reasoning. But what role could be played in science by inferences that are not subject to rules and which allow for irresolvable disputes over whether they are correct or not? We must take these sorts of worries very seriously.

3.2. Vagueness

What has just been gestured at is a cluster of questions between which distinctions must be drawn, even if they are interrelated. Consider first the case in which we see a friend who we know to be in need of money talking with someone we know to be very rich. Without clear rules telling us when immediate explanatory inference can be applied, it is difficult to know whether immediate explanatory inference allows us to conclude that the friend is asking for money.

It is indeed difficult to know what can be inferred in such cases. But I do not see anything troublesome about this. We find ourselves in similar situations with many vague concepts. There are colors for which it is hard to say whether they are red or pink, and people of whom it is hard to say whether they are bald. Being justified is also a vague concept. Imagine a scientific statement that requires a mass of evidence (perhaps 1,000,000 pieces of numerical data) before we can be clearly justified in

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

believing it. Would we be justified in believing it if we only had the first 800,000 pieces of data? It would be bizarre to posit a number n such that in moving from n to $n + 1$ pieces of numerical data we suddenly move from unjustified to justified belief. Surely the most reasonable thing to say is that being justified is a vague concept – one can sometimes be *somewhat* justified (or somewhat unjustified) in a belief, given the evidence at hand. Just as there are borderline cases in which we find ourselves at a loss to say whether someone is bald or whether a color patch is red, so too there will be cases in which we find ourselves at a loss to say whether there is enough evidence before us to be justified in inferring a given explanation. Given the vagueness of justification, this should come as no surprise.

In this regard, it should be noted that there is even something vague about the traditional 'rules' of non-deductive reasoning such as enumerative induction. We are typically not justified in inferring a generalization based on 0 or 1 instances, but are sometimes justified in inferring a generalization when presented with 1,000,000 instances. Between these extremes lie many cases in which it is simply hard to say whether or not we have enough instances to warrant belief in a generalization. There are many reasons to worry about the epistemic status of the principle of enumerative induction, but its vagueness is not among them. Nor do I see any reason to be troubled by the same vagueness in immediate explanatory inference.

3.3. Doubt

There is a certain security that comes with an inference being governed by a rule. So long as we are confident that we have met whatever conditions the rule imposes, we can be confident that we are justified in inferring the conclusion given by the rule. (This of course does not mean that the conclusion is true, but just that we are justified in believing it.) We do not have to worry that our inference is the product of error, bias, or delusion.

In the absence of rules governing our inferences, this security is taken from us. In our main example, suppose that looking at the mess before me, I start to wonder whether in fact I really do have enough evidence to infer that the dog tore up the plant. Might my conclusion not be the result of bias against the dog, who has recently irritated me on several occasions, defecating on my favorite rug? With no rule to act as the ultimate arbiter of what it is and is not reasonable to believe, what is to be done when struck by such doubts?

Much here depends on the exact character of the doubts. If the dog has irritated me on several occasions, and I have good reason to think that my interpretation of the evidence could be biased as a result, then once the evidence before me is broadened to include the fact that the hypothesis in question is the judgment of someone irritated with the dog, I may well *not* have warrant to immediately infer the guilt of the dog. (This is not much different from the fact that my headache might in some cases rob me of my warrant for believing the conclusion of a 3 digit multiplication I have correctly performed.) Or if in general I know that I am inclined to be too quick to jump to conclusions about what my dog has done, that too might serve as counter-evidence that undercuts the immediate explanatory inference. On the other hand, the mere abstract possibility that my judgment might be wrong does *not* seem sufficient to undercut an inference made on the basis of strong evidence, lest we fall into skepticism.[9] In some cases then doubt will undercut the inference in question, while in other cases it will not.

[9]Of course, to give a detailed account of *why* the abstract possibility of error does not undercut such an inference would be to solve the problem of skepticism. I am not claiming to be able to do that – all I am

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

Of course, one might wonder *why* some doubts are sufficient to undercut an immediate explanatory inference, while others are not. I do not see much hope in finding some sort of principle or rule that allows us to answer this sort of question in complete generality. But this is compatible with our being able to judge correctly in many cases that we are justified in believing the conclusion of an immediate explanatory inference, in spite of our doubts. So long as this is true, the phenomenon of doubt poses no particular threat to the possibility of immediate explanatory inference.

3.4. Magic

A further lingering worry is that without a clear rule that licenses us inferring an explanation, it is only by some sort of 'magical faculty' that we can judge whether we have enough evidence to infer the conclusion in question. Imagine again the dog case, but with progressively less evidence – suppose the dog's behavior is less damning, the dirt less scattered, and so on. What makes it the case that at a certain point there is simply not enough evidence to infer the guilt of the dog? A rule of inference might be able to tell us when a certain amount of evidence is too little, or another amount of evidence is sufficient. Without such a rule, however, we seem to be left positing some sort of brute ability to distinguish these cases. Such an ability seems magical, much like the mysterious abilities of the chicken sexers. Can we really be content with such a mysterious view of our basic epistemic faculties?

We must be careful in our comparisons with chicken sexers. What mystifies us about the chicken sexers is their ability to reliably determine the sex of a chicken with an apparent lack of evidence. When I judge that the dog is responsible for the mess before me in the situation I originally described, there is by contrast no lack of evidence. So there is nothing mysterious in that particular case – it is simply the evidence before me that justifies my belief. But what about when I refrain from this judgment in a case in which there is insufficient evidence? What holds me back? The worry is that in the absence of a rule, the ability to judge that our evidence is *insufficient* is magical, and no different from chicken sexing.

This worry however is misplaced. Although I am claiming that in certain cases we can simply recognize without further deliberation that the evidence before us is sufficient for performing an immediate explanatory inference, I am *not* claiming that in cases in which the evidence before us is *insufficient* for performing such an inference that we can also recognize this without deliberation. Typically, coming to see that the evidence before us is insufficient for performing an inference requires deliberation – we recognize, for example, that we have insufficient grounds for eliminating a decently probable rival hypothesis, or that some bias is present that must be checked against, or some such similar problem. The situation then is quite unlike that of the chicken sexers. Perhaps there are cases in which the evidence is so clearly lacking that we can recognize this immediately, but this will not be the typical case. The key point is that even if deliberation and consideration of rival hypotheses is involved in coming to see that the evidence is *not* sufficient for inferring a given explanation, it does not follow that deliberation and consideration of rival hypotheses is involved in recognizing that the evidence *is* sufficient for inferring that same explanation.

claiming is that as a matter of fact the abstract possibility of error on its own does not undercut the inference in question.

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

3.5. Disagreement

Yet another issue concerns matters of disagreement. Suppose a telescope seems to reveal a body in the sky where we did not expect one. One scientist might claim that the obvious explanation for this is that there really is a body (of some sort) there. Another scientist might claim that the obvious explanation is that there has been some sort of observational error, and that there is no body there. One can easily imagine both scientists taking their conclusion to be the result of an immediate explanatory inference. Absent rules that tell us precisely (or even vaguely) when an immediate explanatory inference is justified, how could such a dispute be resolved?

In practice, such disputes are often resolved by gathering more evidence. If the scientist who thinks that there has been observational error can produce further compelling evidence that the telescope is defective, or if the scientist who believes that the body is really there can replicate the observation with many different telescopes in different locations, the dispute will typically be regarded as resolved. In terms of the causal conception of explanation I am working with, this just means that the gathering of further evidence can resolve disputes about what is genuinely causally responsible for the phenomenon in question.

But this does not address the original concern. The original concern was that given only the *initial* evidence, and given two scientists claiming that immediate explanatory inference applied to this evidence leads them in incompatible directions, then with no rule governing this form of inference there is nothing in virtue of which we can judge that one scientist is reasoning well and the other poorly. We presumably do not want to say that reasoning with the same form of inference and the same data can correctly lead to contradictory conclusions. In virtue of what then can we criticize one or another (or both) of our scientists?

There are many potentially complicating factors here. Suppose for instance that the initial evidence includes the fact that there are experts with differing opinions on the cause of the observation. The fact that the evidence before me includes the fact that an expert whom I respect has reached a different conclusion may sometimes function as a piece of counter-evidence that robs me of my warrant to perform the immediate explanatory inference I would otherwise have been inclined to perform. How generally this holds gets into issues in the epistemology of disagreement that have been hotly contested and that I will not try to take sides on here.

To avoid these issues, let us suppose that in our example the initial evidence does not include facts about how other scientists see the situation, but instead only includes the first-order observational evidence itself. In virtue of what might one (or both) of these scientists be reasoning poorly?

There will certainly be situations in which one of the scientists is patently drawing an unwarranted conclusion, and we can alert them to this fact by directing their attention to something in the evidence before them *without* introducing new evidence or pointing to different points of view among experts. Perhaps the scientist has been discounting some of the evidence – either through carelessness, or as the result of some sort of bias. In such a case, although the scientist thinks they are warranted in inferring a particular explanation for the evidence before them, they are simply wrong. By redirecting their attention to some overlooked feature of the evidence, or explicitly pointing out the bias in question, they can be corrected. Nothing in this process presupposes the existence of rules undergirding immediate explanatory inference.

There are also further cases in which although the dispute *could* be resolved in this way, we do not yet see exactly how – that is to say, although one (or both) of the

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

scientists are overlooking something in the evidence or are in the grips of a bias, no-one can pinpoint the precise error with any sort of precision. But what any currently existing scientist can or cannot pinpoint is beside the point. The fact that in these cases the dispute would be possible to resolve were we to appreciate the error in question is enough to address the concern that without laws governing the relevant inferential practices, dispute resolution is in principle impossible.

But might there not be cases in which two scientists reasoning faultlessly from the same body of data nevertheless find themselves drawn to rival explanations? Perhaps examples of this are to be found the early days of a Kuhnian scientific revolution, when science finds itself in crisis. Or perhaps there are cases in which two scientists simply see the data differently – one sees a rabbit, while the other sees a duck, as it were – and neither party is guilty of any sort of epistemic negligence, nor is it possible to draw either party's attention to any feature of the data before them to alter their conviction in their inferred explanation.

Whether such examples really exist is not clear to me. Even the person who sees the rabbit can, after all, typically be made to see the duck, and in the process might well change their justified explanatory beliefs. (I thought Bobby was trying to give me a hint to buy him a pet rabbit when he drew this picture for me, but now I see that he might also be open to the idea of a pet duck!) The example we want must be one in which this sort of reconciliation is in principle impossible. Do such cases exist? Kuhn describes a scientist who in the very early 1900s remains convinced that all the effects of relativity can be explained in purely Newtonian terms, while at the same time a different scientist with the very same data infers the theory of relativity. Kuhn finds no irrationality in either party. An issue I will touch on later is whether there might be a confusion between justified belief and reasonable conjecture at play in this sort of example. For now, it suffices to claim that if such Kuhnian examples exist, they are highly anomalous, and not at all common. That a form of inference not guided by rules might occasionally lead to situations like this might even be regarded as unsurprising. So long as such situations are rare, there does not seem to me to be a particularly strong argument against the very possibility of immediate explanatory inference here.

4. Extending the Paradigm

Thus far I have largely been on the defensive, arguing that certain questions about immediate explanatory inference do not constitute threats to the very possibility of its existence. Perhaps at this point the reader is willing to grant that there is a form of inference in which an explanation is immediately inferred from the first-order evidence supporting it. In the rest of the paper, I will consider what general role such a form of inference might play in our investigations of the world, and in particular in science.

4.1. The gathering of evidence

I have suggested that our everyday lives are filled with examples of immediate explanatory inference. When we infer from the sound of a horn that a car is outside, or that a friend believes something when they assert it seriously, it is immediate explanatory inference that is being employed. Based on these sorts of examples, one might suspect that immediate explanatory inference will only be useful in giving us fairly trivial

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

information about the world, and that it will be of little use in more complex or scientific deliberation about the world.

Let us then turn our attention to cases of non-deductive inference that involve much deliberation. We consider examples both from ordinary life and from science, beginning with an example from ordinary life. If we suppose in our dog example that the plant is upturned and *both* the dog and the cat's faces are soiled, it is much less clear what we may infer. (Perhaps it is not even clear that we can infer that it was either the dog or the cat – perhaps all bets are off when the evidence is sufficiently confusing.) In this case, further investigation is needed before we can draw any substantive conclusions about the cause of the mess. Such further investigations will surely involve the explicit comparison of the hypothesis that the cat did it with the hypothesis that the dog did it, and perhaps further hypotheses still. Given this, one might think that immediate explanatory inference will not be of much use to us, and that something more like IBE will give us our final verdict. Likewise, one might think that the slow and deliberate investigations of science in which hypotheses compete against each other until a clear victor emerges do not lend themselves to immediate explanatory inference, but rather to something more like IBE.

I think however that this conclusion is incorrect. If both the dog and cat's faces are soiled, the problem is that we have insufficient evidence on which to infer the kind of explanation we seek. Thus we will likely search the house for clues, perhaps looking more closely at the upturned plant, the cat, and the dog in hopes of finding new evidence that settles things one way or another. Let us suppose that upon doing so we see that there are dog-like pawprints near the plant but no cat-like ones, and that the soil on the dog's nose is identical to that in the houseplant but that the soil on the cat's nose is utterly different, and more similar to that in the backyard. Going to the backyard, we find the cat's pawprints near a pile of soil identical to that on the cat's nose. At this point, we are surely right to infer the dog's guilt. (The cautious reader can supplement this story with even more damning new-found evidence for the guilt of the dog if the evidence described seems insufficient.)

The conclusion that in this case the dog is guilty is just as impossible to resist as the conclusions in the original dog or torn page examples, and likewise seems to me to be justified purely by the first-order evidence at hand. I claim that the inference in question is then an immediate explanatory inference. The main difference between this example and our earlier examples is that the evidence is not presented to us all at once, but has come to us largely over time as the result of an explicit search for further evidence. In fact, the desire to perform an immediate explanatory inference giving us a particular sort of explanation has presumably guided our search for evidence. For we presumably had a sense of what sort of evidence would suffice to infer various hypotheses by immediate explanatory inference, and our search for evidence has proceeded under the hope of uncovering enough evidence to perform precisely one of these immediate explanatory inferences. It may well be that in the process of deciding how to conduct our search for evidence, we compared and contrasted various conjectures. But such deliberations should not be viewed as constituting premises or sub-arguments in some larger piece of non-deductive reasoning that delivers justified belief in the explanation in question. Rather, the first-order evidence that we so happened to discover is what justifies belief in the explanation in question. Deliberation may have played a significant role in the discovery of this evidence, but such deliberations must be distinguished from evidence itself, on the basis of which we infer the explanation in question. With this in mind, I

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

claim that immediate explanatory inference forms the heart not just of the original dog example, but also the more complex variant we have considered here.

4.2. An example from science

Many examples from science may be treated similarly. Inspired by known phenomena such as stellar aberration (a correction that needs to be introduced into our astronomical measurements due to the velocity of the earth, known already in the 1700s) and Fizeau's experiments examining the propagation of light through various media in motion conducted in 1851, Einstein was led to postulate the theory of special relativity in 1905.[10] It was a while, however, before the theory of special relativity was to become established science. Part of the problem was that in 1905 detailed quantitative confirmation of special relativity's most surprising consequences (such as time dilation) had not occurred. A further problem was that rival approaches to the problem of reconciling the Lorentz invariance of Maxwell's equations of electromechanics with the apparent Galilean invariance of particle mechanics (e.g. the emission theories of Ritz and others) had not yet been refuted, and shared at least some of the virtues of Einstein's theory of special relativity.

In the decades after 1905, however, emission theories floundered as precise numerical confirmation of even special relativity's most striking consequences occurred. Emission theories were generally regarded as refuted by De Sitter's studies of binary stars in the early 1910s. By contrast, precise numerical confirmations for purely special relativistic phenomena such as time dilation were obtained by Ives and Stilwell in 1938 and Rossi and Hall in 1941, amongst others. In the meantime, general relativity – a theory of gravity that in certain ways generalizes special relativity – was also to enjoy considerable empirical success in the 1910s, providing quantitatively precise accounts of the perihelion precession of Mercury and the deflection of light by the sun. Further extremely precise confirmations of both general and special relativity were to follow. In addition, the predictions of modern particle physics, which also presupposes special relativity, were confirmed and continue to be confirmed with stunning accuracy in countless experiments.

What are we to make of these repeated and remarkable successes of the special theory of relativity in such a diverse set of experiments, and the dismal failure of its rivals? The obvious explanation is the truth (or at least, high level of accuracy or empirical adequacy) of the theory of special theory relativity itself. The evidence in question warrants this inference on its own (just as it did in the modified dog example), and we do not need permission from any sort of general rule of inference to be justified in believing it. Given all the evidence we now have, it is an immediate explanatory inference that justifies our belief in special relativity.

Just as with the modified dog example, I claim that immediate explanatory evidence guided the search for what turned out to be evidence for special relativity. Similar stories could be told about Darwin's theory of evolution, or Semmelweis' discovery that the lack of certain basic hygiene practices was a cause of childbed fever. Far from being of limited use in science, I claim that immediate explanatory inference is one of the central tools in the confirmation of theories in general, and governs the way in which we search for evidence in cases in which it is initially lacking or insufficient.

[10]How large a role the results of the Michelson-Morley experiment played in Einstein's formulation of special relativity is the subject of debate, but does not matter here.

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

4.3. Justified beliefs and conjectures

Even though in 1905 Einstein's theory had not been precisely quantitatively confirmed, it would not have been unreasonable for someone to argue at that time that special relativity offered a much simpler and much more elegant explanation of the outcome of the Michelson-Morley and Fizeau experiments than its rivals. Thus, even before the precise quantitative confirmations of special relativity that were to come in the following decades, would one not have been in a position in 1905 to perform an IBE to obtain justified belief in special relativity? If so, this would mean that although in some cases immediate explanatory inference can render belief in a given theory justified, in other cases it is something more like IBE that does this job.

I think however that this question must be answered in the negative. As charming and intriguing the theory of special relativity was in 1905, I do not think we could say that prior to the detailed quantitative confirmation it subsequently enjoyed that it could have been the object of *justified* belief. In 1905, the theory of relativity was certainly a remarkable and promising conjecture. A good part of what made it such a remarkable and promising conjecture was the simplicity and elegance with which it explained, relative to its rivals, things such as the outcomes of the Michelson-Morley and Fizeau experiments. A process in which we explicitly compare the simplicity and elegance of rival theories is a perfectly reasonable way to identify promising conjectures. But absent precise quantitative confirmation, it is hard to see how such a process could ever deliver justified belief. Thus, many putative examples of IBE – such as an attempt to use IBE to justify belief in special relativity in 1905 – are not really examples of 'inference' at all, at least insofar as we think of correct inference as giving us justified belief. They are rather examples of arriving at especially promising or reasonable *conjectures*. This is reminiscent of something like abduction in the Peircian sense. A conjecture can of course guide us as we search for more evidence, and once enough evidence is accrued, immediate explanatory inference can transform a conjecture into justified belief. This is precisely what happened in the decades after 1905 as special relativity came to be accepted (and perhaps also what happens in the modified dog example). But a process of arriving at a promising or reasonable conjecture should not be confused with a process of arriving at justified belief, and in light of this, it is difficult to see any clear and general sense in which IBE can sit alongside immediate explanatory inference as an alternative way of arriving at justified belief.

5. Comparison with the Peirceans

One might wonder how the views developed here compare with both Peirce's views on abduction as well as contemporary discussions of abduction; especially the modern literature on so-called 'creative abduction'. I discuss this now. One immediate difficulty with making comparisons with Peirce is that his views on abduction changed considerably over time. For the sake of simplicity, we focus on the later Peirce. Here, scholars generally agree that the later Peirce did not see abduction as a mode of inference yielding justified belief, but rather as a creative process in which some sort of explanatory suggestion was merely *hypothesized* or *suggested*.[11] In this way, Peircian abduction is quite unlike IBE, in that it is not designed to generate justified beliefs (or 'accepted

[11]For example, in 1901 Peirce writes that 'abduction is, after all, nothing but guessing' (Peirce 1992: 107). In his 1903 Harvard Lectures on Pragmatism, he writes that 'abduction is the process of forming an

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

science') on its own. (This point has been emphasized in Minnameier (2004), Campos (2011) and Psillos (2011), among other works.)

Peirce's view was that after abduction suggests a hypothesis, one must separately test whether the consequences of this hypothesis accord with experience. If after a great deal of experimentation all the consequences of the abductively given hypothesis are seen to accord with experience, it may be accepted. This later process of confirmation of the hypothesis given by abduction Peirce describes as *induction*.[12] On this view, inquiry then involves an abductive and inductive step. Although Peirce tends not to use the terminology of justification in talking about such things, we may say that while abduction generates a hypothesis worthy of closer investigation, induction confirms this hypothesis and gives us justified belief.

The Peircean conception of abduction as something merely delivering a hypothesis worthy of closer investigation has been developed further in the modern literature on so-called *creative abduction*, especially in works such as Magnani (2001, 2009), Schurz (2008), Yu and Zenker (2018) and Feldbacher-Escamilla and Gebharter (2019). Of particular note is the distinction in Magnani (2009) and Schurz (2008) between *creative* and *selective* abduction. In creative abduction, plausible explanations worthy of further explanation are generated, while in the subsequent selective abduction, the candidate explanation(s) most worthy of investigation are selected. Peirce often alternates between talking of abduction as a creative and a selective process, so explicitly distinguishing them and treating them separately as Schurz, Magnini and others do is a genuine advance.

Either way, abduction viewed in any of these loosely Peircean ways does not deliver justified belief, but only a promising hypothesis (or hypotheses) worthy of further investigation. This is in contrast with immediate explanatory inference, which yields fully justified belief. In the dog and torn page examples we do not merely get a hypothesis that deserves closer scrutiny, but a hypothesis supported by evidence which we are justified in believing. Immediate explanatory inference ends an inquiry, and in this sense is quite unlike Peircian or creative abduction.

Indeed, it is not clear how to characterize immediate explanatory inference from such Peircean points of view. In the case of the dog or torn page examples, it is difficult to distinguish any sort of abductive moment from any sort of subsequent inductive process, or any sort of sequence of creative abduction, selective abduction, then inductive confirmation. If anything, it is the first-order evidence that suggests, selects and confirms the hypothesis that the dog tore up the plant all at once. In this way, Peirce's abductive and inductive moments of inquiry seem to coincide. I am not aware of any passage in which Peirce acknowledges the possibility of these various moments coinciding, but it is not clear to me that Peirce prohibits this possibility either.[13] Likewise for the modern literature on creative abduction. Perhaps a

explanatory hypothesis' (Houser and Kloesel 1992: 216). In 1905 he writes that 'abduction … consists in examining a mass of facts and in allowing these facts to suggest a theory' (Peirce 1997: CP 8.209).

[12]Peirce writes that 'the operation of testing a hypothesis by experiment, which consists in remarking that, if it is true, observations made under certain conditions ought to have certain results, and then causing those conditions to be fulfilled, and noting the results, and, if they are favorable, extending a certain confidence to the hypothesis, I call induction' (Peirce 1997: CP 6.526).

[13]Peirce says that 'generally speaking, the conclusions of Hypothetic Inference cannot be arrived at inductively, because their truth is not susceptible of direct observation in single cases. Nor can the conclusions of Inductions, on account of their generality, be reached by hypothetic inference' (Peirce 1997: CP 2.714). Whether this should be viewed as an absolute prohibition against the idea that the abductive

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

comparison could be made with Peirce's view of perceptual judgments, which Peirce views as 'an extreme case of abductive inferences, from which they differ in being absolutely beyond criticism' (Peirce 1997: CP 5.181). But I will not pursue this comparison further, as it takes us too far deep into tangential issues in some of the most perplexing parts of Peirce's philosophical system.

It is however important to note that not all cases of immediate explanatory inference are like the dog or torn page examples in this respect. In the case of special relativity, for example, the relevant hypotheses became compelling hypotheses well worthy of further investigation around 1905 or earlier, while such hypothesis could only be said to have been confirmed over the subsequent decades. Our coming to justified belief in special relativity occurred via an immediate explanatory inference, I claim, but the corresponding creative or selective abductions happened significantly earlier. Thus, not all immediate explanatory inferences involve the moments of creative abduction, selective abduction and inductive confirmation coinciding. So while immediate explanatory inference must be sharply distinguished from any sort of Peircean conception of abduction insofar as the former generates justified belief and the latter does not, exactly how immediate explanatory inference fits (or fails to fit) into broader Peircean accounts of inquiry may vary from case to case.

6. Conclusion

The arguments given here might create the impression that there is no room left for IBE in our inferential tool-box. But this would be an overreaction. In immediate explanatory inference, we infer explanations. Moreover, the fact that belief in such explanations is warranted by the first-order evidence surely makes them the *best* explanations. (How could an explanation *not* warranted by the evidence be better than one that is?) It would thus not be a complete abuse of terminology to think of immediate explanatory inference as IBE re-conceived. The critical distinction, however, is that in immediate explanatory explanation it is not in virtue of an explanation being the 'best' of some larger set of explanations that we are justified in inferring it. It is rather in virtue of it being supported by the first-order evidence that we are justified in inferring it.

In making this move, I avoid some of the puzzles from which I earlier suggested IBE has difficulty extricating itself. At the beginning of the paper I suggested that any account of inference to the best explanation must either be a form of argument in which it having been established that some explanation E is the best of *all* possible explanations, we thereby infer E, or a form of argument in which E is inferred from the fact that E is the best of all explanations *that we can presently think of*. I suggested that both options lead to serious problems. However, insofar as immediate explanatory inference involves nothing more than consideration of the first-order evidence, it takes neither of these forks in the road, but rather creates a different path ahead, avoiding such difficulties. For reasons like these as well as others suggested elsewhere in this paper, I therefore think that immediate explanatory inference is worthy of further attention.

References

Campos D. (2011). 'On the Distinction Between Peirce's Abduction and Lipton's Inference to the Best Explanation.' *Synthese* **180**, 419–42.

moment and inductive moment of an investigation could ever coincide, or whether Peirce is merely claiming that in general these moments do not coincide, is not entirely clear.

https://doi.org/10.1017/epi.2023.19 Published online by Cambridge University Press

Churchland P. and Hooker C. (eds) (1985). *Images of Science: Essays on Realism and Empiricism*. Chicago, IL: University of Chicago Press.
Dellsen F. (2017). 'Reactionary Responses to the Bad Lot Objection.' *Studies in History and the Philosophy of Science* **61**, 32–40.
Dretske F. (1973). 'Perception and Other Minds.' *Noûs* **7**(1), 34–44.
Earman J. (1992). *Bayes or Bust*. Cambridge, MA: MIT Press.
Feldbacher-Escamilla C. and Gebharter A. (2019). 'Modeling Creative Abduction Bayesian Style.' *European Journal for Philosophy of Science* **9**, art. 9.
Hempel C. (1942). 'The Function of General Laws in History.' *Journal of Philosophy* **39**(2), 35–48.
Hempel C. and Oppenheim P. (1948). 'Studies in the Logic of Explanation.' *Philosophy of Science* **15**(2), 135–75.
Houser N. and Kloesel C. (1992). *The Essential Peirce, Selected Philosophical Writings, Volumes 1 and 2*. Bloomington, IN: University of Indiana Press.
Ladyman J., Douven I., Horsten L. and van Fraassen B. (1997). 'A Defence of van Fraassen's Critique of Abductive Inference: Reply to Psillos.' *Philosophical Quarterly* **47**(118), 305–21.
Lipton P. (2004). *Inference to the Best Explanation*, 2nd edition. Abingdon: Routledge.
Magnani L. (2001). *Abduction, Reason, and Science*. Dordrecht: Kluwer.
Magnani L. (2009). 'Creative Abduction and Hypothesis Withdrawal.' In J. Meheus and T. Nickles (eds), *Models of Discovery and Creativity*, pp. 95–125. Dordrecht: Springer.
Minnameier G. (2004). 'Peirce-Suit of Truth: Why Inference to the Best Explanation and Abduction Ought Not to Be Confused.' *Erkenntnis* **60**(1), 75–105.
Norton J. (2010). 'There are no Universal Rules for Induction.' *Philosophy of Science* **77**(5), 765–77.
Norton J. (2021). *The Material Theory of Induction*. Calgary: University of Calgary Press.
Okasha S. (2001). 'What did Hume Really Show About Induction?' *Philosophical Quarterly* **51**(2–4), 307–27.
Peirce C.S. (1878). 'Deduction, Induction and Hypothesis.' *Popular Science Monthly* **13**, 470–82.
Peirce C.S. (1992). 'On the Logic of Drawing History from Ancient Documents.' In N. Houser and C. Kloesel (eds), *The Essential Peirce, Selected Philosophical Writings, Volumes 1 and 2*, vol. **2**. pp. 75–114. Bloomington, IN: University of Indiana Press.
Peirce C.S. (1997). *Collected Papers of Charles Sanders Peirce*. London: Thoemmes Continuum.
Psillos S. (1996). 'On van Fraassen's Critique of Abductive Reasoning.' *Philosophical Quarterly* **46**(182), 31–47.
Psillos S. (2011). 'An Explorer Upon Untrodden Ground: Peirce on Abduction.' *Handbook of the History of Logic* **10**, 117–51.
Salmon W. (2006). *Four Decades of Scientific Explanation*. Pittsburgh, PA: University of Pittsburgh Press.
Schupbach J. (2014). 'Is the Bad Lot Objection Just Misguided?' *Erkenntnis* **79**(1), 55–64.
Schurz G. (2008). 'Patterns of Abduction.' *Synthese* **164**, 201–34.
Scriven M. (1959). 'Truisms as the Grounds for Historical Explanation.' In *Gardiner, Theories of History*, pp. 443–75. Glencoe: Free Press.
Scriven M. (1962). 'Explanations, Predictions, and Laws.' *Minnesota Studies in the Philosophy of Science* **3**, 170–230.
Stanford P.K. (2006). *Exceeding our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.
Thagard P. (1978). 'The Best Explanation: Criteria for Theory Choice.' *Journal of Philosophy* **75**(2), 76–92.
van Fraassen B. (1989). *Laws and Symmetry*. Oxford: Oxford University Press.
Yu S. and Zenker F. (2018). 'Peirce Knew Why Abduction Isn't IBE – A Scheme and Critical Questions for Abductive Argument.' *Argumentation* **32**, 569–87.

Kevin Davey is an Associate Professor in the Department of Philosophy at the University of Chicago.

Cite this article: Davey K (2024). On Inferring Explanations and Inference to the Best Explanation. *Episteme* **21**, 1120–1137. https://doi.org/10.1017/epi.2023.19