What Is Empirical Testing?

Michael Strevens

Draft of February 2008

Abstract

Science is epistemically special, or so I will assume: it is better able to produce knowledge about the workings of the world than other knowledge-directed pursuits. Further, its superior epistemic powers are due to its being in some sense especially empirical: in particular, science puts great weight on a form of inductive reasoning that I call *empirical confirmation*. My aim in this paper is to investigate the nature of science's "empiricism", and to provide a preliminary explanation of the connection between empirical confirmation and epistemic efficacy. I will try to convince you that the place to find an account of empirical confirmation is the dusty, long-neglected instantialist account of scientific inference offered by mid-century logical empiricists. Some revision of instantialism will be required. As for what is advantageous in empirical confirmation, I propose that it is an unusual degree of independence from background belief.

> Take nothing on its looks; take everything on evidence. There's no better rule.
>
> Dickens, *Great Expectations*

1. Empirical Exceptionalism

Why do we—or some of us, at any rate—think that science is epistemically privileged, that it has a special power to produce knowledge, or to find the truth, about the way the world works? A power more special than, say, that of folk medicine, or conspiracy theorizing, or astrology, or systematic metaphysics, or various forms of religious activity, such as the interpretation of sacred texts?

A familiar answer is that science has a special epistemic connection to the world's workings: it is *empirical* in a way that the other knowledge-directed activities enumerated above are not. You might say that this connection consists in science's being guided above all by the empirical evidence, or its being concerned above all with empirical testing. The implication, then, is that there is a certain kind of data, the *empirical evidence*, and a certain way of drawing conclusions about hypotheses from this data, which you might call *empirical confirmation*, that is first, endemic to science, and second, laden with epistemic goodness. (Note that I am using the term *confirmation* here in its technical philosophical sense, in which a hypothesis is said to be confirmed by a datum if the datum provides any degree of support at all for the hypothesis; to say that a hypothesis is confirmed by the evidence, then, is not to say that we have good reason to believe it, only that we have better reason to believe it than before.)

Consider, then, the thesis that science is epistemically efficacious because it is especially empirical. The aim of this paper is, first, to find a meaning for the term *empirical* on which the thesis is substantive and plausible, and second, to provide at least a preliminary defense of the thesis.

1.1 A Test Case

The contrast that I draw between empirical testing in science on the one hand, and the kinds of reasoning found in various non-scientific pursuits on the other, is unavoidably diffuse and is further muddied by the inevitable disagreement about the methodology of scriptural hermeneutics, philosophy, folk medicine, and so on. It will be helpful to focus the discussion on a far more precise contrast, a particular pair of inductive procedures, only one of which is empirical in the scientific sense that interests me here. There is of course a danger that my chosen example—a toy case, really—is not a paradigm of the contrast between empirical testing and other inductive methodologies; I think there is not much to be gained by attempting to ameliorate this worry in advance.

Suppose that Empiricus and Lector are inquiring into the color of ravens. Empiricus observes a selection of ravens in the wild, and notes that they are black, inductively reaching the conclusion that all ravens are black. Lector heads to the library, where the contents of a reliable ornithology book provide him with good reason to believe that all ravens are black.

Both Empiricus's and Lector's inferences to the truth of the hypothesis *All ravens are black* are paradigms of good inductive reasoning.[1] Lector's is in many ways more reliable than Empiricus's, and is likely more efficient as well. It is Empiricus's, however, that strikes us as a properly scientific investigation. This is not to say that scientists learn nothing from books and journals—on the contrary, they no doubt learn much more from the scientific literature

1. Or so I will suppose. Some epistemologists treat knowledge gained through testimony differently; whether the expert testimony found in scientific texts counts as such, I am unsure.

than they could ever learn in the lab. But it is not book-learning that makes science special.

Indeed, it is universally acknowledged that the epistemic status of the written word in science depends ultimately on the empirical testing that supports what is written— we take the ornithology book seriously, for example, because we think that someone once did the same sort of testing that Empiricus does, and that the conclusions of the book are supported by the tests in question.

1.2 What Is Empirical?

My computer says that *empirical* means *based on, concerned with, or verifiable by observation or experience rather than theory or pure logic*. And a naive observer might indeed conjecture that science owes its special status to its paying close attention to observed phenomena (the empirical data) and drawing whatever conclusions these phenomena inductively warrant (empirical confirmation). Philosophers will not, however, be satisfied with such a thesis: the sacred texts are just as much observed as the outputs of particle accelerators, and observations of the movements of the planets serve as the basis for both astronomical and astrological reasoning. Or confining your attention to my test case, both the contents of an ornithology text and the color of ravens in the wild are learned from experience.

We need a narrower notion of what is empirical, then, to comprehend the scope and limitations of the doctrine of empirical exceptionalism. One suggestion sometimes heard is that a datum is empirical if it is produced by a repeatable, controlled scientific experiment. Such a definition would, however, rule out vast swathes of scientific evidence as empirical: the fossil record, the cosmic microwave background and all other cosmological data, just about all of social, political, and economic history, and so on. On a smaller scale, most observations of raven color are not experimentally produced.

Another suggestion is that certain hypotheses are themselves "unempirical". The logical empiricists, for example, proposed that scientific hypotheses should have "empirical content", or very roughly, "observational consequences". Or consider a criterion sometimes put forward in the debate about the scientific merits of creationism and intelligent design, that hypotheses should be "naturalistic". It is difficult to give non-question begging substance to these notions; Hempel (1965), for example, provides an exhaustive (and exhausting) rebuttal of various criteria for empirical contentfulness. But in any case, proposals of this sort seem too broad to capture the kind of distinction I am pursuing in the present paper: in my test case, the empirical and the non-empirical methods lend support to the very same hypothesis about the color of ravens, and more broadly, astrological predictions and the claims of conspiracy theorists may be quite empirical and quite naturalistic.

It is better, I think, to focus not on a certain kind of data or a certain kind of hypothesis, but on a certain kind of reasoning connecting data and hypothesis: science is especially empirical, I will suggest, because it places special emphasis on a particular inductive method, what I have been calling empirical confirmation. The "empirical data" can then be defined as any phenomena that play the role of premises in empirical confirmation, and the "empirical hypotheses" as any propositions that play the role of conclusions (or something like that). The questions this paper must answer, then, are *What is empirical confirmation?* and *What are the advantages of empirical confirmation?*

A historical note: My project can be viewed as a contribution to what is called the "demarcation problem" in the philosophy science, that is, the problem of picking out what is methodologically distinct about scientific as opposed to other forms of inquiry into the workings of the world. The demarcation problem is often regarded as defunct, a dead end in philosophy of science. On this view there is no magical wonder method that is the sole preserve of science: any technique used in science is (or can be) used elsewhere, and no (or almost no) technique is off limits to scientists. The dissimilarities among various knowledge-seeking endeavors are not qualitative epistemic differences, then; they may be differences of emphasis (as perhaps in the interpretation of scriptural hermeneutics contrasted with science), differences in social structure (as perhaps in philosophy contrasted with science), or differences in competence (as in pseudoscience contrasted with science). If science has been more successful than these other enterprises, it is because scientists are harder working, more generously funded, better organized, less tolerant of wishful thinking and other epistemic sins, perhaps even intellectually more talented.

Yet questions about demarcation have not gone away, and they continue to hinge, as they always have, on the importance of science's distinctly empirical methodology. In the debate about intelligent design, for example, some ID skeptics run an argument of the following form:

> There is something about the ID movement, quite aside from the motivations of its members, that is profoundly unscientific. Although any particular piece of ID reasoning might have its analogs in real science, the intellectual structure of the program as whole is missing something that is essential to science, namely, a core comprising a substantial body of empirical testing.

Whether or not this is an accurate characterization of ID, it points to a strong sense among practitioners of science that their endeavor is epistemically special in some way, and that its specialness is related to its empirical aspect.

Although the present paper is not focused on the demarcation problem, then, I accept that the fate of my project is linked to that project. While it is likely true that a certain extraordinarily ambitious version of the problem is intractable—a version on which the solution would declare for any isolated piece of reasoning (or datum, or hypothesis) whether it was "scientific" or not, and on which a knowledge-seeking endeavor would qualify as a science just in case it dealt exclusively in such individually "scientific" items—this hardly precludes any number of other less totalizing attempts to find inherently epistemic differences between science and other intellectual pursuits. I will suggest, in particular, that real science, though it is not strictly limited to any particular style of reasoning or data, must have empirical testing at its heart.

2. What Is Empirical Confirmation?

Empirical confirmation is one kind, but not the only kind, of inductive support. Thus a theory of empirical confirmation is not a theory of inductive support, not a complete inductive logic. To put it another way, a theory of *empirical* confirmation is not a theory of *confirmation*, if "confirmation" is defined to occur whenever one proposition lends support to another.

Let me assume for expository purposes that there is a complete, correct account of inductive support.[2] Perhaps it is none other than the currently dominant theory of confirmation in the philosophy of science, the Bayesian theory, which aims to provide an understanding of inductive warrant that applies not only to certain forms of reasoning in science, but to every variety of justifiable inductive reasoning.[3] In any case, whatever it is, this universal inductive theory will tell us which inferences are inductively valid, or which propositions inductively support which other propositions, and why.

The role of an account of empirical confirmation is to tell us which of the inductive support relations distinguished by the universal theory are empirical. It will, then, distinguish a subclass of the support relations, the relations of empirical confirmation, and it will offer an account of what is

2. The alternative: the inductive support relations can be divided into two or more categories having little to do with one another; the nature of support within each category will be analyzed separately—much as deductive and inductive support are analyzed separately.

3. Nothing about the mechanics of Bayesianism is assumed in this paper, but in case you would like your memory jogged: In Bayesian confirmation theory, one proposition e provides inductive support for another proposition h if the subjective probability for h conditional on e is greater than the unconditional probability of h, that is, informally, if taking e into account raises the probability of h.

special about that subclass—not of why the empirical confirmation relations are inductive, since that is the purpose of the more general theory of support, but of what is special about the particular kind of inductive support that is properly regarded as empirical.

One way to pursue a theory of empirical support or confirmation is to begin with a universal theory of support and to use the same tools that brought out what all inductive inferences have in common to find distinctions between them. If Bayesianism is the universal theory, then—so that inductive support is just subjective probabilification—you would try to find a special kind of subjective probabilification that corresponds to the empirical support relation. The presupposition of such an approach is that, although Bayesian confirmation theory in its usual form does not care to make a distinction between empirical and other kinds of support, the distinction is nevertheless inherent in the Bayesian notion of inductive warrant.

My strategy is quite different. In the course of the twentieth century, a number of different theories of confirmation were constructed, many of which now seem too narrow to capture the inductive support relation in all its richness—because, for example, they lend too little weight to epistemic context in deciding what confirms what. Narrowness in a universal theory of inductive support is a vice, but it is precisely what is needed in a theory of empirical confirmation, where the goal is to focus on just one kind of support. Further, the source of the narrowness in twentieth century confirmation theory, a perhaps parochial exultation of observation-based science as an epistemic model, seems to promise exactly the right kind of narrowness.

Sure enough, it turns out that a particular, once-influential variant of confirmation theory, instantialist or instance-based confirmation, does a remarkably good job of capturing our pretheoretical intuitions about what is and is not empirical. In what follows, then, I will take a version of instantialism—Hempel's—and modify it to produce the beginnings of an account of empirical confirmation.

2.1 Hempel's Instantialism

The guiding idea of instantialism is that scientific hypotheses are *directly confirmed* by their positive instances and *directly disconfirmed* by their negative instances,[4] and that they are *indirectly confirmed or disconfirmed* if some other hypothesis to which they are appropriately related is directly confirmed or disconfirmed. The nature of the appropriate relation is spelled out by what might be called a *transmission rule*, a rule prescribing the way in which confirmation is transmitted from directly confirmed to other hypotheses.

As an example, let me explain the workings of Hempel (1945)'s instantialism. First, a preliminary: for Hempel, confirmation is a relation between a hypothesis and an observation statement, where an observation statement is a proposition about the observable properties of some entity, set of entities, or system, which might be called the observation set. When he talks of the evidence, or of the instances of a hypothesis, Hempel is talking about observation statements, not the things in the world that the observation statements describe, that is, not the entities in the observation set or their properties. I adopt the same idiom.

Second, the account of instantiation, and thus the account of direct confirmation. According to Hempel, a hypothesis of the form *All Fs are G* has as a negative instance an observation statement entailing that at least one of the entities in the corresponding observation set is an *F* that is not *G*. Consider, for example, the hypothesis *All ravens are black*. An observation statement that mentions only a single object is a negative instance of the raven hypothesis just in case it entails that the object is a non-black raven. If the statement mentions more than one object, it is a negative instance if it entails that at least one of the objects is a non-black raven. An observation statement is, by contrast, a positive instance of such a hypothesis if it entails

4. For statistical hypotheses, direct confirmation depends on the ratio of positive to negative instances; I will ignore this complication in what follows, effectively assuming that we are dealing with deterministic hypotheses only.

that no member of the observation set is a negative instance. The raven hypothesis is positively instantiated, then, by a statement concerning a single object if the statement entails that the object is *not* a non-black raven, which is to say, if it entails that it is either a black raven or not a raven at all. An observation statement that meets neither of these conditions is not a positive or a negative instance of the hypothesis.[5] For Hempel, then, the raven hypothesis is directly confirmed by any observation statement entailing that no members of its observation set are counterexamples to the hypothesis, and directly disconfirmed by any observation statement entailing that one or more members of its observation set is a counterexample.

Third, the transmission rule, and thus the account of indirect confirmation: if an observation statement directly confirms a hypothesis (or set of hypotheses), then it indirectly confirms any logical consequences of the hypothesis (or set). My knowledge that the most recently observed raven in Canada was black, then, directly confirms the raven hypothesis and indirectly confirms both the prediction that the next raven observed in Canada will be black and the hypothesis that all ravens in Mexico are black (since both are entailed by the directly confirmed raven hypothesis). On Hempel's theory, you will see, instance confirmation flows first along the instantiation relation to the instantiated hypothesis (direct confirmation) and then down the entailment relation from the hypothesis to its logical consequences (indirect confirmation).

2.2 Instantial Confirmation as Empirical Confirmation

I posed the following problem in section 1.1 above: what is the difference, if any, between the kind of inductive inference in which Empiricus comes to believe the raven hypothesis on the basis of many observations of black (and

5. I have given instantiation conditions only for simple if/then hypotheses. Hempel's theory gives instantiation conditions for any hypothesis that can be formulated in first order logic. There is no need for the details here.

no non-black) ravens, and the kind of inductive inference in which Lector comes to believe the same hypothesis on the basis of sentences inscribed in a trusted ornithological treatise?

Hempel and other instantialists have a ready answer: Empiricus's inference, but not Lector's, is a case of direct instance confirmation.[6] Lector's inference does not seem even to be a case of indirect instance confirmation; it is a kind of inductive inference that is not addressed by the instantialist theory of confirmation at all. Assume, without any further argument, that this is correct: Empiricus's inference is an example of instance confirmation, while Lector's inference is an example of some other kind of rational inductive reasoning. Could a theory of instantial confirmation, such as Hempel's, then be regarded as an account of empirical confirmation—that is, an account of the kind of inductive support that makes science special? Does science owe its empirical grip on the world to the emphasis it puts on instance confirmation?

The kind of reasoning that goes on in philosophy or scriptural hermeneutics does not look very much like instantial confirmation: many of the empirical inputs to these endeavors, such as inscribed sentences and intuitions about the nature of things, are not instances of generalizations to which they are supposed to lend support. But the paradigms of empirical testing in science fit the instantial account rather well. Even scientific observations that are not experimentally produced, such as fossils, are instances of hypotheses that they support: the configuration of a fossilized tooth might, for example, confirm the hypothesis that *Tyrannosaurus rex* had serrated teeth.

Clearly, a more systematic survey of empirical tests in science would be required in order to satisfy you that instantial confirmation and empirical confirmation go hand in hand. There are three issues I need to confront before I consider such a test.

First, let me emphasize that I am not so much reviving as repurposing the

6. Note to confirmation aficionados: as you know, there is serious obstacle to Hempel's giving this answer, to be discussed in note 9 below.

instantialist theory of confirmation: although instantialism is traditionally regarded as a general theory of inductive reasoning, I deploy it rather as a theory of empirical confirmation, that is, a theory of what distinguishes the empirical from other relations of inductive support. As such, my instantialism is intended to be compatible with any general theory of support, with any "theory of confirmation".[7]

In particular, it is at least prima facie compatible with a Bayesian account of inductive warrant. If Bayesianism is correct, then instantial confirmation should then be understood as one way in which a hypothesis can come to enjoy Bayesian inductive support, that is, one way in which an observation can increase your subjective probability for a hypothesis. Further work would be required to mesh the instantialist and Bayesian accounts consistently (see below); my point here is that in understanding instantialism as an account of empirical confirmation, you are not thereby committing yourself to any particular general account of inductive warrant.

Second, I want to raise briefly the possibility that there are several kinds of empirical confirmation, of which instantial confirmation is only one variety. This option will hover unacknowledged in the background for much of the discussion to follow; it will then be considered near the end of the paper.

Third, there are certain limitations and flaws in Hempel's account of instantial confirmation. If instantialism is to be considered an adequate theory of empirical confirmation, the problems must either be corrected or shown to create difficulties only for a universal account of inductive reasoning, as opposed to an account of empirical confirmation only. That will be the business of the next two sections.

7. It is possible that in some moods, Hempel himself intended instantialism in this way, that is, as an account of inductive reasoning in science rather than across the board. He does, however, appear to have universalistic aspirations for instance confirmation.

2.3 Hempelian Difficulties

A first, and considerable, limitation of Hempel's instantialism is that it allows the confirmation only of those hypotheses framed entirely in terms of observable properties. Consider, for example, a hypothesis of the form *All Fs are Gs* where *F* is a "theoretical" property, that is, a property that will not appear in any observation statement. The hypothesis cannot be directly confirmed, because Hempel's account of instantiation implies that *All Fs are Gs* has instances only if *F* is observable. To see why: recall that an observation statement is an instance of the hypothesis, positive or negative, only if it entails something about the *F*-ness or otherwise of the observation set. This is possible only if the presence or absence of *F* can be entailed by observable facts, which pretty much amounts to the observability of *F* itself. But for more or less the same reasons, the hypothesis cannot be indirectly confirmed: a hypothesis is indirectly confirmed just in case it is entailed by directly confirmed hypotheses, but only hypotheses concerning observable properties can be directly confirmed, and such hypotheses cannot entail facts about theoretical properties such as *F*. On Hempel's account, then, a theoretical hypothesis, such as *The force of gravity falls off with the inverse square of distance*, cannot be confirmed in any way.

Observe that this limitation of Hempel's instantialism is in part due to the fact that indirect confirmation never flows "upward", but only "downward". An instance of a hypothesis indirectly confirms the logical consequences of the hypothesis, but it never indirectly confirms a theory that itself has the hypothesis as a logical consequence, even in those cases where it seems that the theory should be confirmed.[8] Consider, for example, a series of observations that instantiate a hypothesis about the cosmic microwave background, that is, about the distribution of microwave radiation across and around the celestial sphere. Intuitively, such observations should confirm not only the

8. On Hempel's view, the instance might *directly* confirm the hypothesis's entailer, if it is strong enough to instantiate not only the hypothesis, but some stronger statement.

microwave distribution hypothesis but also the theoretical posits that we invoke to explain the distribution, in particular the fact that the universe underwent an epoch of "inflationary expansion" shortly after the big bang. On Hempel's account, the only way for the microwave data to confirm the hypothesis of inflation is by instantiating the hypothesis, which would require (given Hempel's account of instancehood for existential claims) that the observations report an actual episode of inflation. Even if instances were allowed to be theoretical, this is impossible: inflationary expansion concluded many billions of years ago (lasting, in fact, only a fraction of a second at the very beginning of time). Inflation itself cannot be observed, directly or indirectly; its existence must be inferred from its lingering effects, that is, from phenomena or patterns of phenomena that it explains. We would want to count at least some such inferences as empirical, but they seem not to fit Hempel's framework.

A second difficulty with Hempelian instantialism, also noted by Hempel himself, is its having the following untoward consequence: an observation statement saying of some particular object that it is a white shoe is a positive instance of the raven hypothesis, since it entails that the object is not a non-black raven, so not a negative instance. A white shoe therefore confirms the raven hypothesis. Much attention has been devoted to the question whether the observation of a white shoe can provide some inductive support for the raven hypothesis. If instantialism is understood not as a general theory of induction but as a theory of empirical confirmation, this is not quite the right concern. What needs to asked is whether the white shoe provides some inductive support for the raven hypothesis *qua* positive instance, that is, because it instantiates the hypothesis. And to this question the answer is clearly no, simply because a white shoe is obviously not in any known sense of the term an instance of the raven hypothesis. That Hempel's theory of instantiation implies otherwise is a serious flaw.[9]

9. The problem of the white shoe, or as Hempel called it, the raven paradox, in fact

Third and finally, in some circumstances a positive instance of a hypothesis will fail to lend it inductive support, and may even count against it. I. J. Good constructed a famous example in which the observation of a black raven leads the rational observer to think the raven hypothesis far less likely than it was before.[10] This is not, I think, a consequence of a deficiency of the machinery of Hempel's instantialism; it is, rather, a sign that instantialism should be understood not as a general theory of inductive warrant but rather as an account of a particular kind of inductive argument—as an account of empirical confirmation. In cases such as Good's the confirmatory power of a datum that accrues from its being an instance of a hypothesis is swamped, neutralized, or undermined by other inductive implications that the datum has for the hypothesis independently of its instancehood.

---

plays all sorts of havoc with the proposal to understand empirical confirmation as instantial confirmation. Consider Lector's inference from a sentence in an ornithology book to the truth of the raven hypothesis. On Hempel's account, any non-raven instantiates the raven hypothesis. But Lector's sentence—the sequence of letters on the printed page—is a non-raven. Thus it instantiates, and so empirically confirms, the hypothesis. Of course, Lector's inference is quite independent of the putative fact of instantiation: it is not because the sentence is (according to Hempel) an instance of the hypothesis that Lector treats it as powerful evidence for the hypothesis. On this basis, a Hempelian might attempt to maintain a distinction between Empiricus's and Lector's inferences—but Hempel's eccentric account of instancehood makes the going uphill all the way.

10. The case: Suppose you know that you have just been transported randomly, with 50/50 probability, to one of two worlds. Nine-tenths of the first world's birds are ravens, and nine-tenths of these are black. The chance of a random bird's being a black raven in this world is therefore 81%. In the second world, one bird in ten is a raven, and all of these are black. The chance of a random bird's being a black raven in this world is therefore 10%. You open your eyes. The first bird you see is a black raven. This is much more likely in the first world than the second world, so the observation gives you good reason to think that you are in the first, not the second world. In the first world, however, not all ravens are black. Thus your observation of a black raven has given you good reason to disbelieve that all ravens in your new world are black. As Good rightly reasons, the black raven, uncontroversially an instance of the raven hypothesis, provides strong inductive support for the hypothesis's negation (Good 1967).

2.4 A New Instantialism

I want to suggest an approach to instancehood that is quite different from Hempel's, and that avoids some of the problems noted in the previous section. There is a price for this deproblematization: the account of instantiation that I offer applies only to causal hypotheses of the form *All Fs are G* or *In conditions, Z, all Fs are G*. I understand *In conditions Z, all Fs are G* as a causal hypothesis when it is interpreted to mean something along the following lines: there is a single causal mechanism by way of which something about *F*s causes *G*-ness; among the conditions required for the operation of the mechanism are *Z*. For example, the raven hypothesis *All ravens are black* is a causal hypothesis because it is understood to mean that there is a single causal mechanism in virtue of which something about ravens causes blackness. (There is, I think, considerable psychological evidence for this interpretation of taxon/trait generalizations such as the raven hypothesis; see Strevens (2000).)

A causal hypothesis is always associated, I will assume, with a putative causal mechanism, the hypothesis's *underlying mechanism*, that is described, at least partly, in what I call the hypothesis's *associated theory*. (Of course this is a little too strong: we may posit a causal mechanism connecting two properties but have no substantive beliefs about the nature of the mechanism. You may think of such hypotheses as associated with a null or empty theory. And what if we have in mind several candidate mechanisms for a hypothesis? I will later consider such a case; for now, let me keep things simple.)

The associated theory will state or imply that certain conditions are required for the operation of the mechanism it describes. A causal hypothesis may make some of these operating conditions explicit, spelling out the fact that they must hold if *F*-ness is to cause *G*-ness by way of the hypothesis's underlying mechanism. The explicit conditions will occupy the place of *Z* in *In conditions Z, all Fs are G*. Typically, they are a rather small subset of the operating conditions enumerated by the theory.

The causal hypothesis *In conditions Z, all Fs are G* has as its positive instances, I propose, all *F*s that are *G* in conditions *Z*, and as its negative instances, all *F*s that are not *G* in conditions *Z*. Or perhaps it is better to follow Hempel in insisting that instances are observation statements, in which case, for an observation statement concerning a single entity *x*:

1. The statement is a positive instance of the causal hypothesis *In conditions Z, all Fs are G* just in case it entails that conditions *Z* hold (or held) for *x*, and that *x* is both *F* and *G*.

2. An observation statement is a negative instance of the causal hypothesis *In conditions Z, all Fs are G* just in case it entails that conditions *Z* hold (or held) for *x*, and that *x* is *F* but not *G*.

As on Hempel's account, a (deterministic) causal hypothesis is directly confirmed by its positive instances and directly disconfirmed by its negative instances.

Two remarks. First, an instance of *All Fs are G*, positive or negative, must be an *F*; white shoes and so on do not instantiate the raven hypothesis.[11] Second, you will observe that the causality of the causal hypothesis plays no role in the conditions for instancehood; it is central, however, in the account of indirect confirmation, as you are about to see.

I propose that when a positive instance directly confirms a causal hypothesis *In conditions Z, all Fs are G*, it indirectly confirms the following propositions:

1. All elements of the hypothesis's associated theory,[12]

2. All other causal hypotheses causally explained by the associated theory,

11. Hempel was prevented from taking this line on the shoe problem by his view that *All Fs are Gs* has the logical form $\forall x : Fx \rightarrow Gx$, and is thus equivalent to *All non-Gs are non-Fs*. The equivalence does not hold for causal hypotheses, which assert a causal directionality from *F* to *G*.

12. I take it that, by definition, all elements of the associated theory play an essential part in characterizing the hypothesis's underlying mechanism.

3. The particular initial conditions invoked by the associated theory to explain the observed instance (apart from those specified in Z, which are of course entailed by the observation statement),

4. Of any other F, that provided conditions Z hold, it is G.

Whereas in Hempel's theory, confirmation is transmitted downwards along logical relations, in this new *causal instantialism*, it is transmitted both upwards and downwards along causal-explanatory relations.

Let me work through an example. Suppose that the theory associated with the raven hypothesis has two parts: first, an empirical claim, that all ravens have a certain physiological property T, and second, a causal claim, that T causes blackness. Call the first part the *basing generalization* and the second part the *causal-mechanical law*.

You observe that a certain bird x is both black and a raven, directly confirming the raven hypothesis. The following propositions, corresponding to (1–4) above, are indirectly confirmed:

1. The associated theory's basing generalization and causal-mechanical law, that is, both the proposition that all ravens have T and the proposition that T causes blackness (along with the best understanding of why T causes blackness).

2. Any hypothesis explained by the basing generalization and causal-mechanical law. Suppose, for example, that your observed black raven x is Canadian. Then a number of specific raven hypotheses not instantiated by x are indirectly confirmed by x, for example, *All Australian ravens are black*. (Note that not just any logical consequence of the raven hypothesis can be indirectly confirmed: *If ravens are not black, then the moon is made of green cheese* is not confirmed, because it is not causally explained by the raven theory.)

3. The propositions that (a) x has T, and (b) whatever background conditions necessary for T to cause blackness held in x's case.

4. Of any other raven, that it is black.

This rather brief sketch of what I am calling causal instantialism will suffice for my principal purposes in this paper; I will sketch ways to fill some obvious gaps in section 2.6.

Let me conclude with a few remarks on the place of causal instantialism within the inductive hierarchy.

First, causal instantialism tells you which instances support which hypotheses *in virtue of their instantiating the hypotheses*. It is possible that a piece of evidence impacts a hypothesis by more than one "inductive path": it might support a hypothesis in virtue of instantiating it, while undercutting the very same hypothesis by some other, separate inductive argument (as in I. J. Good's famous example; see the end of section 2.3 and note 10).

Second, some forms of inductive inference that are not strictly speaking empirical may partake of the goodness of empirical confirmation. A paradigm is Lector's inference from the ornithology text: because the book reports the outcome of empirical testing, it is empirical itself at one remove. Empirical testing is the epistemic lifeblood of science; given suitable vessels, it may flow very far from the heart.

Third and finally, I am writing as though all instance confirmation in science is of causal hypotheses, though I have given you no reason to believe this. Let me simply leave open the possibility that there are some kinds of instance confirmation, hence some kinds of empirical confirmation (on the view that all instance confirmation is empirical confirmation) that remain unaccounted for by what I have said in this section.

What is the nature of empirical confirmation, then? If empirical confirmation is instance confirmation, and the nature of instance confirmation is captured by causal instantialism, then empirical confirmation is inference from a pattern of coinstantiation to the underlying causal mechanism that brings it about—from the G-ness of Fs in conditions Z to the existence of a causal mechanism by way of which, when Z holds, something about Fs causes *G*-ness. Not every good inductive inference has this form, but it is the inferences that do—the empirical inferences, or tests—that make science exceptional.

2.5 Theoretical Inference

Does causal instantialism avoid the pitfalls of Hempel's account enumerated in section 2.3? As explained earlier, it has no difficulties with white shoes and the like. The more significant challenge is to accommodate the paradigms of theoretical discovery in science—to account for cases of empirical inference in which we get from observable evidence to theoretical conclusions. I will show that theoretical propositions can be both indirectly and, in a certain sense, directly confirmed, according to the tenets of causal instantialism.

Indirect Confirmation On Hempel's account a directly confirmed hypothesis can be no theoretically deeper than the observation statement itself: whatever properties the hypothesis deals in, must be mentioned in or entailed by properties mentioned in the observation statement. Further, since an instance indirectly confirms only the logical consequences of the instantiated hypothesis, what is indirectly confirmed cannot be logically stronger than what is directly confirmed; there is no prospect of theoretical deepening here. But on causal instantialism, an instantiated causal hypothesis's associated theory is indirectly confirmed, a theory that is invariably deeper and stronger than the hypothesis itself. Unlike Hempel's instantialism, then, causal instantialism allows us to go from observable consequences to underlying mechanisms; in so doing, it captures the characteristic inferential movement in science from observed patterns to unobserved explanations of those patterns.

Consider how causal instantialism allows the indirect confirmation of the hypothesis that the universe went through a period of inflationary expansion. The occurrence of an inflationary event plays a part in the (putative) causal explanation of a certain feature of the cosmic microwave background, namely, its relative homogeneity. Observations of the microwave background instantiate a hypothesis to the effect that the distribution is relatively homogeneous, and so indirectly confirm the elements of the theory that explain the homogeneity, and thus the occurrence of inflation.[13]

Note, however, that not just any inference from an observation to facts that might explain it qualifies as instance confirmation: no doubt, the fact that all ravens are black plays a role in explaining why a sentence to this effect came to be inscribed in Lector's ornithology text, but this does not make Lector's inference a case of empirical confirmation, for two reasons. First, Lector's inference does not go by way of the direct confirmation of some pattern of the form *All Fs are G*. To put it another way, the explanation is not of a pattern of which the evidence—the sentence—is just one instance; it is rather a "one off" explanation of the inscription of the sentence as a singular event. Second, all, or almost all, of the inductive support that the observed sentence lends to the raven hypothesis is due, not to the hypothesis's playing a role in some putative mechanism, but to the inferrer's knowledge that the book in which the sentence appears is reliable.

Direct Confirmation For everything that I have said so far, causal instantialism still appears to share with Hempelian instantialism a certain limitation: there can be no direct confirmation of a hypothesis that deals in theoretical properties. Why not? A positive instance of *All Fs are G* is an *F* that is also a *G*. To observe a positive instance, then, it would seem that you would need to observe an *F* that is also a *G*, implying that *F* and *G* are observable properties.

There is no need, however, to require that an instance's *F*-ness and *G*-ness (or lack thereof) be directly observed. You might infer from other observa-

13. The distribution hypothesis is not of the causally directional if/then variety: it does not connect an antecedent cause to a consequent effect. Rather, it correlates two effects of an unstated common cause. I have not discussed such "common cause" hypotheses in the main text, but they too are causal hypotheses, and their instantiation conditions and resulting role in causal instance confirmation are analogous to those of directional if/then hypotheses—how so, I leave as a (straightforward) exercise for the reader.

tions that the object *x* before you is both *F* and *G*. You would thereby come to know that *x* is an instance of, and so directly confirms, *All Fs are G*. The hypothesis that all electrons have a charge of 1.602×10^{-19} C, for example, is instantiated by any electron with that charge. To know that you have such an electron, and thus a direct confirmer, you must somehow get from what you see—say, oil drops suspended in the Millikan apparatus—to the fact of electronhood and so on. Are there ways to do so that are themselves empirical? In other words, can the fact that you have an electron before you be instance-confirmed by what you observe?

It cannot, of course, be directly confirmed, but it can be indirectly confirmed: the fact that an oil drop is playing host to extra electrons is a part of the causal explanation of what you see, the fact that a drop produced in a certain way experiences acceleration in a certain direction in an electric field. (Compare the indirect confirmation of the observed raven's having *T* in the previous section.) To unpack each step in the inference (and assuming, for simplicity's sake, that the existence of the electric field is observable):

1. Your observation of the motions of a particular drop instantiate, and so directly confirm, the hypothesis that oil drops produced by such-and-such an apparatus will behave in a certain way in an electric field.

2. And so they indirectly confirm initial conditions invoked in the explanation of the hypothesis, including the presence of excess electrons in the oil drop,[14]

3. Which presence partly constitutes an instance of, and so directly confirms, the hypothesis about the magnitude of the charge on the electron.

14. Note that the explanation does not assume any particular magnitude for the charge on the electron, except perhaps some kind of lower bound. Millikan was able to infer from the behavior of the drops that they had excess electrons, without knowing the magnitude of the charge on the electrons.

It is precisely this sort of inference that I claim is mostly or wholly absent in non-empirical knowledge-seeking activities such as scriptural hermeneutics and analytic metaphysics.

2.6 *Competing Theories*

The sketch of causal instantialism presented in the previous two sections contains only the broadest outlines of an account of empirical confirmation, just enough to give you a sense of how an account of empirical confirmation, as opposed to inductive support, might turn out. Let me now give you some indication as to how one of the more yawning gaps should be filled.

I have so far assumed any given causal hypothesis has a single associated theory describing the causal mechanism in virtue of which the hypothesis is putatively true, or in other words, that there is a single putative causal explanation for any causal hypothesis. What if there is more than one such theory? What if there are two plausible stories as to how F-ness causes G-ness, and thus as to why all Fs are G?

The short answer is that both associated theories are empirically confirmed; this doctrine raises several questions, however, to be considered in what follows.

Quantifying Confirmation How is empirical confirmation to be apportioned among competing associated theories? The answer, I think, must be as follows: all matters relating to quantities of empirical confirmation are settled by the general theory of inductive support. This general theory, then, answers all questions of the form: to what degree does this proposition support this other proposition? The role of a theory of empirical confirmation is not to question these judgments, but simply to indicate which are judgments about purely empirical support, as opposed to non-empirical support or to a mix of empirical and non-empirical support.

More generally, all quantitative aspects of empirical confirmation are to be referred to the general theory of inductive support. It could not be any other way, I think (unless there is no general theory). If the theory of empirical confirmation offered its own quantitative judgments, these would either reflect the general theory's judgments, in which case they would be redundant, or they would contradict them, in which case the general and empirical theories of support would be inconsistent.[15]

Defeat Contrary to the simple picture of indirect confirmation given in section 2.4, a datum that instantiates a hypothesis does not always confirm its associated theory. More generally, such a datum may confirm only one of two or more rival associated theories.

Why? I proposed earlier that when a hypothesis is instantiated, the instance indirectly confirms not only the associated theory but the particular initial conditions in conjunction with which the theory would explain the instance. To recap my earlier example, suppose that the theory associated with the raven hypothesis *All ravens are black* posits a property T such that all ravens have T and T causes blackness. When the blackness of a particular raven x instance-confirms the raven hypothesis, the T-hood of x is confirmed, since the associated theory requires this initial condition to explain x's blackness. But what if for some reason, you know that your particular black raven x does not have T? Then not only is the initial condition (x's T-hood) not

15. Okay, it's not that simple. A theory such as the Bayesian theory of confirmation clearly quantifies the total support offered by one proposition to another, but it is not obvious how these judgments can be decomposed into judgments about the support offered along empirical as opposed to non-empirical pathways. For example, in the Good case (note 10), it is easy to calculate the posterior probability for the ravens hypothesis due to the observation of the black raven—a judgment of total (in this case, negative) support—but not so easy to extract facts about the degree of (positive) support offered to the hypothesis in virtue of the raven's being a positive instance. Presumably, if this question is going to be answered at all, it must be with reference to a counterfactual epistemic background in which the reasoner does not have knowledge of the composition of the two worlds. The theory of empirical support may play a role in determining the relevant background, that is, the relevant counterfactual prior probability distribution.

confirmed, the associated theory that posits *x*'s *T*-hood is not confirmed either. This holds whether or not there is a rival theory that explains blackness without reference to *T*. If there is, that theory alone gets confirmed; if not, no theory gets confirmed—confirmation stops at the hypothesis itself.

Discovering Causal Directionality I have so far assumed that whatever hypothesis is undergoing testing incorporates a definite assumption about causal directionality: *All Fs are G* implies that *F*-ness, or something about *F*-ness, causes *G*. Associated theories for hypotheses of this sort must agree on directionality; they can disagree only on the question how the causation gets done.

Frequently in science, however, the causal directionality of the mechanism underlying a hypothesis is unclear: you have an apparent correlation, expressed in the *All Fs are G* form, but you do not know whether *F*-ness causes *G*-ness or whether they are both effects of some common cause.[16] Suppose, in particular, that such a hypothesis has two associated theories, of differing directionality: one theory proposes a mechanism by which *F* causes *G*, the other a mechanism by which a common cause *T* causes both *F* and *G*. How can purely empirical confirmation distinguish the two theories, and so discover the true causal directionality behind the *F*/*G* correlation?

The usual experimental procedure would be to control for the putative common cause *T*. In the simplest case, you find some suitable systems that lack *T* and induce them to have *F* (using some method that does not independently induce *T*). If you get *G*, you confirm the *F*-as-cause theory at the expense of the common cause theory. Can my account of empirical confirmation make sense of such an inference as purely empirical? Yes; in the controlled experiment, you observe *F*/*G* pairs and so directly confirm *All Fs are G*. The confirmation flows on up to the theory that names *F* as a cause,

16. Typically, you can rule out the possibility that *G* causes *F*, on the grounds that *G* appears after *F*.

but not to the common cause theory, for the reason given in the discussion of defeat above: to explain the instances in question, the common cause theory must posit the presence of T, but the experimenter knows that T is absent. The indirect confirmation of the common cause theory is thus defeated.[17]

3. The Epistemic Virtues of Empirical Confirmation

What is epistemically special about empirical confirmation? Or on the assumption that I am right in equating empirical confirmation and causal instance confirmation, what is epistemically special about causal instance confirmation?

If you need to find out about the color of ravens fast and efficiently, you would do well to follow Lector's lead, making a non-empirical visit to the nearest ornithology text, rather than emulating Empiricus's painstaking fieldwork. More generally, if library-based learning is a form of non-empirical confirmation, it will be difficult to argue that as a matter of principle empirical testing has the quantitative edge when it comes to inductive support—that empirical evidence by its very nature lends a greater degree of inductive support to a hypothesis than information that bears on the hypothesis non-empirically, or is cheaper or easier or faster to obtain.

The superiority of empirical testing to other forms of inductive inference, if any, must therefore be qualitative: it must be the kind, rather than the amount or availability, of support that gives science its extra epistemic oomph.

Let me motivate a possible answer to the question what makes empirical

17. What if you are already certain of the correlation between F and G? How do you extract confirmational goodness from your experiment to confer on the F-as-cause hypothesis. Referring all the complex questions about quantifying confirmation to the general theory of inductive support, I give this answer: you are not yet certain of the following hypothesis: *In the absence of T, all Fs are G.* Your experimental data instantiates this hypothesis, so initiating a chain of confirmation that flows on up to the F-as-cause hypothesis.

confirmation better by examining a weakness of non-empirical confirmation. Suppose that Empiricus and Lector develop a taste for metaphysics. Empiricus has has empiricist leanings, and believes that on metaphysical questions, the doctrines of the philosopher David Lewis are usually or always correct. Lector was not admitted to the Princeton Ph.D. program, and so believes that everything written by Lewis is garbage. Both investigators visit the library and begin to read the collected works of Lewis. They accumulate a great deal of observational evidence, that is, knowledge of the various views espoused by Lewis. Although they base their metaphysical views on the same evidence, however, their beliefs diverge as the evidence comes in: wherever there is some metaphysical issue on which both investigators were previously agnostic, and they discover that Lewis's position on the issue is contained in the doctrine p, Empiricus comes to believe p and Lector comes to believe $\neg p$. The more that Lewis insists on p in print, the greater the resulting disagreement between Empiricus and Lector.

If the purpose of philosophy is to acquire knowledge about the world, something has gone badly wrong. Empiricus's and Lector's epistemic states are diverging, and it seems that there is nothing they can do about this divergence as long as they rely on their chosen method of non-empirical inference, namely, reading and learning from the greats. (Of course, metaphysics has other methods that you might hope will resolve the question of Lewis's reliability, and so bring Empiricus and Lector closer together; history, however, provides slender grounds for optimism.)

This example points to a significant feature of a great deal of non-scientific inquiry, I hope you will agree. Though such inquiries typically include inferences based on information acquired by the senses—empirical evidence in the broad sense—the import of the evidence depends heavily on background beliefs. The Lewis case is an extreme, since the background makes the difference between the evidence's supporting a hypothesis and its supporting the hypothesis's negation. More commonly, the background makes the difference between the evidence's supporting and its having no implications for a hypothesis—as when the background concerns the reliability of intuitions of plausibility in philosophy, or the divine inspiration of a text in scriptural hermeneutics. Either way, the background belief acts like a self-contained rule of inductive inference in itself: *From Lewis's asserting that p, infer p*. For this reason, the background beliefs are able to establish almost any inferential relation between observation and hypothesis that you like—or in other words, they are able to inductively connect a given observation to a tremendously wide range of different hypotheses.

Does science, and more particularly empirical testing, avoid this biasing effect of the background? The logical empiricists tended to believe that it does, or at least, that it should and could; thus their search for an objective "logic of justification". Later philosophers of science—names such as Hanson, Kuhn, and Feyerabend spring to mind—were skeptical, discerning a certain circularity or holism in empirical inference that is not so different from that of the Lewis case.[18]

In the matter of direct confirmation—the confirmation of a causal hypothesis by its instances—the logical empiricists appear to have the upper hand. No background belief is required to determine which observation statements instantiate which causal hypotheses. Scientists must therefore agree on the facts of instantiation, and thus on the facts of direct confirmation. The beliefs of a scientific community driven by direct confirmation will for this reason converge, and, because direct confirmation is insulated from defective background beliefs, they will tend to converge on the truth. (I have not laid out all the assumptions required for a convergence argument; it is enough for my purposes here to persuade you that background beliefs do not stand in the way of convergence driven by direct confirmation. There are other impediments to convergence, notably, problems of empirical underdetermination and grue.)

18. Hanson (1958); Kuhn (1962); Feyerabend (1975).

In other words, the instances of a hypothesis, qua instances, speak directly and unequivocally about the hypothesis (as indeed the hypothesis speaks about its instances). They are not interpreted in the light of, or with the aid of, or in the context of, the epistemic background. Thus differences in background cannot lead to disagreement among different inquirers, and flaws in the background cannot lead to falsehood. It is to this directness, this inductive context-independence, that I attribute the special epistemic virtue of direct confirmation.

As far as empirical testing is a matter of direct confirmation it has a very clear advantage over many other kinds of observation-based inductive inference. But that is not very far: direct confirmation cannot take us from the observable to the theoretical, and so cannot span the most profound gap in empirical inquiry. The logical empiricists were prepared to solve this problem by dispensing with the theoretical altogether; they would be left with a purely observational science driven largely or completely by direct confirmation. But this will not satisfy us.

What, then, to say about indirect confirmation? I will focus on what I take to be the central element of indirect confirmation, the confirmation of a causal hypothesis's associated theory, that is, the theory that purports to describe the hypothesis's underlying mechanism.

Indirect confirmation licenses an inference from an instance of a hypothesis to a possible causal explanation of the hypothesis. Because the elements of an associated theory are as explicit about their effects as the causal hypothesis that they putatively explain, there is no role in indirect confirmation for very broad interpretive principles acting, like the Lewisian reliability principle, as inductive laws in themselves. The range of possible causal theories explaining a given hypothesis will be quite limited, the more so the stricter your conception of the causal explanation of regularities. (I develop a strict theory of explanation in Strevens (2008).) Thus the indirect confirmatory connections between evidence and theory are tightly constrained, and furthermore, are a matter of agreement among all those who understand causal language in the same way.

It is possible, nevertheless, for scientists' beliefs about certain theoretical matters to diverge as more evidence arrives. Such divergence is a consequence of the fact—a post-empiricist dogma—that the elements of a hypothesis's associated theory achieve their explanation of the hypothesis in concert rather than singly. This soupçon of holism is sufficient to create a mild version of the Lewis effect, in the following way.

Suppose that Empiricus and Lector receive the same evidence, a stream of positive instances of a causal hypothesis *All Fs are G*. Both scientists agree that the hypothesis is increasingly confirmed. They differ, however, on the likely mechanism explaining the hypothesis, and thus causing the *G*-ness of all those *F*s. Empiricus thinks one mechanism is very likely; Lector that another, quite different mechanism is very likely. Both mechanisms explain the hypothesis equally well. Empiricus regards her mechanism as well confirmed in the light of the evidence; Lector regards his mechanism as equally well confirmed. Suppose that some regularity p must hold for Empiricus's mechanism to operate, but must *not* hold for Lector's mechanism to operate (thus, Lector's mechanism requires an alternative to p). Then though the two scientists might start out agnostic about p, by the time *All Fs are G* becomes highly confirmed, they will disagree: Empiricus believes p on the basis of the evidence, while Lector believes $\neg p$.[19]

In contrast to the case of Lewis scholarship, however, it is possible for Em-

19. You might model the learning process along Bayesian lines. Let h stand for *All Fs are G*. Empiricus starts out with a high subjective probability for p conditional on h; Lector with a correspondingly low probability. Neither thinks that h is at all likely, so their differing conditional probabilities make very little contribution to their overall subjective probabilities for p, which might be roughly equal. As the evidence for h comes in, the conditional probability starts to contribute more and more of the total probability for p; eventually it dominates. The conditional probability itself does not change. Thus, after h is sufficiently confirmed, Empiricus will have gone from being agnostic to having a high subjective probability for p, while Lector will have gone from the same state to having a high subjective probability for $\neg p$.

piricus and Lector to resolve their differences through more of the same sort of inference, that is, through empirical testing. The associated theories that they hold in such differing regards themselves consist of causal hypotheses. (They also contain what I referred to above as "basing generalizations", a fact that I will conveniently ignore here.) They may thus be put to the test, either directly or indirectly.

A direct test looks for instances, positive or negative, of the causal hypotheses that constitute the two theories (perhaps including *p* itself). An indirect test is a kind of triangulation, testing other causal hypotheses that are thought to share some elements of the putative mechanism that is in dispute. There are, of course, methodological complications to all of this; I do not claim that resolving Empiricus's and Lector's disagreement will be straightforward. My point is rather that it can be resolved in principle by the same kind of inference that exposed the disagreement: the differing "background" beliefs that drive the divergence (in this case, Empiricus's and Lector's differing opinions on the most likely mechanism by which *F*s cause *G*-ness) can themselves be put to the test empirically. By contrast, differences of opinion concerning the reliability of Lewis's and other metaphysicians' beliefs cannot be resolved by consulting the writings of the metaphysicians themselves. And this is not an accidental feature of a toy example, but a familiar and persistent problem in philosophy, revealed religion, literary theory, and so on.

It is of course possible to imagine an unfriendly world in which both direct testing and triangulation are impossible; in such a world the predicament of the empirical scientist is closer to that of the metaphysician and the hermeneuticist. Perhaps there is such an unfriendly stretch in our own and in many other generally epistemologically genial worlds, namely in fundamental physics, the arena in which battles about the underdetermination of theory by evidence are almost always fought. But it would be a mistake, in the kind of inquiry I am undertaking here—an assessment of the qualities of various kinds of inductive inference—to equate epistemic virtue and happiness, to insist that every vindication be a proof of the inevitability of success. The problems of induction are not so easily solved.

Let me sum up. Empirical confirmation offers two advantages over the sort of observation-driven inference typified by the Lewis case. First, what counts as a good inductive inference is not relative to a set of background beliefs that act as protean inductive norms in their own right. Second, what divergence empirical confirmation allows because of differences in epistemic context can itself be resolved by empirical testing.

The relative virtues attributed to empirical confirmation in this section do not explicitly invoke its being a kind of causal inference. Might other, non-causal varieties of inductive inference have some of the same qualities? This seems to me to be entirely possible. Whether these other kinds of inference are characteristic of our science, I do not know.[20] In short: causal instantialism may be one among several kinds of empirical confirmation—though it predominates in science as we know it.

20. The testing of basing generalizations would be somewhere to look.

References

Feyerabend, P. K. (1975). *Against Method*. Verso, London.

Good, I. J. (1967). The white shoe is a red herring. *British Journal for the Philosophy of Science* 17:322.

Hanson, N. R. (1958). *Patterns of Discovery*. Cambridge University Press, Cambridge.

Hempel, C. G. (1945). Studies in the logic of confirmation. *Mind* 54:1–26, 97–121.

———. (1965). Empiricist criteria of cognitive significance: Problems and changes. In *Aspects of Scientific Explanation*, chap. 4. Free Press, New York.

Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. Chicago University Press, Chicago.

Strevens, M. (2000). The essentialist aspect of naive theories. *Cognition* 74:149–75.

———. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.