

International Phenomenological Society

Why Are the Laws of Nature So Important to Science?

Author(s): Marc Lange

Source: *Philosophy and Phenomenological Research*, Vol. 59, No. 3 (Sep., 1999), pp. 625-652

Published by: [International Phenomenological Society](#)

Stable URL: <http://www.jstor.org/stable/2653785>

Accessed: 13/11/2013 00:46

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



International Phenomenological Society is collaborating with JSTOR to digitize, preserve and extend access to *Philosophy and Phenomenological Research*.

<http://www.jstor.org>

Why Are the Laws of Nature So Important to Science?

MARC LANGE

University of Washington

Why should science be so interested in discovering whether p is a law over and above whether p is true? The answer may involve the laws' relation to counterfactuals: p is a law iff p would still have obtained under any counterfactual supposition that is consistent with the laws. But unless we already understand why science is especially concerned with the laws, we cannot explain why science is especially interested in what would have happened under those counterfactual suppositions consistent with the laws. It is argued that the laws form the only non-trivially "stable" set, where "stability" is invariance under a certain range of counterfactual suppositions not itself defined by reference to the laws. It is then explained why science should be so interested in identifying a non-trivially "stable" set: because of stability's relation to the best set of "inductive strategies".

Praise the Lord, for he hath spoken
Worlds his mighty voice obeyed;
Laws, which never shall be broken,
For their guidance he hath made.

English Hymnal (1796),
Foundling Hospital Collection, no. 535

1. Introduction

The laws of nature *govern* the universe. For instance, a given emerald's color is *governed* by the law that all emeralds are green, comets *follow* the same laws of motion as the planets *obey*, and gravity was believed to act *in accordance with* (or *under*) Newton's inverse-square law. The laws governing various phenomena may themselves be governed by higher-order laws. For instance, it is a law that all laws of motion are Lorentz-invariant, and it is a law that for each kind of elementary particle, there is a law specifying its rest mass. So we have a hierarchy:

⋮
LAWS
that govern

LAWS
that govern
 LAWS
 ⋮

At the bottom of this hierarchy are the facts about the actual world (i.e., *not* counterfactual conditionals) that are governed by laws but do not themselves govern anything (or describe what governs something). These are the “non-nomic facts”—e.g., the facts about particles in motion that are governed by the laws of *motion*, as contrasted with the facts that are governed by laws *about laws*. Whereas $\blacksquare p$ (“It is a law that p ”), $\neg\Box q$ (“It is not physically necessary¹ that q ”), and $(r \supset \blacksquare r)$ are (if true) nomic facts, the non-nomic facts include that all emeralds are green (a law) and that all gold objects are smaller than one cubic mile (an accidental generalization). A “non-nomic claim” (a claim that, if true, expresses a non-nomic fact) does not concern the laws; it is not made true or false by which counterfactual conditionals are correct, and its truth-value does not depend in any obvious way on whether or not some fact is a law (or physically necessary).² Let U be a language for science containing exactly the non-nomic claims; identify a language with the set of its sentences.

I take for granted science’s interest in discovering the non-nomic facts. But for $p \in U$, why should science care about discovering p ’s lawhood insofar as this goes beyond discovering p ’s truth? By discovering a regularity to be a matter of natural law, we learn nothing more about the non-nomic facts than we already knew simply from discovering the regularity itself.³ Why, then,

¹ As I use the term “physically necessary”, $\Box h$ if and only if h holds in every logically possible world with *exactly* the same laws as the actual world—that is, if and only if h follows from the laws’ lawhood and the non-laws’ non-lawhood.

² Although I presume that “Humean regularities” (e.g., that all emeralds are green) are non-nomic, I am not prepared to say whether a fact concerning causal relations, explanatory relations, or objective chances is non-nomic, since this would require analyses of these difficult concepts, which it is not my aim here to provide. Facts concerning causal relations, explanatory relations, and objective chances can perhaps be physically necessary or accidental. That each of the atoms in this vial has a 50% chance of decaying in the next 8.1 days is, presumably, an accident, whereas it is a law that each atom of iodine 131 has such a chance. Likewise, that each of my mishaps while driving a car was caused by carelessness is presumably an accidental generalization, whereas it is perhaps a law that every bolt of lightning causes a clap of thunder.

³ To defend this claim properly, I would have to argue that the distribution of \blacksquare ’s and $\neg\blacksquare$ ’s over the truths in U fails to supervene on those truths. This seems to me especially evident in the case of a highly impoverished possible world. For example, consider a world in which nothing happens in its entire history except that a single lonely electron moves uniformly forever. Intuitively, the actual laws could hold in this world (dictating what the force between the given electron and another electron would have been, had there been another electron), or the laws could make the gravitational force twice as strong as in actuality. In contrast, consider Lewis’s account (1983) of law, according to which the laws are (roughly) the generalizations in the logical system of truths having the

should science concern itself with discovering whether some regularity is accidental or a matter of natural law? Of course, if science's goal is *stipulated* to be the discovery of *all* facts, including nomic facts, then science must be interested in discovering the natural laws. But this answer merely begs the question; it cannot persuade anyone who did not already accept that the natural laws matter to science.

One way to explain science's interest in discovering which truths in U are laws is to grant that science is interested in non-nomic facts *in a broad sense* (broader than U): not only in what actually happens, as far as facts in U are concerned, but also in what (non-nomically) would have happened—that is, in which subjunctive conditionals (with antecedents and consequents in U) hold.⁴ Call these the facts in U* (which is U supplemented by >, where “*p* > *q*” is the subjunctive conditional “Were *p* the case, then *q* would be the case”). Let Λ be the set of truths in U that are laws along with their logical consequences in U.⁵ Laws and counterfactuals are related by the following principle:

best combination of simplicity and informativeness regarding the non-nomic truths. On this view, the laws supervene on the non-nomic facts; in a lonely-electron world, it is presumably a law that all particles are electrons and not a law that all copper objects are electrically conductive. This seems counterintuitive. Having set exactly the actual laws, God (to put the point colorfully) could have arranged the initial conditions so as to generate nothing in history but a single electron. See also Carroll (1994, pp. 60ff.).

⁴ An alternative suggestion for explaining science's interest in the laws is that unless we believe that it is a law that all F's are G, we cannot be justified in believing that all F's are G while believing that there remain F's that we have not yet examined for their G-ness. In other words, we cannot justly project the pattern exhibited by our observations (every examined F has been found to be G) onto unexamined cases if we believe this pattern to be coincidental rather than perhaps physically necessary. Such a view was famously defended by Goodman (1983, pp. 20–22, 73), and has likewise been advanced by J. S. Mill (1893, p. 230; Bk. III, ch. 4, sect. 1), C. S. Peirce (1934, p. 66), G. E. Moore (1962, p. 12), R. B. Braithwaite (1927, pp. 467 and 473), Hans Reichenbach (1947, pp. 359ff. and 368), William Kneale (1952, esp. pp. 45, 52, and 65), P. F. Strawson (1952, pp. 199f.), J. L. Mackie (1962, pp. 71–73), Israel Scheffler (1981, pp. 225ff.), and F. I. Dretske (1977, pp. 256–60). But this view is mistaken, as has been emphasized recently by Elliott Sober (1988, pp. 95ff.) and Bas van Fraassen (1989, pp. 134ff., 163ff.). I shall go into this at the start of section III. (There I consider an example in which we become justified in believing that all of the pears now on my tree happen to be ripe, though we know that we have examined the ripeness of only some of the pears now on the tree.) Still, I shall ultimately argue that this view contains a key grain of truth: there is an important kind of confirmation (which I shall call “inductive” confirmation) that we can make available only to those hypotheses that we believe may express physical necessities. One way to see this paper (and my forthcoming-b) is as finally elaborating the sort of “projection” that Goodman *et al.* had in mind, but failed to specify properly, thereby leaving themselves vulnerable to the criticisms of van Fraassen and Sober. In order to explain what distinguishes “inductive” confirmation, why it is important, and why its availability is related to whether the hypothesis may express a physical necessity, I must first discuss the relation of laws to counterfactuals, which I am about to do.

⁵ In my (forthcoming-a,b), I argue that Λ contains exactly the facts *p* in U where $\Box p$.

Λ is preserved under every subjunctive supposition in U that is consistent with Λ ,

where a set Γ of truths is “preserved under p ” exactly when $p > q$ holds for any $q \in \Gamma$.⁶

Counterfactual conditionals are notoriously context dependent. (For example, as Bennett (1984, p. 71) remarks, our concerns influence which of these counterfactuals we would be correct in asserting: “Had I jumped from this window, I would have suffered serious injury”, “...I would have arranged for a net to be in place below, and so I would not have suffered serious injury”, “...the window would have to have been much closer to the ground.”) Principles like the above are intended to cover *all* contexts, since they are supposed to be logical truths—to reflect what it is for a claim to express a natural law in a given possible world. So the above principle demands that *in any context*, $(p > m)$ is correct for any $m \in \Lambda$ and any $p \in U$ that is consistent with Λ .

This principle explains how science can discover that some counterfactual conditional is correct by discovering that a given non-nomic fact is physically necessary. For example, by discovering that it is a law that all copper objects are electrically conductive, we may ascertain (if we believe it physically possible for me to hold a copper object in my hand) that were I holding a copper object in my hand, then I would be holding an electrical conductor. So having already discovered that r (where $r \in U$), science remains *interested* in discovering that $\Box r$.

However, this account does not explain why it is so *important* in science to discover that $\Box r$ over and above discovering r (where $r \in U$). Admittedly, the members of Λ are the *only* non-nomic claims that are preserved under every $p \in U$ that is consistent with Λ . (This is easily shown: If $q \notin \Lambda$ because q is false, then any truth $p \in U$ is consistent with Λ , but $\neg(p > q)$. And if $q \notin \Lambda$ and q is a truth in U , then $\neg q$ is consistent with Λ (since otherwise q follows from members of Λ , so $q \in \Lambda$, contrary to our stipulation), so if we make $p = \neg q$, we again have $\neg(p > q)$ where p is consistent with Λ .) So to explain not merely the scientific *relevance*, but also the special scientific *importance* of identifying Λ 's membership (insofar as this goes beyond discovering truths in U), we must explain why it is especially important to science that some truth in U is preserved under every $p \in U$ that is consistent with Λ . But why is preservation under *this* particular set of subjunctive antecedents in U more important than preservation under some *other* set of

⁶ Among the many who have made similar suggestions are Bennett (1984), Carroll (1994), Chisholm (1946; 1955), Goodman (1947; 1983), Horwich (1987), Jackson (1977), Mackie (1962), Pollock (1976), and Strawson (1952). While I shall take this view as my starting point in this paper, it actually requires very careful defense against a variety of challenges. (Lewis is a notable dissenter.) See my forthcoming-b.

subjunctive antecedents in U ? If we cannot say why, then we cannot explain why identifying the natural laws is more important to science than identifying the truths that are preserved under some other, arbitrarily chosen set of subjunctive antecedents. (After all, no special title is bestowed upon those truths that are preserved under every subjunctive antecedent logically consistent with, say, “George Washington was the first President of the United States”; such preservation lends a truth no special importance.) On the other hand, if we say that preservation under every $p \in U$ that is consistent with Λ is especially important because science especially cares about counterfactual suppositions that are physically possible, then we have just traversed a very tight circle: the laws are important because they are preserved under a certain range of counterfactual suppositions, where this range is important because it is delimited by the laws and the laws are important.

The purpose of this paper is to explain why the truths of the form $\Box r$ and $\neg\Box r$, insofar as they go beyond the truths in U , are so important to science. In section 2, I shall present a respect in which Λ is distinguished from all other logically closed sets of truths in U . I shall not use the concept of a natural law in specifying the property by which this set is distinguished. I can then explain, in a non-question-begging way, why it is especially important for science to identify the laws—namely, by showing (in sections 3 and 4) why science must take a special interest in identifying the claims that form the set possessing this special property.

Of course, I would be begging the question if I explained the particular scientific significance of this property by appealing to the scientific importance of identifying the members of Λ (insofar as this goes beyond discovering truths in U). It will turn out that the reason it is so important for science to discover which set of claims possesses this property is that science is interested in *predicting* truths in U^* —getting to know them in advance of having directly observed them to obtain. (In the case of truths in U^* that concern what *would have* happened under various *counterfactual* circumstances, their observation is out of the question; any knowledge of them must be a “prediction”.) How can we successfully make such predictions in a case where (as I shall explain) *our prior opinions fail to supply us with a good reason to regard any observation as confirming any such prediction?* To discover how is (as it turns out) to discover which set of claims in U possesses the special property distinguishing Λ . To seek the best way to arrive at such predictions “starting from scratch” is to seek the natural laws.

I do not mean to suggest that this is necessarily the only reason that the truths of the form $\Box r$ or $\neg\Box r$ (where $r \in U$), insofar as they go beyond the truths in U , are important for science to discover. Whereas I took for granted science’s interest in discovering the correctness of various counterfactual conditionals with non-nomic antecedents and consequents, another philosopher might just as well begin by presuming science to be interested not merely in

predicting but in explaining why the non-nomic facts obtain. But any account of the laws' value to science risks traversing a very tight circle in the same manner as our initial appeal to counterfactuals did. For instance, if we say that a non-nomic fact's lawhood has special scientific importance because science is interested in explaining the non-nomic facts and the laws have a special explanatory power, then we must say *why* the laws have this power. If we say that a scientific explanation essentially involves subsumption under the laws, then in attributing the laws' importance to their explanatory power, we have really done little more than *stipulate* the laws' importance; we have not explained why science should especially care about subsuming non-nomic facts under laws, just as the earlier proposal fails to explain why science should especially care about counterfactual conditionals in U^* with antecedents consistent with Λ . To explain what is so special about that particular range of counterfactual suppositions, or what it is about laws that makes it explanatory (and hence important) for us to subsume non-nomic facts under *them* in particular, we need to specify a relation to the non-nomic facts (in the broad sense) that is borne uniquely by the laws and whose scientific importance can be appreciated without presupposing that we already (for some reason) take special interest in the laws. Such a property has (to my knowledge) never before been identified. That is what I now propose to do.

2. Non-Nomic Stability

The set Λ is preserved under every $p \in U$ that is consistent with Λ , whereas for any accidental truth $r \in U$, no set of truths in U containing r is preserved under every such p . But as we have just seen, this fact cannot help to explain the scientific importance of Λ without begging the question, since Λ itself is used to fix the range of p 's where preservation under p is at issue. This seems like rigging the game. If we allowed some accidental truths to fix the relevant range of subjunctive antecedents, then we might well find that Λ fails to be preserved under some relevant p . Would we find that some $\Gamma \subset U$ of truths that contains an accident is preserved under every relevant p ?

Let's avoid using Λ to fix the relevant range of subjunctive antecedents. Let's say that a set Γ has "non-nomic stability" exactly when

- (i) $\Gamma \subseteq U$;
- (ii) For any $q \in U$: If $q \in \Gamma$, then q is true;
- (iii) For any $q \in U$: If Γ logically entails q (i.e., if q is entailed by some members of Γ), then $q \in \Gamma$; and

- (iv) For any $p \in U$ that is consistent with Γ and for any $q \in \Gamma$: $p > q$ is true.⁷

That is, a set of non-nomic truths that is logically closed in U is “non-nomically stable” exactly when it is preserved whenever (non-nomically speaking) it *could* be preserved—that is, under any subjunctive antecedent (in U) that is consistent with each of its members. When discussing a set’s “non-nomic stability”, then, *the set itself* determines the range of p ’s where preservation under p is at issue.

We have already seen that Λ is non-nomically stable. Two other non-nomically stable sets are the set of logical truths in U and the set of all truths in U . These two are the sets that possess non-nomic stability *trivially*. The set of logical truths in U is stable because for any p and any logical truth q in U , trivially p logically entails q , and so automatically $p > q$ (at least where p is consistent with the logical truths). The set of all truths in U is preserved under each subjunctive antecedent in U consistent with each of its members because $p > q$ is trivially true if p and q are true, and no false p in U is consistent with each member of the set.

Now I will argue that the only non-nomically stable set besides these two—in other words, the only *non-trivially* non-nomically stable set—is Λ . Non-trivial non-nomic stability is the property distinguishing Λ that I promised in section 1.

Let’s begin with a lemma. Suppose that both Γ and Γ' are non-nomically stable, and neither is a subset of the other. Take any

$$r \in \Gamma, r \notin \Gamma', r' \in \Gamma', r' \notin \Gamma.$$

Then

$$\neg r \vee \neg r' \text{ is consistent with } \Gamma,$$

since otherwise, some members of Γ must logically entail $\neg(\neg r \vee \neg r')$, i.e., (r & r'), but then (from (iii) in the definition of “non-nomic stability”, recalling that Γ is non-nomically stable) $r' \in \Gamma$, contrary to our initial assumption. Likewise

$$\neg r \vee \neg r' \text{ is consistent with } \Gamma'.$$

Therefore, by the stability of Γ and Γ' , respectively, it follows that

$$(\neg r \vee \neg r') > r, (\neg r \vee \neg r') > r'.$$

⁷ Actually, (iv) renders (ii) superfluous: Let p be a logical truth in U . Then (iv) requires that for any $q \in \Gamma$: $p > q$ is true. Since p is a logical truth, ($p > q$) iff q .

But these are mutually inconsistent, since the former implies

$$(\neg r \vee \neg r') > \neg r'.$$

We have, then, a *reductio*. We have shown that if there are two distinct, non-nomically stable sets, then one must be a proper subset of the other. Since Λ is non-nomically stable, it follows that if there is any other non-trivially non-nomically stable set Γ , then either $\Lambda \subset \Gamma$ or $\Gamma \subset \Lambda$. By eliminating these two possibilities, I shall show that Λ is the only non-trivially non-nomically stable set.

Let's first consider the case where $\Lambda \subset \Gamma$; let Γ satisfy requirements (i), (ii), and (iii) in the definition of "non-nomic stability". (That is, Γ is a set of non-nomic truths, logically closed in U .) So Γ contains some truth $a \in U$ such that $a \notin \Lambda$ —in other words, some accidental truth a . To see why Γ fails to be stable, let's begin by considering an example: let Γ be the logical closure in U of "All of the matches now in this book remain forever unlit" (the accident a in this example, after Goodman 1983) together with all of the truths $p \in U$ where $\Box p$. Suppose that all of the matches in the book are dry and well-made, oxygen is present, and so on, but as it happens, none of the matches is ever struck. Consider the counterfactual antecedent p : "Had one of them been struck". Now p is consistent with Γ , since the laws of nature plus the fact that one of these matches is struck does not logically entail that one of them lights; for the match to light, oxygen must also be present, the match must be dry and well-made, and so on—which is not contained in Γ . Since p is consistent with Γ , Γ is non-nomically stable only if $p > a$, i.e., only if all of the matches would still have remained unlit even if one of them had been struck. But when standard conditions prevail, then (in an ordinary conversational context, it is correct to say that) had one of the matches been struck, oxygen would still have been present, the matches would still have been dry and well-made, etc., and so the match would have lit— $\neg(p > a)$. Hence, Γ lacks non-nomic stability.

We could have shown this result in a different way. Let b be some accidental truth that is unrelated to a , such as "All gold cubes are smaller than one cubic mile" (after Reichenbach 1947, p. 368). Consider $(\neg a \vee \neg b)$ as a counterfactual antecedent: "Had either one of the matches in the book been lit or there been a gold cube exceeding one cubic mile". Now $(\neg a \vee \neg b)$ is consistent with Γ . But in a great many contexts, we would be correct in denying $(\neg a \vee \neg b) > a$, though this counterfactual is required by Γ 's non-nomic stability—since we would be correct in denying $(\neg a \vee \neg b) > \neg b$ ("Had one of the matches in the book been lit or had there been a gold cube exceeding one cubic mile, then there would have been a gold cube exceeding one cubic mile"). Indeed, in a great many contexts, I daresay we would be correct in denying $(\neg a \vee \neg b) > \neg b$ and correct in denying $(\neg a \vee \neg b) > \neg a$.

This argument can be generalized to show the non-nomic instability of any set Γ where $\Lambda \subset \Gamma$ and (to preclude Γ 's *trivial* stability) b is a truth in U where $b \notin \Gamma$. (It follows that b is accidental.) If Γ satisfies requirements (i), (ii), and (iii) in the definition of “non-nomic stability”, then there is a claim a such that $p \in \Gamma$ just in case p follows from Λ and a . (It follows that a is an accidental truth.) Now $(\neg a \vee \neg b)$ is consistent with Γ , since otherwise some members of Γ must logically entail $\neg(\neg a \vee \neg b)$, i.e., $(a \& b)$, and so (by (iii) in the definition of “non-nomic stability”) $b \in \Gamma$, contrary to our supposition. So Γ 's non-nomic stability requires that $(\neg a \vee \neg b) > a$, since $a \in \Gamma$, and hence requires that $(\neg a \vee \neg b) > \neg b$. It requires, in other words, that under this counterfactual antecedent, b should always be sacrificed for the sake of preserving a . But if there are some contexts in which we would be correct in doing this, there are other contexts in which (at least for some such b) we would be correct in denying this counterfactual and instead asserting $(\neg a \vee \neg b) > \neg a$, because a closer possible world is reached by sacrificing a . (The sacrifice of a may not succeed in preserving b , as when b logically entails a .) In addition, there may well be contexts in which neither of these counterfactuals can be correctly asserted. Since there are contexts in which a is not preserved under a counterfactual antecedent that is consistent with Γ , Γ is non-nomically unstable. In short, Γ is non-nomically stable only if a 's preservation is more important than b 's in *every* conversational context, for any $b \notin \Gamma$ —which is highly implausible.⁸

A similar argument suggests the non-nomic instability of any set Γ where $\Gamma \subset \Lambda$ and (to preclude Γ 's *trivial* stability) Γ contains some logically contingent truth. Consider a counterfactual antecedent in U that is consistent with every member of Γ , though inconsistent with some member of Λ —so Λ 's non-nomic stability does not require Γ 's preservation under this counterfactual antecedent. Then Γ will not be preserved under this counterfactual antecedent. Let $a \in \Lambda$ where $a \notin \Gamma$. Let $b \in \Gamma$ where b is contingent. Now $(\neg a \vee \neg b)$ is consistent with Γ , since otherwise some members of Γ must logically entail $(a \& b)$, and so (since Γ must be closed in U or else lack non-nomic stability) $a \in \Gamma$, contrary to our stipulation. Hence, Γ 's non-nomic stability requires that $(\neg a \vee \neg b) > b$, since $b \in \Gamma$, and so requires that $(\neg a \vee$

⁸ We could put this claim to the test by selecting as our a an accidental truth that is preserved under a tremendously broad range of counterfactual antecedents consistent with it, the members of Λ , and their logical consequences. For instance, let a be “Sometime in the history of the universe, there exists some matter.” It is perhaps initially difficult to find any counterfactual antecedent consistent with a , the members of Λ , and their logical consequences, under which a might not be preserved. But let b be “The energy of the universe is insufficient to return the universe to a Big Crunch in much less than 15 billion years [the current age of the universe].” I see no reason to say that $(\neg a \vee \neg b) > (a$ and $\neg b)$ —that is, “Had there either been no matter ever or else so much energy as to close the universe in much less than 15 billion years, then there would still have been matter sometime, and there would have been sufficient energy to close the universe so soon.”

$\neg b) > \neg a$ —in other words, requires that under this counterfactual antecedent, a should always be sacrificed for the sake of preserving b . But it is implausible that b 's preservation takes precedence over a 's in every context, for any such b . For example, suppose a is that “Hooke’s law” is true and b is that “Snell’s law” is true. It is not at all plausible that in every context, it is correct to assert that had either “Hooke’s law” or “Snell’s law” been false, then Snell’s would still have held and Hooke’s would have been false.⁹ In many contexts, both this counterfactual and “...Hooke’s would still have held and Snell’s would have been false” are rightly denied.

So non-trivial non-nomic stability is to be found neither in a proper subset of Λ nor in a superset of Λ ; Λ is the only logically closed set of non-nomic truths that is non-trivially preserved whenever—non-nomically speaking—it logically possibly could be: under every non-nomic counterfactual antecedent that is logically consistent with each of its members.¹⁰ The range of non-nomic counterfactual antecedents, under which the members of some logically closed set of non-nomic truths must be preserved for that set to be non-nomically stable, is fixed not by appealing to nomic concepts, but rather entirely by the set’s members themselves. (This is what we originally sought, and will be crucial to the argument of the following section.) But we have yet to see why science should especially care about identifying the members of the only set that non-trivially possesses non-nomic stability.

3. Inductive Confirmation

To see why, I must first highlight a familiar feature of confirmation: When some evidence confirms a hypothesis (i.e., justly raises—by some increment—our confidence in its truth) and that hypothesis makes a prediction, then the evidence need not confirm that prediction.¹¹ To reinforce this point,

⁹ Of course, what to make of a counterfactual like this (a “counterlegal”) is a vexing issue on which reasonable people differ. Those who regard the natural laws as conceptual necessities will take every counterlegal as trivial (though they may disagree on whether they are all vacuously true or all trivially false). Obviously, I am working here to elaborate the idea that physical necessity constitutes a “grade” of necessity between logical (or conceptual) necessity and no necessity at all. (See my forthcoming-a,b.) Accordingly, I do not take all counterlegals as trivial.

¹⁰ Another version of this argument (differing in some details that are inessential for my purposes here) appears in my (forthcoming-a). There I use it to explain the sense in which the laws have a kind of “necessity” but an accident does not—even if that accident is preserved under a broad range of counterfactual suppositions. I explicate the sense in which physical necessity is “between” logical or conceptual necessity and no necessity at all, and what it would be for *multiple* “grades” of necessity to lie between these extremes. Since that paper is not concerned with explaining the laws’ scientific importance, it does not include any version of the argument in sections 3–5 below.

¹¹ This point was made forcefully by Carnap (1950/1962, pp. 462ff.), though he did not elaborate it as I shall now do—by regarding certain counterfactual conditionals as just like predictions regarding actual cases. In this respect, my approach is in the spirit of Goodman (1983)—see note 4.

consider four brief examples, each involving the confirmation of a hypothesis, but differing in the *range* of its predictions that are confirmed.

1. We believe that a given die is fair. The hypothesis is that each of its next three tosses lands six. Accordingly, our initial confidence in the hypothesis is $1/6 \times 1/6 \times 1/6 = 1/216$. We toss the die once, and it lands six. Since this evidence eliminates one way in which the hypothesis might have been falsified, the hypothesis is confirmed; our confidence in it is raised to $1/6 \times 1/6 = 1/36$. But none of its predictions regarding actual unexamined cases (the outcomes of either of the next two tosses) is confirmed.¹²

2. The hypothesis is that all emeralds are green and all rubies are red. The first emerald we examine turns out to be green. This would typically be relevant, confirmation-wise, to unexamined emeralds but not to unexamined rubies. In other words, this evidence would typically confirm the hypothesis and, unlike the previous example, would confirm *some* of the *predictions* made by the hypothesis (such as that the next emerald I check will be green). But typically, there are other predictions that it would fail to confirm (such as that the next ruby I check will be red).

3. The hypothesis is that all of the pears on this tree are now ripe. By checking a pear from the tree and finding it to be ripe, we might confirm, of each actual unexamined pear on the tree, that it is ripe. In that event, the evidence not only confirms the hypothesis but also, unlike the previous example, confirms *each* of the predictions that the hypothesis makes regarding *actual* unexamined cases. Furthermore, the evidence typically also confirms the ripeness of certain *counterfactual* pears. For instance, it confirms that had there been a pear on the third branch of the tree, then it would now have been ripe. But the evidence typically fails to confirm the predictions that the hypothesis makes concerning certain other counterfactual cases—e.g., that had there been a pear on the tree whose environmental conditions (e.g., temperature, length of day, plant hormones) were experimentally manipulated to differ from those actually experienced by the pears on the tree, then it would now also have been ripe.

4. Newton regarded a successful prediction made by his putative gravitational-force law as bearing upon each of the predictions (that had not already been accepted) that the hypothetical law makes regarding actual unexamined cases. Moreover, unlike the preceding example, Newton took this evidence as bearing upon *each* of the predictions (that had not already been accepted) that the hypothesis makes regarding cases that do not actually exist but (roughly speaking) *might have happened to exist*—e.g., as confirming that no matter what the Earth-Moon separation might have been, their mutual gravitational attraction would still have accorded with his hypothesis.

¹² Ken Gemes refers to confirmation of this sort as “mere content cutting.”

Thus, when we confirm hypothesis h where $h \in U$, we may confirm a broader or narrower range of h 's predictions regarding actual and counterfactual states of affairs. Let me define more explicitly what I mean by h 's "predictions": they are all and only the subjunctive conditionals $p > q$, where $p, q \in U$, such that $p \& h$ logically entails q . A prediction $p > q$ made by h specifies what h says would happen if p were the case; it is a prediction "regarding p ".

Now consider some set $\Gamma = \{h_1, h_2, h_3, \dots\}$ of hypotheses $h_i \in U$; let $Cl(\Gamma)$ be its logical closure in U . When we discover that some prediction $p' > q'$ made by a member h of Γ is borne out, what other predictions $p > q$ made by h are we willing to take this discovery as confirming? The answer differs in different cases, as examples 1–4 illustrate. Suppose that for any $h \in \Gamma$, we are willing to confirm each of h 's predictions in a very broad range—as Newton did in example 4. Specifically: suppose that for any $h \in \Gamma$, we are prepared to regard the discovery of any $p' > q'$ predicted by h as confirming each of h 's predictions $p > q$ that satisfies the following two constraints:

- (i) our subjective $\text{pr}(p > q) \neq 1$ (else we could not confirm $p > q$; our confidence in it cannot rise beyond 1)
- (ii) it is logically possible for all of the predictions regarding p that are made by Γ 's members to be true—that is: p is consistent with $Cl(\Gamma)$.

I am *not* claiming that for *any* set Γ of hypothesis, we are willing to pursue this policy. Nor am I claiming that there is always *some* set Γ of hypotheses that we are prepared to confirm in this special manner. For the moment, I am simply imagining that as a matter of fact, there is a certain set Γ of hypotheses that we are prepared to confirm in this manner—in which case, let's say that we are prepared to regard any successful prediction made by any $h \in \Gamma$ as confirming h *inductively*.

The label "inductive" emphasizes what is distinctive about such confirmation: that when h is so confirmed, each of its predictions regarding each circumstance p in a certain broad range is confirmed. This is why inductive confirmation is like the confirmation in example 4, rather than the confirmation in examples 1–3. It is perhaps intuitively more perspicuous to express "inductive" confirmation in terms of Goodman's (1983) notion of "projection".¹³ To be willing to "project" a given h onto some actual or counterfactual circumstance p (e.g., to be willing to project "All emeralds are green" onto the circumstance where I hold an emerald in my hand) is to be willing to take a successful prediction made by h as confirming each of h 's

¹³ Indeed, inductive projection is the kind of confirmation to which I alluded in note 4. This will become clearer in the next section.

predictions $p > q$ regarding that circumstance (e.g., as confirming that were there an emerald in my hand, it would be green). When we take a hypothesis as confirmed by a successful prediction, we may project it over a broader or narrower range of p 's, as examples 1–4 illustrate. Intuitively, “inductive” confirmation involves projection over a very broad range. In particular, when we confirm $h \in \Gamma$ inductively, we project h as far out onto p 's as it can go without running up against the projections of the other h_i (or bumping into its own projection)—in the sense that we confirm each of h 's predictions $p > q$ (where our $\text{pr}(p > q) \neq 1$) for any p where the predictions regarding p that are made by the various h_i are mutually consistent. That is, we are willing to project *each* of the h_i as far onto p 's as they can *all* logically possibly go.

From the definition of “non-nomic stability”, it follows immediately that

It is non-trivially the case that every prediction $p > q$ made by Γ 's members, regarding each p that is consistent with $Cl(\Gamma)$, is true

exactly when

$Cl(\Gamma)$ non-trivially possesses non-nomic stability.¹⁴

Thus, $Cl(\Gamma)$'s non-trivial non-nomic stability is a necessary and sufficient condition for the non-trivial truth of all of the predictions, made by Γ 's members, that are confirmed when Γ 's members are all confirmed inductively.¹⁵ But we saw in the previous section that $Cl(\Gamma)$ non-trivially

¹⁴ Let's do this slowly. First, let's show that if $Cl(\Gamma)$ is non-nomically stable, then every prediction $p > q$ made by Γ 's members, regarding each p that is consistent with $Cl(\Gamma)$, is true. Consider any such prediction; let it be made by $h \in \Gamma$. By the definition I have given of “prediction”, $p \in U$ and $p \ \& \ h$ logically entails q . So $p > q$ holds if $p > h$ holds. And if $Cl(\Gamma)$ is non-nomically stable, then by (iv) in the definition of “non-nomic stability”, $p > h$ holds.

Now let's show the reverse direction. The set $Cl(\Gamma)$ is stipulated as satisfying (i), (ii), and (iii) in the definition of “non-nomic stability”. Suppose that every prediction $p > q$ holds that is made by Γ 's members where p is consistent with $Cl(\Gamma)$. Then in particular for any such p and any $h \in \Gamma$, $p > h$ holds. Hence, $Cl(\Gamma)$ satisfies (iv) in the definition of “non-nomic stability”.

It is *trivially* the case that every prediction $p > q$ holds that is made by Γ 's members, regarding each p that is consistent with $Cl(\Gamma)$, exactly when (a) $Cl(\Gamma)$ is the set of logical truths in U , since then every prediction is a logical truth, or (b) $Cl(\Gamma)$ is the set of all truths in U , since then p must be true (since no falsehood in U is consistent with $Cl(\Gamma)$) and $p > q$ is trivially true if p and q are true. This is exactly when $Cl(\Gamma)$ *trivially* possesses non-nomic stability.

¹⁵ Actually, this is a slight oversimplification: I have ignored the fact that we can confirm a prediction $p > q$ only if the subjective probability we assign to it, prior to confirming it, is not equal to 1. Notice, then, that even if some of the predictions $p > q$ made by $h \in \Gamma$, where p is consistent with $Cl(\Gamma)$, are false, it can happen that each of h 's predictions that we confirm, in confirming h inductively, is true—so long as we (mistakenly) already assign the maximal degree of confidence to each of the (unbeknownst to us) false predictions. When we are willing to confirm each of the h_i inductively (if the right evidence

possesses non-nomic stability if and only if $Cl(\Gamma)$ is Λ , i.e., if and only if the h_i span exactly the logical consequences of the laws in U . It follows, then, that

It is non-trivially the case that every prediction that is confirmed, when all of the h_i are confirmed inductively, is true

just in case

$Cl(\Gamma)$ is Λ .

Intuitively, Λ is the only logically closed set of non-nomic claims such that when each of the claims in the set is projected inductively—i.e., is projected as far onto p 's as the claims in the set can *all* logically possibly go—then non-trivially, each of the predictions thereby confirmed is true.

Perhaps inductive projection has some particular importance in science. If so (and I shall pursue this possibility in a moment), then it is especially important for science to identify a set Γ of claims whose inductive confirmation does not mislead us, but rather leads (non-trivially) to the confirmation of truths alone. For any such set Γ , $Cl(\Gamma)$ is Λ . So in identifying the laws, science identifies what it would be best to confirm inductively. This would explain why it is so important for science to discover what the laws of nature are.¹⁶

4. Inductive Strategies as Free Electives

Why might inductive confirmation be especially important in science? Let me make a very rough proposal. Return to the pear example above (example 3). The hypothesis h , “All of the pears now on the tree are ripe”, is projected onto certain p 's; by finding a pear on the tree to be ripe, we confirm h 's predictions regarding p_i = “There is a pear on the third branch of the tree”. But

comes along), then each of the predictions that we are thereby willing to confirm is true, no matter which (if any) of the predictions $p > q$ where p is consistent with $Cl(\Gamma)$ we already assign the maximal degree of confidence, if and only if $Cl(\Gamma)$ possesses non-nomic stability.

¹⁶ I shall not consider here what would make the h_i belonging to one set Γ , where $Cl(\Gamma) = \Lambda$, better to confirm inductively than the h_i in another such set. (This distinction ultimately corresponds, I believe, to the distinction between the natural laws and the physical necessities that are not natural laws, such as “All signals travel slower than twice the speed of light”, “All things that are emeralds or rubies are green if emeralds, red if rubies,” and “All non-green things are non-emeralds”. See Fodor (1981, p. 40) and my (forthcoming-b).) I shall ignore this question here, and presume that each set Γ , where $Cl(\Gamma) = \Lambda$, contains exactly the h_i that it is best for us to confirm inductively. (Accordingly, perhaps “optimal” would be a better word than “best”.) Notice that for each of these sets, the logical closure in U^* of the predictions that are confirmed, in confirming every member of that set inductively, is the same. So it makes no difference, as far as these predictions are concerned, which of these sets is the set of hypotheses that we are prepared to confirm inductively.

we are not willing to project h onto certain other p 's, such as $p_2 =$ "The tree holds a pear that experienced environmental conditions that were experimentally manipulated to be ... [differing greatly from the conditions actually experienced by the pears on the tree]". What justifies our projecting h so far but no farther? When h alone is considered, the limits of its projection appear arbitrary; h itself provides no basis for treating p_1 differently from p_2 . Indeed, even when we consider h along with various other relevant claims that we have projected (e.g., $j =$ "All pears on the same tree that have experienced roughly the same environmental conditions are ripe to about the same degree", $k =$ "All of the pears now on the tree have experienced roughly the same environmental conditions"), we still find no ground for projecting h onto p_1 but not onto p_2 . In particular, it is logically possible for $h, j,$ and k to have all obtained together with p_1 , and likewise $h, j,$ and k could all have held together with p_2 . But of course, the limits of h 's projection are not arbitrary. They are determined by our *reason* for taking the ripeness of one actual pear on the tree as relevant confirmation-wise to the ripeness of another actual pear on the tree. For example, those counterfactual circumstances p regarding which we have already highly confirmed that both j and k hold are among those onto which we are prepared to project h . These include p_1 but not p_2 (onto which we have not projected k). In short, then, our prior empirical work justifies our taking the discovery that a given pear actually on the tree is ripe as confirming, of each actual or counterfactual pear in a certain range, that it is ripe; without this empirical work, that range would be arbitrary, unmotivated.

In contrast, when a hypothesis h is confirmed *inductively*, the limits of its projection are fixed by h itself along with the other hypotheses that we are committed to confirming inductively (should we discover one of their predictions to be borne out). No appeal need be made to other background opinions that we have already arrived at through prior empirical work. From the viewpoint of the various hypotheses h_i (including h) that we are committed to confirming inductively, the range of h 's projection is not arbitrary; we are prepared to project *each* of the h_i as far across the p 's in U as they can *all* logically possibly go. Inductive projection is *indiscriminate*; it draws no distinction among the p 's across which we *could* project all of the h_i . Therefore, it is not beholden to prior empirical work to supply a good reason for drawing such a distinction—for projecting the hypothesis so far but no farther.

Here is why this is important. We pursue various "strategies" for arriving justly at predictions. Part of carrying out such strategies is categorizing objects or situations in certain ways for the purpose of seeking regularities of certain sorts—e.g., choosing to regard copper, Cepheid-type variable stars, and autism as perhaps constituting natural kinds of certain sorts. Once we have tentatively adopted such a categorization and decided provisionally to seek a regularity of some specified kind covering some category, we are pre-

pared to take (say) the electrical conductivity of one copper object as relevant confirmation-wise to another copper object's electrical conductivity, or the response of one autistic patient to a given drug as relevant confirmation-wise to another autistic patient's response. As relevant to *which* others? As the four examples given earlier (the die, the pear tree, and so on) illustrated, the range of projection in different examples can vary considerably. The limits of *inductive* projection are not fixed by considerations beyond the various hypotheses that we are prepared so to project; nothing more is needed to justify including certain of *h*'s predictions but not certain others in the range to be confirmed by *h*'s making some accurate prediction. Therefore, inductive confirmation is the only sort of projection that we can justly make when we have nothing more available to justify such discrimination among *h*'s predictions. The only sort of strategy that we need no prior empirical work to entitle us to carry out—that makes no discrimination requiring justification from prior empirical work—is an *inductive* strategy, so-called because we are prepared to confirm “inductively” the hypotheses generated by the inductive strategies that we are carrying out. That is why inductive confirmation is so important in science.

I shall now explain this proposal.

To understand the scientific importance of inductive projection, we must back up a bit to consider a fundamental question: What entitles us to regard one discovery as confirming the truth of various other claims? The answer, of course (albeit rather unilluminating at this high level of generality), is that the other opinions we hold when we take our discovery into account determine what we should take that discovery as confirming. For instance, in the pear case, we start with considerable confidence that all of the pears actually on the tree have experienced roughly the same environmental conditions, and that these conditions (whatever they may be) are responsible for the pears' degree of ripeness. It is this opinion that gives our discovery of a ripe pear on the tree the power to confirm, of every other pear actually on the tree, that it is ripe. Our earlier observations justify our holding this background opinion.

But how can we use one discovery to confirm the truth of various other claims when we lack any considerable body of relevant prior opinion—when we have little in the way of relevant past observations and already well-supported theories? In such circumstances, how can we be justified in taking one claim's truth as confirming another's—in holding the prior opinions (i.e., conditional probabilities) required for such confirmation?¹⁷ Let's look at an actual scientific example.

¹⁷ It might be suggested that such confirmation is sometimes virtually automatic according to Bayesian conditionalization, since if *h* entails *e* (as when *h* is “All of the pears on the tree are ripe” and *e* is “The first pear selected from the tree is ripe”), then $\text{pr}(el|h) = 1$, and so when $0 < \text{pr}(h) < 1$ and $0 < \text{pr}(e) < 1$, then $\text{pr}(h|e) > \text{pr}(h)$. But even presupposing Bayesian confirmation theory, the above argument avails us not. Our question concerned

Psychiatry today is much like chemistry was in the late eighteenth century, when Kant bemoaned its lack of any overarching theory. To be sure, psychiatry has justified various generalizations, allowing predictions to be made (with some confidence) from various observations. But it largely lacks any comprehensive theory systematizing or explaining these generalizations. These generalizations, then, must have been arrived at largely on the basis of observations—various case histories—not supplemented by any sophisticated theoretical considerations. How, then, have psychiatrists managed to confirm the predictive accuracy of these generalizations, when they have had no basis for the prior opinions apparently needed to justify regarding their observations as confirming those predictions?

To appreciate the difficulty here, consider some list A, B, C... of psychiatric symptoms, such as “unreasonable insistence on following routines in precise detail”, “delusions”, and “absence of imaginative activity”. Suppose that some patient exhibiting symptoms A, B, and C, but none of the other symptoms on the list, is discovered to respond in a certain way to the administration of some new drug. Should psychiatrists regard this discovery as confirming (by some increment), of any other patient exhibiting A, B, and C but none of the other symptoms on this list, that she would have the same response to the drug? What about a patient exhibiting A, B, C, and only one of the other symptoms on the list? What about a patient exhibiting at least two of A, B, and C, and no more than one of the other symptoms on the list? In the absence of considerable background information, psychiatrists apparently can have no good reason for pursuing one of these ampliative policies rather than another; the range of confirmed predictions appears to be fixed arbitrarily. Yet apparently, psychiatrists are not on these grounds precluded from justly regarding their evidence as confirming some range of predictions rather than another. Psychiatrists appear to be justified in pursuing a hunch: simply deciding to take a certain class of evidence as relevant confirmation-wise to a certain class of predictions, even though their opinions prior to making this hunch fail to supply them with a good reason for imposing this particular limitation on the range of their projections.¹⁸ What sort of hunch is a scientist permitted to pursue, when her prior opinions are not determining

how we justify taking e as confirming some of h 's predictions (e.g., “Were we to select a second pear from the tree, then it would be ripe”). No such prediction $p > q$ entails e , and so it remains unclear on what ground we can justify $\text{pr}(elp > q) > \text{pr}(e)$, which is necessary (according to Bayesian conditionalization) for e to confirm $p > q$.

¹⁸ By saying that psychiatrists are “permitted to” (i.e., “entitled” to, “justified” in deciding to) pursue this hunch (a certain inductive strategy), I obviously do not mean that they have good reason to believe that the gamble will pay off, since their predicament is precisely that they lack any such reason (as I shall explain further). Rather, I mean that in pursuing the hunch, the psychiatrists are not acting in an epistemically irresponsible fashion, but are faithfully carrying out their epistemic duties, acting within their rights, subject to no epistemic reproach.

which predictions she would then be prepared to regard certain evidence as confirming?¹⁹

This is obviously part of the classical “problem of induction.” What I am about to suggest is not meant to address this entire problem or to exclude other approaches. But certain other approaches do not work. For example, it might be supposed that psychiatrists justify grouping together the patients exhibiting certain combinations of symptoms on the basis of having already found some similarities among these cases—perhaps in their reactions to some drugs or in other features suggestive of a common mechanism responsible for these symptoms. But to have already discovered that all patients with a certain combination of symptoms also have in common the disposition to respond in a certain way to the administration of a given drug, psychiatrists presumably must have already checked some patients in this category and then projected the result across the category. Such projection requires that psychiatrists already regard the examined cases as relevant confirmation-wise to the unexamined cases. It presupposes, in other words, the categorization that it was supposed to justify. It might be insisted that psychiatrists tentatively form certain categories of patients, and then support this categorization by appealing to the results reached by projecting across these categories. I agree, but for these results to support anything, they must themselves be warranted, and that depends upon the initial categorization’s being warranted. I shall argue that one way science gets off the ground is with some free moves. But not every sort of policy for regarding a certain class of evidence as bearing confirmation-wise on a certain class of predictions is the kind of policy that scientists can be entitled to pursue for free. For example, to group together all *actual* patients exhibiting symptoms A, B, and C but none of the *counterfactual* patients of this sort involves drawing an arbitrary distinction; if our *reason* for regarding one actual patient’s response to a given drug as bearing upon another’s is that both fall into this diagnostic category, then we must regard the actual patient’s response as bearing on some possible patients in this category regardless of whether they are ever actualized. I am suggesting that only strategies involving *inductive* projection are indiscriminate in a manner that allows them to be free moves.

¹⁹ In terms of Bayesian confirmation theory, I am suggesting that scientists are sometimes permitted simply to decide to adopt some new probability distribution, rather than arriving at it by updating their prior distribution by Bayesian conditionalization upon the receipt of new evidence. Their new probability distribution then supplies the conditional probabilities expressing their willingness to regard some class of evidence as relevant confirmation-wise to some class of predictions. But I am not suggesting that scientists are permitted to decide to adopt just any new probability distribution at all. So which non-Bayesian shifts of opinion are permissible? Which would be arbitrary or biased in a sense that would render a shift to them unjustified? That is the question I am asking. Obviously, non-Bayesian shifts of opinion raise a great many issues (e.g., regarding Dutch book arguments that Bayesian conditionalization is the only rational way to change opinions) that are beyond the scope of this paper.

A psychiatrist might begin an “inductive strategy” by grouping together any collection of symptoms she wishes as a single diagnostic category (e.g., “autism”). Then she decides to seek a truth in U that would supply certain information regarding the cases in this category—e.g., a truth specifying how any autistic patient would respond to administration of a certain drug. Her next step is to observe various autistic patients after they have received this drug. It may happen that this evidence suggests a unique hypothesis. By this, I do not mean something psychological—e.g., that the scientist’s observations of these autistic patients lead her to come up with some hypothesis. Rather, I mean that scientists observe what it would be for unexamined autistic patients to respond in the same way to the drug as the examined autistic patients have done. In contemplating some way in which unexamined cases might depart from the hypothesis suggested by the evidence, scientists observe that this departure would involve the unexamined cases’ behaving differently from the cases heretofore examined. That this is an observation means that the scientists’ belief possesses a certain special kind of justificatory status; in the manner characteristic of observation reports, this status requires that there be widespread agreement, among qualified observers who are shown the data, regarding which hypothesis expresses what it would be for unexamined cases to go on in the same way as the cases already examined.²⁰ The scientist who is gambling on this risky strategy then regards any successful prediction made by the salient hypothesis as confirming that hypothesis inductively—as I shall explain momentarily.

The history of science is replete with the pursuit of inductive strategies. Balmer and Rydberg carried out inductive strategies in confirming various regularities in the locations of lines in the spectrum of hydrogen. Leavitt pursued an inductive strategy in arriving at the period-luminosity relation governing Cepheid-type variable stars. (She could thereby use observations of only a few, nearby Cepheids to justify predictions regarding the period-lumi-

²⁰ What we believe this would be depends on our background beliefs. Inductive strategies cannot be employed without *some* background beliefs; the observation that such-and-such behavior by unexamined cases would represent a departure from the way that examined cases have been found to behave depends (for its status as an *observation*) on the putative observer’s justly holding certain prior opinions, just as any other observation does. But a psychiatrist’s justified opinions, prior to engaging in an inductive strategy, are insufficient to justify her in taking one autistic patient’s response to a given drug as relevant confirmation-wise to any other’s regardless of which autistic symptoms (if any) the two patients share. I discuss salience at greater length in my (1998) and (forthcoming-b). In connection with this belief’s status as an observation, I find it useful to bear in mind remarks like those of Dulong and Petit, as they present a table of experimental values to justify their proposed law (that an element’s atomic weight, multiplied by the quantity of heat needed to raise a given mass of that element by a given temperature, is the same constant for all elements): “Mere inspection of these numbers reveals a relation so remarkable in its simplicity that in it one immediately recognizes the existence of a physical law capable of being generalized and extended to all the elements” (quoted by Nash 1956, p. 100).

nosity relation of every other Cepheid, even those in the Andromeda “nebula”, which allowed the nebula’s extra-galactic distance to be determined for the first time.) Einstein employed an inductive strategy (from about 1905) in regarding the success of any prediction made by the light-quantum hypothesis (e.g., equations governing the black-body spectrum, the photoelectric effect, the Volta effect, specific heats, etc.) as confirming each of the others. The current edition of the American Psychiatric Association’s *Diagnostic and Statistical Manual* codifies the diagnostic categories across which U.S. psychiatrists have collectively decided to gamble in projecting their hypotheses. For instance, it classifies as “autistic” any patient possessing at least eight of sixteen designated symptoms, where at least two of the eight belong to the first category of five symptoms, one to the second category of six symptoms, and one to the third category of five symptoms. One autistic patient’s response to a given drug is then supposed to be taken as bearing on what another autistic patient’s response would be, even if the two patients share no particular symptoms at all.

When a scientist pursues some inductive strategy, the evidence suggesting h thereby justifies her in treating one of h ’s predictions as able to confirm another. The scientist’s prior opinions play no part in this. Therefore, the scientist has no basis for taking one of h ’s predictions rather than another as confirming h , or (with one restriction I shall explain momentarily) for regarding one of h ’s predictions rather than another as thereby confirmed. Any such discrimination among h ’s predictions would beg the question: “If *this* prediction is confirmed by the evidence, then why isn’t *that* one?” Without appealing to relevant prior opinions, the scientist has no grounds for such partiality. (In the pear example, scientists’ prior opinions give them a good reason for regarding the discovery of a ripe pear on the tree as confirming some but not all of the predictions made by “All of the pears now on the tree are ripe”.) The scientist pursuing an inductive strategy is not entitled to discriminate even against h ’s predictions concerning various unrealized circumstances. Her reason for regarding her discovery as confirming some claim concerning an *actual* circumstance is (roughly) that this discovery and that claim are both predicted by h . So, on pain of inconsistency in applying this reason, she must recognize her discovery as (roughly) confirming each of h ’s predictions, even a counterfactual conditional.

A scientist needs no justification for launching an inductive strategy. It is a free move—she is entitled to pursue it even if she lacks any relevant prior opinions—*precisely because* in making a discovery bear confirmation-wise upon some other claim’s truth, an inductive strategy does not depend upon any prior opinions. (If the scientist’s prior opinions supplied the justification for pursuing an inductive strategy, then that justification would depend upon the warrant for those prior opinions, which would depend upon certain empirical work having already been done. The strategy would then fail to be a free

elective.) So the kind of confirmation figuring in an inductive strategy must draw no distinction among h 's predictions that could be motivated only by such prior opinions. The inductive strategy's status as a free move thereby determines the particular range of projection that characterizes inductive confirmation. Specifically: in pursuing some set of inductive strategies, we are gambling that they will succeed in taking us to accurate predictions. We are hoping, in particular, that the hypotheses h_i currently salient on these strategies (one per strategy, together forming the set Γ) will receive enough inductive confirmation to warrant their acceptance and the acceptance of all of their predictions confirmed in confirming them inductively, and that all of those predictions are true. So it must be logically possible for all of those predictions to be true, else we are hoping for what we know to be impossible. If p is inconsistent with $Cl(\Gamma)$, then not all of the predictions regarding p that are made by the h_i can be true. So (in order for our hope to be realizable) it cannot be that each prediction that a salient hypothesis makes regarding p is confirmed when that hypothesis is projected inductively. Rather, the inductive strategies must select which of the predictions regarding p are to be confirmed during inductive projection. But the inductive strategies lack precisely the basis they would need to justify any such selection; from the viewpoint of the h_i , any such selection would be arbitrary. So in confirming h inductively, none of its predictions regarding p is confirmed.²¹ In this way, the distinctive range of h 's inductive projection (specified in the previous section as h 's predictions $p > q$ where p is consistent with $Cl(\Gamma)$ and where $p > q$ has not already been ascertained) is explained by the inductive strategy's status as a free move.

Inductive strategies are especially important in science because scientists are entitled to pursue them even before otherwise becoming entitled to hold any relevant background opinions; inductive strategies can be pursued right from the outset of empirical work in some area. The range across which the salient hypotheses are to be inductively projected is fixed not by our prior opinions, but by those hypotheses themselves. The inductive strategies are collectively *self-limiting*, each change in Γ —as when a different hypothesis becomes salient on a given strategy—affecting for each of the other h_i 's the range of its predictions that we regard as having been confirmed in the course of its inductive confirmation.

Here, then, is the moral of my story: (i) We can presume, without begging the title question, that science is interested in predicting the non-nomic facts *in a broad sense*—the facts in U^* ; (ii) To achieve this goal, it is

²¹ Someone carrying out these inductive strategies could nevertheless confirm some of these predictions, but not in virtue of carrying out these inductive strategies. That is, someone may regard a discovery as confirming a certain range of predictions in virtue of carrying out some inductive strategy, and in addition, as confirming various other predictions for other reasons.

especially valuable for science to find a set of inductive strategies for which it is non-trivially the case that as these strategies are carried out, the predictions that become confirmed sufficiently to justify their acceptance are all true; and (iii) This occurs if and only if the logical closure in U of the various hypotheses that are ultimately generated by these strategies is Λ .²² In short, to discover which truths in U state laws is to discover which inductive strategies form the best set for us to carry out. Therefore, it is especially important for science to identify the laws in U .²³

5. Conclusion

Let's take stock. I began section 1 by suggesting that science is *interested* in discovering whether $\Box p$ or $\neg\Box p$ (for $p \in U$), insofar as this goes beyond discovering whether or not p , because science thereby learns something about what would have happened under various unrealized non-nomic circumstances. Since the subjunctive conditionals that science thereby ascertains to be correct have antecedents and consequents in U , we can grant science's interest in ascertaining their truth-values without begging the question (i.e., without

²² This does *not* mean that scientists can justly pursue some set of inductive strategies only if they believe that they already know all of the laws, or only if they believe that their set of inductive strategies is *complete*—i.e., that the hypotheses ultimately rendered salient on these strategies span Λ . Scientists can inductively project the hypotheses in Γ , believing that these *may* span Λ , but without having tremendous confidence that they do. This is precisely the case in psychiatry today; psychiatrists do not believe that their current diagnostic categories are correct, or even probably correct, but they justly carry out the inductive strategies corresponding to this classificatory scheme. When scientists learn more about the laws—as when new inductive strategies are initiated or old ones abandoned, or different hypotheses become salient on strategies already being pursued—then scientists amend the range of predictions $p > q$ that they take to be confirmed, in projecting inductively some hypothesis that they were already projecting inductively. (Recall the closing sentence of the previous paragraph.) Surely, scientists often do change their minds concerning the relevance of a given observation to a given prediction (as when psychiatrists amend the boundaries of some of their diagnostic categories in revising their *Diagnostic and Statistical Manual*, as they frequently do). Moreover, I can justly believe that Λ is consistent with p —e.g., that my now holding a copper object in my hand is consistent with the natural laws—even if I do not believe that I have identified all of the natural laws. I may therefore confirm some prediction regarding p , in confirming various hypotheses inductively, even before I believe that I have found all of the laws.

²³ Inductive confirmation is available only to hypotheses that logically follow from currently salient hypotheses. (Roughly speaking, then, inductive confirmation is confined to inductive strategies.) For if h does not so follow, then $\neg h$ is consistent with the closure of the currently salient hypotheses, and so to confirm h inductively, we must confirm each of its predictions (that we do not yet believe) regarding $\neg h$. But by the definition of “prediction”, h makes some logically impossible predictions regarding $\neg h$ (e.g., $(\neg h > h)$, $(\neg h > p \ \& \ \neg p)$), and these cannot be confirmed. Note, then, that if we must believe that the currently salient hypotheses may be physically necessary—may be what is arrived at by carrying out the best set of inductive strategies—then inductive confirmation is available only to hypotheses that we believe may be physically necessary. The intuitions of Goodman *et al.* are thus vindicated. (Recall note 4.)

thereby merely stipulating science to be interested in ascertaining whether $\Box p$ or $\neg\Box p$).

I have now gone one step further. I have just argued that a reason why discovering whether $\Box p$ or $\neg\Box p$ (for some truth $p \in U$) is so *important* to science is because this is discovering how best to go about discovering the truth-values of various claims in U^* . In aiming to learn which set of inductive strategies is best for us to carry out, science is aiming to identify the members of Λ .

If we select the right inductive strategies to carry out, then we can ascertain some of the laws' accurate predictions in U^* without depending upon prior empirical work to have already given us a good reason for picking out those predictions, rather than others, as the ones to be confirmed.²⁴ In this respect, the laws differ in their range of predictive accuracy from the non-nomic facts that lack physical necessity, such as that all of the pears on the tree are ripe. For us to confirm exclusively the *accurate* predictions that some accidental generalization makes, we must have already conducted enough empirical work to have warranted the background opinions needed to privilege that particular range of projection. In the pear example, for instance, we must have already become confident that a pear's degree of ripeness is determined partly by the amount of sunlight received in the preceding weeks. Otherwise, in confirming the accidental generalization, we might confirm that all of the pears on the tree would still have been ripe even if they had received much less sunlight in the preceding weeks. It may, then, be far less difficult to justify projecting a hypothesis that turns out to be a natural law across a range in which it turns out to make only accurate predictions than to justify projecting a hypothesis that turns out to be an accidental truth across a range in which it turns out to make only accurate predictions, since the latter projection demands that we motivate drawing some otherwise arbitrary distinction among the various predictions made by the hypothesis.

Here I am taking issue with a remark of Ayer's:

If on the basis of the fact that all the A's hitherto observed have been B's we are seeking for an assurance that the next A we come upon will be a B, the knowledge, if we could have it, that all A's are

²⁴ Of course, we may not pursue the right inductive strategies; we cannot know in advance which inductive strategies are best. But neither can we know the laws of nature *a priori*. Furthermore, our inductive strategies need not have rendered salient *all* of the laws in order for the predictions they yield that are relevant to our project all to be accurate; it is sufficient that the salient hypotheses entail all of the *relevant* laws. For instance, we do not need to have rendered salient all of the laws of elementary particle physics (or psychology) in order for our inductive strategies to render salient the period-luminosity relation governing Cepheid variables, and so to result in the confirmation of "Were a Cepheid variable star to exist in the Andromeda nebula, then its period-luminosity ratio would be ..." but not in the confirmation of "Were a Cepheid variable to be accelerated from rest to beyond the speed of light, then its period-luminosity ratio would be...."

B's would be quite sufficient; to strengthen the premise by saying that they not only are but must be B's adds nothing to the validity of the inference. The only way in which this move could be helpful would be if it were somehow easier to discover that all A's must be B's than that they merely were so.

But this is not possible:

It must be easier to discover, or at least find some good reason for believing, that such and such an association of properties always does obtain, than that it must obtain; for it requires less for the evidence to establish. (1976, pp. 149f.)

How can you acquire a good reason for regarding some observed A's B-ness as confirming that the next A to be examined will be B? If you are not pursuing an inductive strategy, then you need to have already done some empirical work in order to justify regarding your observation as relevant confirmation-wise to the next A's B-ness. (In the pear case, you typically need to have amassed sufficient evidence to justify believing that the examined and unexamined pears developed under environmental conditions that are similar in those respects that influence the pears' ripeness.) It may be far easier—require much less prior empirical work—for the observed A's B-ness to confirm the next A's B-ness in the course of some inductive strategy on which “All A's are B” becomes salient. The decision to carry out that inductive strategy is a free move, and it may then take relatively little evidence to render “All A's are B” salient. It is hard to imagine how psychiatrists today could justly regard a given autistic patient's response to a drug as relevant confirmation-wise to how any other actual autistic patient would respond to that drug except through carrying out an inductive strategy. The same could be said for any of the other historical examples of inductive strategies that I gave earlier. For instance, how could Leavitt in 1908 have justified taking the periods and luminosities of nearby Cepheid-type variable stars as bearing upon the period-luminosity ratio of Cepheids in the Andromeda “nebula” except as a risky inductive gamble that she was entitled to elect in virtue of the indiscriminate character of the projection involved? Utterly lacking any theories concerning the internal structure of stars, much less of Cepheid variables, she was in no position to offer any other sort of justification for the requisite conditional probabilities. The same thought is echoed in Planck's 1913 recommendation that Einstein be admitted to the Prussian Academy. Concerning Einstein's inductive strategy of regarding any successful prediction made by the light-quantum hypothesis (e.g., equations governing the black-body spectrum and the photoelectric effect) as confirming any other prediction made by the light-quantum hypothesis (e.g., equations concerning the

Volta effect and specific heats), Planck incorrectly judged that Einstein had failed. But he understood that a scientist is permitted to elect to take a gamble of this kind:

That [Einstein] may sometimes have missed the target in his speculations, as, for example, in his hypothesis of light-quanta, cannot really be held too much against him, for it is not possible to introduce really new ideas even in the most exact sciences without sometimes taking a risk. (Kirsten and Korber 1975, p. 201)

Now if “All A’s are B” becomes salient on one of your inductive strategies, and you confirm that claim inductively to a very high degree, and you ultimately adopt it, then you believe that it is physically necessary that all A’s are B, since you believe that “All A’s are B” belongs to a set possessing non-nomic stability. Contrary to Ayer, then, it can be easier to justify believing that it is physically necessary that all A’s are B than to justify believing that all A’s are B without believing that it is physically necessary that all A’s are B.

Of course, Ayer is correct to point out that since “All A’s are B”’s lawhood entails its truth, our degree of confidence in its lawhood cannot exceed our degree of confidence in its truth. But this does not show that it is harder to justify believing that it is a law than to justify believing that it is true *without thereby justifying the belief that it is a law*.²⁵ Ayer’s argument pre-

²⁵ Foster (1983, p. 88) appears to be making roughly the same objection to Ayer. But he contends that “if extrapolative induction [i.e., an extension to all or some of the unexamined cases—for example, to all “nomologically possible” cases—of what we have found to hold for the examined cases] is the only form of inference [to $\blacksquare h$], then Ayer is clearly right.” I have argued that this is mistaken; inductive projection (a kind of “extrapolative induction”) can justify acceptance of $\blacksquare h$ but not of h without $\blacksquare h$. Foster instead defends inference to the best explanation as playing this role. As I mentioned in section 1, I have no problem with using science’s interest in explanations rather than its interest in counterfactuals in order to account for its interest in the laws. But this strategy requires an argument that laws bear a special relation to explanations, just as I have had to argue that laws bear a special relation to counterfactuals. Foster argues that accidental generalizations are not explanatory because “[i]n subsuming the past regularity under a universal regularity [that we do not believe to be physically necessary] we would not be diminishing its coincidental character, but merely extending the scope of the coincidence to cover a larger domain” (p. 91). But some accidental generalizations are not utter “coincidences.” In my pear example, we explain why all of the examined pears from the given tree are ripe by subsuming this fact under the regularity that all of the pears on the tree are ripe. This regularity is (modestly) explanatory despite its lack of physical necessity; it shows that it was no coincidence that all of the pears we picked for examination were ripe. That is, this accidental regularity’s explanatory power derives partly from its preservation under a certain relevant range of counterfactual suppositions; for example, it would still have held had we checked different pears from the tree. On this view, a law’s preservation under counterfactual suppositions partly accounts for its explanatory power. (For more discussion of explanatory non-laws and their invariance under counterfactuals, with particular attention to explanatory generalizations in the

supposes that some evidence would, as it were, *first* bear upon whether various unexamined A's *are* B, and *then* you would need some *further reason* for regarding that evidence as bearing upon whether those unexamined A's *must be* B. I suggest, on the contrary, that you may be required to believe that various unexamined A's *must be* B in virtue of your reason for believing that various unexamined A's *are* B.²⁶

A basic presupposition of scientific research is that we do not need to *observe* whether a claim is true in order to ascertain this fact. (Indeed, to ascertain the correctness of some non-trivial counterfactual conditional, we do not have the option—even in principle—of observing whether it is true.) Science is very much interested in knowing how it can use its observations to make accurate predictions when it begins without any reason to regard its observations as bearing confirmation-wise upon any predictions. In seeking the best way to proceed from such ignorance to knowledge of non-nomic facts (in the broad sense) beyond the limited range of past observations, science seeks to identify the laws. Beliefs about the laws, over and above beliefs about facts in U*, must therefore be acknowledged as playing an important role in scientific investigation of the non-nomic facts.²⁷

social sciences, see Woodward (forthcoming) and my (forthcoming-b).) Foster's approach, I think, fails to explain why laws as distinct from explanatory non-laws are so important to science.

Foster favors inference to the best explanation over any sort of extrapolation because he believes that "[w]hen rational, an extrapolative inference can be justified by being recast as the product of two further steps of inference, neither of which is, as such, extrapolative. The first step is an inference to the best explanation—an explanation of the past regularity whose extrapolation is at issue. The second is a deduction from this explanation that the regularity will continue or that it will do so subject to the continued obtaining of certain conditions" (p. 90). I am unsure of this. Why couldn't we extrapolate some past regularity (to some unexamined cases or to all "nomologically possible" cases) without adopting any explanation of that regularity, so long as the past regularity (and our background beliefs) give us good reason to believe that the explanatory factors (whatever they are) will continue to hold? For instance, scientists expect the roughly 22-year sunspot cycle to persist even though they do not understand why it holds, only that it has held steady for a long time.

²⁶ Likewise, your reason for believing that various *actual* unexamined A's are B may apply just as well to belief that various *counterfactual* (and perforce unexamined) A's would have been B. Counterfactual conditionals are confirmed empirically in the course of confirming non-counterfactual claims; there is no need for some *special, new* kind of reason for *going beyond* our beliefs about the actual world and undertaking various beliefs about counterfactual cases.

²⁷ I would like to thank Harold Hodes, Lawrence Bonjour, and anonymous referees for this journal for their helpful suggestions.

References

- A. J. Ayer (1976), *The Central Questions in Philosophy* (London: Pelican).
- Jonathan Bennett (1984), "Counterfactuals and Temporal Direction," *Philosophical Review* 93: 57–91.
- R. B. Braithwaite (1927), "The Idea of Necessary Connexion (I)," *Mind* 36: 467–77.
- Rudolf Carnap (1950/1962), *The Logical Foundations of Probability* (Chicago: University of Chicago Press).
- John Carroll (1994), *Laws of Nature* (Cambridge: Cambridge University Press).
- Roderick Chisholm (1946), "The Contrary-to-fact Conditional," *Mind* 55: 289–307.
- _____ (1955), "Law Statements and Counterfactual Inference," *Analysis* 15: 97–105.
- Fred I. Dretske (1977), "Laws of Nature," *Philosophy of Science* 44: 248–68.
- John Foster (1983), "Induction, Explanation, and Natural Necessity," *Proceedings of the Aristotelian Society* 83: 87–102.
- Nelson Goodman (1947), "The Problem of Counterfactual Conditionals," *Journal of Philosophy* 44: 113–28.
- _____ (1983), *Fact, Fiction, and Forecast*, 4th edition (Cambridge: Harvard).
- Paul Horwich (1987), *Asymmetries in Time* (Cambridge: MIT Press).
- Frank Jackson (1977), "A Causal Theory of Counterfactuals," *Australasian Journal of Philosophy* 55: 3–21.
- G. Kirsten and H.-G. Korber ((1975), *Physiker uber Physiker* (Berlin: Akademie-Verlag).
- William Kneale (1952), *Probability and Induction* (Oxford: University Press).
- Marc Lange (1998), "Salience, Supervenience, and Sellars", *Philosophical Studies*.
- _____ (forthcoming-a), "Laws, Counterfactuals, Stability, and Degrees of Lawhood", *Philosophy of Science*.
- _____ (forthcoming-b), *Natural Laws in Scientific Practice* (New York: Oxford University Press).
- David Lewis (1983), "New Work for a Theory of Universals," *Australasian Journal of Philosophy* 61: 343–77.
- J. L. Mackie (1962), "Counterfactuals and Causal Laws," in *Analytic Philosophy*, ed. R. S. Butler (New York: Barnes and Noble), 66–80.
- John Stuart Mill (1893), *A System of Logic*, 8th ed. (New York: Harper and Bros.).
- G. E. Moore (1962), *Common-Place Book 1919–1953*, ed. C. Lewy (London: George Allen and Unwin).

- Leonard K. Nash (1956), *The Atomic-Molecular Theory* (Cambridge, Harvard University Press).
- C. S. Peirce (1934), *Collected Papers*, Volume 5, ed. C. Hartshorne and P. Weiss (Cambridge: Harvard University Press).
- John Pollock (1976), *Subjunctive Reasoning* (Dordrecht: Reidel).
- Hans Reichenbach (1947), *Elements of Symbolic Logic* (New York: Macmillan).
- Israel Scheffler (1981), *The Anatomy of Inquiry* (Indianapolis: Hackett).
- Elliott Sober (1988), "Confirmation and Law-Likeness," *Philosophical Review* 97: 93–98.
- P. F. Strawson (1952), *Introduction to Logical Theory* (London: Methuen).
- Bas van Fraassen (1989), *Laws and Symmetry* (Oxford: Clarendon Press).
- James Woodward (forthcoming), "Causation and Explanation in Linear Models", in *The Reliability of Economic Models*, ed. Daniel Little (Dordrecht: Kluwer).