Why Confirm Laws?

Barry Ward

University of Arkansas

Abstract

We argue that a particular approach to satisfying the broad predictive ambitions of the sciences demands law confirmation. On this approach we confirm non-nomic generalizations by confirming there are no actually realized ways of *causing* disconfirming cases. This gives causal generalizations a crucial role in prediction. We then show how rational judgements of relevant causal similarity can be used to confirm that causal generalizations themselves have no actual disconfirmers, providing a distinctive and clearly viable methodology for inductively confirming them. Finally, we argue that for agents in our epistemic position using this methodology to confirm causal generalizations of adequate breadth will commonly demand law confirmation. We make a *prima facie* assessment of the methodology's fit with scientific practice and briefly consider the prospects for an associated analysis of laws.

Keywords: Laws, Induction, Confirmation, Prediction, Causal Inference, Invariance

1. Introduction

Many philosophers have given law confirmation a central role in rational induction. While they acknowledge that we can confirm generalizations we hold to be accidental, in the sense that we can raise our probability that they are true, we cannot inductively confirm them i.e., raise our probability that unexamined cases conform. We're merely engaged in content cutting.

Nelson Goodman (1983, 73) restricts inductive confirmation to lawlike generalizations. David Armstrong (1983, 52-59) and Fred Dretske (1977, 256-258) each offer distinct arguments that rational induction must involve confirmation of specifically non-Humean laws. Plausibly, however, they are mistaken. Elliott Sober, Bas van Fraassen and others, have provided examples that clearly demonstrate that we can reasonably inductively confirm accidental generalizations without any appeal to laws.[1]

Should we conclude that laws have no distinguished role in prediction? Perhaps they're of importance for science because they support explanations but play no special role in the formation of expectations. That would be to embrace a false dichotomy: either all rational inductive confirmation must involve laws or laws bring nothing distinctive to the table.[2]

We shall argue for a different view. Laws are at least entailed by—and perhaps identical with—causal generalizations of a certain kind, and so, they may be inductively confirmed by methodologies that specifically apply to causal claims. However, unlike more parochial causal generalizations, laws are fit to serve the broadest predictive ambitions of the sciences. Hence, they merit a distinctive place in our understanding of scientific prediction and our confirmation theory.

---

[1] We shall not detail these examples. They are precisely and compactly described in Sober (1988, 95-96) and van Fraassen (1989, 136-137). Salmon (1989, 49) provides an example that's also compelling, but not so clearly decisive as the other two.

[2] Michael Tooley (1987, 135) has argued that confirmation of non-Humean laws is needed to rationally assign non-zero probability to generalizations with a potentially infinite set of potential disconfirmers based on finite evidence. Unlike Armstrong and Dretske he's not arguing that the formation of *any* expectations demands appeal to non-Humean laws, only those expressed by confidence in such generalizations. While we don't agree with Tooley's argument, it illustrates one way the dichotomy might be evaded.

Sections 2, 3, and 4 develop the case that laws can play this distinctive role. Section 5 assesses the confirmational story's fit with scientific practice. Section 6 briefly assesses the prospects for an associated analysis of laws.

2. The Causal Methodology for Confirming Generalizations

Our discussion shall be largely qualitative. However, we'll also make use of a quantitative, Bayesian approach, and we introduce the essential formalism and terminology here. The probability of a hypothesis, h, conditional on evidence, e, and background beliefs, K, is given by Bayes's theorem:

$$P(h|e \& K) = P(e|h \& K) . P(h|K) / P(e|K)$$

On learning e, an agent updates her personal probability in h from her prior probability, P(h|K), to her posterior probability, P(h|e & K). Thus, e *incrementally confirms* h, or simply *confirms* h, relative to K if, and only if, P(h|e & K) > P(h|K). The *absolute* degree of confirmation of h is just h's probability. We shall say that h is *absolutely well confirmed*, or just *well confirmed*, for an agent if she gives h a high probability. What counts as a high probability is inevitably vague, and none the worse for that. We can now proceed.

When we seek to confirm a generalization for broad predictive use, we ideally seek to well confirm it. If it's not well confirmed, we're not highly confident that it's generally reliable as a predictor. Of course, our predictive ambitions might be restricted to some domain or some degree of approximation or both. Nevertheless, even for these more circumscribed predictive goals we ideally seek to well confirm *some* generalization—one that is suitably restricted and/or posits an approximate fit between the predictions made and the world. So, well confirming is a legitimate ideal for generalizations a scientist might rely upon for prediction. How might we approach it?

So as not to prejudice our discussion, let's consider a generalization that we're confident is not a law, Reichenbach's generalization:

All solid spheres of gold have a diameter of less than 1 mile.[3]

It's plausibly true. On what basis do we make that judgement? We haven't gone looking for large gold spheres, and even if they were out there, we couldn't detect them at intergalactic or even interstellar distance scales. Nor do we need to rely on geological surveys of the Earth's crust to hold it plausible. So, our methodology doesn't resemble random sampling or just blindly sampling from the universe at large, not even in some highly idealized sense. To the extent that we have justifiable confidence in the generalization's truth, it's by other means.

We can reason as follows. We inventory the various kinds of processes that might *cause* gold spheres of diameter of a mile or more and evaluate whether they're at all likely to be actual. One possibility is gravitational accretion. However, we know the interstellar medium was and is predominantly hydrogen, and our knowledge of how stars *work*—i.e., more causal knowledge, notably that stars do not form gold by nuclear fusion in the non-catastrophic parts of their lifecycles—allows us to confidently predict that medium will not contain suitably large aggregates of reasonably pure gold in the future either. So, gravitational accretion will not cause colossal gold spheres. Our inventory of possible causal processes also includes ones that synthesize gold. However, the neutron star collisions and supernova explosions in which gold is formed by nuclear fusion are incredibly violent and cause the resultant gold to be rapidly dispersed. So, they will not cause colossal gold spheres. Moreover, we know of no other fundamental physical, geophysical, or chemical processes that would remotely likely cause suitably colossal aggregates

[3] Van Fraassen (1989, 27, note 12) notes that this example was informally discussed at UCLA in the 1960s, and suspects it derives from similar examples due to Reichenbach.

of gold of any shape. Given that we're confident that our knowledge of the kinds of processes that could cause colossal gold spheres is comprehensive, we can rationally well confirm this accidental generalization.

The above methodology, let's call it *the causal methodology*, can be justified as follows. We can well confirm a generalization if we can well confirm it has no actual disconfirmers. We can do that if we can individuate a suitably comprehensive set of ways of potentially causing disconfirmers and well confirm that set does not yield actual disconfirmers. As instanced by the above example, that's something we can do. So, the nomic—broadly construed to include the causal—can play a useful and distinctive role in achieving our predictive goals.

To be clear, we've not argued that no possible alternative methodology could do this job. For our purposes, it suffices that there's a clearly rational methodology for well confirming generalizations that essentially depends upon causal commitments.

Now this example does *not* directly involve induction. We well confirmed Reichenbach's generalization by reasoning from antecedently well confirmed nomic claims and boundary conditions. However, since rationally implementing the methodology ultimately demands those nomic claims are well confirmed from the scant data that we can gather, it will evidently involve the inductive confirmation of nomic claims. So, we've identified a distinctive role for the inductive confirmation of at least one species of nomic claim, causal ones, in our pursuit of predictive power.

We did indeed appeal to several laws, the law of gravitation and—implicitly—those concerning the strong nuclear and electromagnetic interactions that cover nuclear fusion. Both in compiling our inventory of the kinds of causal processes that might yield colossal gold spheres and evaluating the outcomes of the processes for various boundary conditions we would naturally employ these laws. However, while causal beliefs are clearly integral to this methodology—at least for compiling the inventory—it's not evident that laws are.

Laws are fit for purpose: if we know the relevant laws and other salient facts, we can infer what causes what. On the other hand, they go beyond the actual, extending to what causes what in all nomologically possible cases. Since our predictive goal merely demands well confirming generalizations regarding the actual, it's reasonable to ask: why invoke laws?

3. From Confirming Causal Generalizations Regarding Actual Cases to Confirming Laws

To well confirm Reichenbach's generalization we invoked nomic generalizations that involve complicated causal dependencies. For instance, the gravitational force between a pair of masses causally depends on three parameters, the magnitudes of the two masses and the distance between them. In exploring the inductive confirmation of nomic generalizations in service of prediction of the actual, we shall begin with a simpler case. We'll return to more complex ones later.

Suppose we're entertaining the causal generalization "ignited sodium fluoride *causes* a yellow flame / *burns* yellow". Since our concern is prediction of the actual, we obviously shouldn't assume that we need to confirm that in all nomologically possible cases ignited sodium fluoride burns yellow. We only need well confirm "*in all actual cases*, ignited sodium fluoride causes a yellow flame". Moreover, we shouldn't assume we need to confirm that all actual sodium fluoride is relevantly causally similar i.e., that the same feature invariably causes the flame to be yellow. Sodium fluoride may well occur in different forms and even if it should turn out that all indeed burn yellow, different forms might do so for different reasons. Minimally, we need to well confirm "*for whatever reasons*, in all *actual* cases, ignited sodium fluoride causes a yellow flame". How might we do that?

While we shouldn't assume all actual sodium fluoride is relevantly causally similar, if we thought that claim had some *prima facie* plausibility, we might reasonably attempt to well confirm it, and thereby, well confirm our targeted causal generalization. Let's explore this possibility.

Presumably, early in our research we're only acquainted with samples of sodium fluoride that burn yellow. However, we recognize that it may actually occur in forms that might relevantly differ. It might occur in different crystal structures, and such differences might cause a different color flame if ignited. It might also occur in a glass form, say, and that variation might be causally relevant to flame color. And so on. Let's assume we have an inventory of such sub-kinds of sodium fluoride. Further, assume we're confident that cases that belong to each such sub-kind are relevantly causally similar—we're confident these kinds need not be further subdivided given that our concern is the color of the caused flame. High confidence in relevant causal similarity does not mean being confident that all cases of that sub-kind invariably burn *yellow*—they could all burn green—only that they all burn the same color and for the same reason. However, granted such confidence, an inductively rational agent can well confirm that all cases of that sub-kind burn yellow by observing only a modest number of samples that uniformly do so.[4] In particular, granted high confidence that all cases of the kind *actual sodium fluoride glass* are relevantly casually similar, we can rationally well confirm the causal generalization "being an ignited piece of sodium fluoride glass *or something generally associated with that* causes a yellow flame". Moreover, in well confirming the causal generalization, we thereby well confirm that every unexamined case conforms to it.[5] So, what we have is a distinctive methodology for *inductively* confirming specifically causal claim*s*.

Of course, for such confirmation to be rational, we must be able to justify high confidence in the relevant causal similarity of cases of that kind. Can we? One potential concern can be dismissed. In individuating the different sub-kinds, we're taking into account every *known* factor we suspect might be causally relevant to actual flame color. Thus, we account *actual sodium fluoride glass* as such a sub-kind

---

[4] Inductively rational in the minimal sense that, excepting situations where the agent has very unusual background beliefs, merely learning that the samples uniformly burn yellow will not prompt her to lose confidence that all (actual) sodium fluoride glass is relevantly causally similar.

[5] On pain of incoherence, the probability of the generalization must be less than or equal to the probability that any particular case fails to conform.

only if there are no additional known factors that we suspect might undermine relevant causal similarity among those cases. If there were further known factors, then as we already did for (actual) sodium *fluoride*, we would have subdivided *actual sodium fluoride glass* into kinds individuated by them. So, known factors that could prompt suspicions of relevant causal dissimilarity are not a concern.

What about unknown factors? Absence of evidence is not generally evidence of absence, but in cases where we've done *due diligence* in identifying kinds of factors that might be causally relevant—in this case, having done sufficient general chemical research to have identified the kinds of factors that might plausibly actually be relevant to flame color—we're justified in giving little credence to unknown causal dependencies we've no grounds to suspect. High confidence in relevant causal similarity when we've done due diligence is commonplace in inductive reasoning and holding it generally problematic is a form of inductive skepticism that would lay waste to the sciences. Of course, there are numerous generalizations for which our background causal knowledge at a time is too impoverished for us to individuate kinds for which we're confident the subsumed cases are relevantly causally similar. Those are generalizations we cannot well confirm using this methodology, at least at that time. We shall assume the sodium fluoride generalization, and hence, the sodium fluoride glass generalization, is not one of those.

Thus far, we've detailed how we can rationally well confirm "*for whatever reasons*, all *actual* ignited sodium fluoride glass causes a yellow flame" by well confirming "being an ignited piece of sodium fluoride glass *or something generally associated with that* causes a yellow flame". However, given our actual epistemic position, there's a further consideration. Sodium fluoride, whether in glass form or not, is not confined to the Earth, nor indeed our local group of galaxies, nor any time-slice of our universe we might care to individuate. Since we can't survey or even randomly sample the entire actual population of sodium fluoride glass, we've no direct evidence that would allow us to reasonably restrict the set of actually occurring additional factors that might cause (some) sodium fluoride glass to burn non-yellow beyond those that are nomologically possible. Thus, absent some kind of esoteric theoretical reason for holding that a nomologically possible sub-kind of sodium fluoride glass just couldn't be actual, the due diligence required to justify judgements of relevant causal similarity is due diligence regarding all *nomologically possible* factors that might be causally relevant: rational confidence that all *actual* sodium fluoride glass is relevantly causally similar demands a similar degree of confidence that all *nomologically possible* sodium fluoride glass is relevantly causally similar. So, we rationally well confirm the causal generalization for actual sodium fluoride glass only if we well confirm "being an ignited piece of sodium fluoride glass *or something necessarily associated with that* causes a yellow flame". And in well confirming that causal generalization we well confirm "*it is a law* that sodium fluoride glass burns yellow". The same goes *mutatis mutandis* for the other causally individuated sub-kinds of sodium fluoride.

Let's note that we're not identifying "being an ignited piece of sodium fluoride glass or something necessarily associated with that causes a yellow flame" with "it is a law that sodium fluoride glass burns yellow". Whether we should and whether we should embrace a corresponding analysis of lawhood are concerns we'll briefly address later, but they're orthogonal to our primary, confirmational concerns. Whatever the prospects for such an analysis, the causal generalization clearly entails the law. Hence, well confirming it automatically well confirms the law.[6]

Two aspects of our methodology led to law confirmation. If we'd confirmed that pieces of sodium fluoride glass uniformly cause a yellow flame when ignited, but for all we knew they did so for a diversity of reasons, then we shouldn't well confirm the corresponding law. Indeed, if we *knew* they did so for a diversity of reasons, we should reject it. Thus, the use of judgements of relevant causal similarity as a basis for making inductions was crucial. The other aspect was imposed by our impoverished epistemic position. Since we can't reasonably rule out the actual occurrence of any nomologically possible factor that might cause a sodium fluoride glass to burn non-yellow, confirmation that the causal generalization

[6] On pain of incoherence, if ϕ entails ψ, the probability of ψ must be greater than or equal to that of ϕ.

held for all actual cases demanded we confirm it held for all nomologically possible ones. Together, these led to law confirmation.

Of course, our initial target was "for whatever reasons, in all actual cases, ignited sodium *fluoride* causes a yellow flame". Unlike the case of sodium fluoride *glass* where, by hypothesis, we're antecedently confident that all the glass is relevantly causally similar, we're explicitly concerned that all sodium fluoride is not. That's why we divided it into causally individuated kinds of potential disconfirmers—the glass form, the different crystal forms, and so on. Are we led to confirm a law in this case too?

A *causally individuated kind of Potential Disconfirmer* (hereafter, *kind of PD*) for a given causal generalization is a kind for which (i) we have *a high degree of confidence* that the subsumed cases are relevantly causally similar, and (ii) there is a single individuating factor that we *believe* might be causally relevant to the production of disconfirming cases. To well confirm the causal generalization for sodium fluoride, we minimally need to well confirm that each kind of PD could not be a source of disconfirming cases. That is the process we explicitly followed for sodium fluoride glass—one kind of PD for the sodium fluoride generalization—that demanded the confirmation of an associated law. However, for several reasons, well confirming a law corresponding to every kind of PD does not inevitably lead to well confirming an overarching sodium fluoride law.

If we could justify holding that some nomologically possible kind of PD is non-actual, we wouldn't need to well confirm the sodium fluoride law to well confirm the causal generalization for actual sodium fluoride. However, just as our actual epistemic position precluded us from restricting actually occurring additional factors that might cause non-yellow burning sodium fluoride glass beyond those that are nomologically possible, it likely precludes us from excluding any nomologically possible kind of PD from actuality. In any case, we shall assume that to be so. To well confirm the sodium fluoride causal generalization, we must include every kind of PD we hold to be *nomologically possible* in our set of kinds of PDs. Hence, well confirming the sodium fluoride causal generalization at least demands well confirming the corresponding nomological necessity.

Easier said than done. We can coherently have high confidence that each member of a set of kinds of PDs does not yield disconfirmers but low confidence that the set does not. As witnessed by the Preface paradox, it's often rational to combine low confidence in a conjunction with high confidence in each conjunct. Moreover, even if we can well confirm the necessity, we won't well confirm a corresponding law if the laws associated with different kinds of PDs support different causal explanations of the same behavior**.**

For the devotee of law confirmation, there's a pleasing synergy here—the nomic character of the hypotheses in question plays a crucial role. If we don't have some good reason to *believe* that there are distinct kinds of causes of yellow burning among the nomologically possible kinds of sodium fluoride, then we should be willing to *countenance the possibility* that all sodium fluoride is relevantly causally similar. If we then well confirm that each law holds for each kind of PD, we in effect eliminate all our explicit reasons for *suspecting* that any kind of sodium fluoride in our entire set of kinds of PDs is relevantly causally dissimilar. Thus, by something very much like an application of Mill's (1868) method of agreement, we may reasonably be highly confident that all, in fact, are relevantly similar i.e., something common to all the sodium fluoride covered in our set of kinds of PDs causes yellow burning. And since that set is the set of nomologically possible kinds of PDs, we thereby well confirm the unified causal generalization for all the nomologically possible kinds, and hence, the law.

At this point, we can clearly see why, *for predictive reasons alone*, scientists might reasonably be interested in confirming laws. In our epistemic position, if we thought it unlikely that all sodium fluoride glass, say, is relevantly causally similar, we'd likely have little reason to expect "all ignited sodium fluoride glass causes a yellow flame" to be true, and so, scant reason to research it. By contrast, if we strongly suspect all such glass is relevantly causally similar, we've an excellent *prima facie* reason to suppose that some informative generalization regarding the color with which it burns is true, and moreover, *confirmable*, since high confidence in relevant causal similarity provides a rational basis for making inductions from modest data. So, the scientist's predictive goals make it reasonable to target broad generalizations we strongly suspect subsume cases that are relevantly causally similar; and given our epistemic position, only those for which we suspect similarity across all nomologically possible cases i.e., suspected laws. On a grander scale, it makes sense to target "all sodium fluoride burns yellow", given that we suspect—to a more modest degree, given its greater scope—that all sodium fluoride is relevantly causally similar across all nomologically possible kinds of cases i.e., again, given that we suspect a corresponding law. As per our discussion, the confirmation of such a broader law is more demanding, and the prospect of failure inevitably higher. Nevertheless, scientists have an excellent predictive motivation to confirm laws.

4. Confirming Laws More Generally

We shall now consider some tougher cases. We'll ultimately discuss Newton's law of gravitation, which involves both complex causal dependencies and continuous variables, but we begin with a law that subsumes the sodium fluoride one, the law that all sodium salts burn yellow. This might seem like a trivial variation. However, it exhibits a practical problem that isn't present in the sodium fluoride case: there's a limitless array of chemically distinct sodium salts, each of which, at least initially, might reasonably be a kind of PD.

For instance, you can take any carboxylic acid and treat it with sodium hydroxide to form a sodium salt, and there's an unlimited variety of the former. A carboxylic acid is a molecule with a COOH group (i.e., a suitably bonded group consisting of one carbon atom, two oxygens, and one hydrogen) bonded to an alkyl, a radical that consists of an alkane (roughly, a molecule of the form C_nH_{2n+2}, for arbitrary integer values of $n \geq 1$) with one hydrogen atom removed. We can't, even in principle, gather data that individually covers every chemical variant in the family of sodium salts generated by carboxylic acids. So, how could we well confirm the sodium salt law that covers this, and indeed, other unlimited families of chemically distinct salts?

Again, rational commitments regarding relevant causal similarity come to our rescue. The similarities between different kinds of alkane molecules and the resultant similarities of the associated carboxylic acid molecules give us some reason to *suspect* the resultant sodium salts might have some causal similarities, but they're insufficient to *assume* that all will burn similarly. However, as we gather data regarding the increasingly complex sodium salts formed from increasingly complex carboxylic acids, and find they invariably burn yellow, at some point we can very reasonably cease to believe that chemical differences among that family of sodium salts provide a reason to suspect disconfirmation. That is to say, the remaining chemically distinct salts in that family no longer qualify as kinds of PDs, since they no longer satisfy condition (ii) for such kinds. Also, where initially we weren't confident that all sodium salts in that family are relevantly causally similar in regard of color of burning, we can now reasonably be confident that they are. Hence, members of that family now satisfy condition (i). We've consolidated an unlimited set of kinds of PDs into one set for which we're rationally highly confident that the subsumed cases are relevantly causally similar, and indeed, given our observations, all burn *yellow*.[7] We might in a similar way consolidate the set of kinds of PDs initially associated with sodium salts formed from sulfonic acids, and so forth. More generally, we can practically well confirm a law that, relative to our initial background beliefs, subsumes an unlimited variety of kinds of PDs by a series of such consolidations. An initially intractable research program evolves, through unification of kinds, into a manageable one.[8]

[7] Some might think we need not go through this elaborate process to well confirm that sodium salts burn yellow, since sodium characteristically burns yellow and one thing all these compounds have in common is a sodium constituent. But they would be wrong. Granted that sodium in isolation burns yellow, we still need evidence that being bonded into such molecules doesn't relevantly affect the sodium, and also, that other parts of the molecules don't make their own differently colored contributions to the resultant flame color.

[8] Such consolidation might be related to a notion of dynamical consilience, although not one subsumed by the model in Thagard (1978).

We can apply a similar kind of story to Newton's law of gravitation, $F = GMm/r^2$. An agent's set of kinds of PDs depends on her background beliefs. For simplicity, assume she thinks only factors mentioned in the equation are likely causally relevant to the attractive force between masses i.e., for all nomologically possible cases, she is confident that only the values of M, m and r, are relevant. Further, for each triple of values of M, m, and r, she (i) has a high degree of confidence that all cases with the same triple of values will cause the same force, but (ii) believes that differences in any of those values might be causally relevant to the production of disconfirming cases. So, for her, each triple of values specifies a kind of PD, and observation of a modest number of cases of a given kind for which F is indeed specified by Newton's equation will well confirm that it holds for that kind of PD. A rational agent may hold that well confirming that the law holds for each of a significant variety of her kinds of PD licenses rejecting the belief that a novel triple of values provides reason to suspect disconfirmation and become correspondingly confident that all such triples are relevantly causally similar. Such an agent consolidates her entire set of kinds of PDs into one class of cases for which she's highly confident that the equation holds i.e., she well confirms Newton's law.

A rational agent might, of course, be *vastly* more skeptical. It will depend upon her background beliefs, and we can certainly understand rational disagreement among scientists on this score. Some might have reason to favor consolidation only over some proper subsets of the kinds; and all who resist wholesale consolidation will presumably only well confirm that the law holds to some degree of approximation in the domains for which they hold it successful—if a law like this is wrong in some cases, it's plausibly wrong in all, or almost all, even if the magnitude of error is often miniscule. Nevertheless, given appropriate background beliefs about potential causal dependencies, we can see how the two mechanisms, well confirming laws associated with individual kinds of PDs, and consolidation of kinds of PDs via changes of background belief, might accommodate the confirmation of laws that invoke equations. Such confirmation in this case broadly resembles Mill's (1868) method of concomitant variation.

We don't wish to understate the difficulty of law confirmation on this methodology. Establishing highly general scientific claims *is difficult*. Individual scientists propose laws, but their confirmation is often the work of legions over decades or even centuries. While we might quite rapidly confirm the sodium fluoride law, well confirming the sodium salt law demands the accumulation of a wealth of causal knowledge regarding an immense variety of forms of sodium salts and other related compounds. The same goes for the variety of interestingly different systems relevant to the law of gravitation. Consolidating kinds of PDs renders well confirming such laws a practical possibility, but it may be a *very* long road indeed.

We've only applied our methodology to a few cases but can reasonably be optimistic about its extension to others e.g., laws confirmed using something like Mill's method of difference. However, we clearly can't attempt a general treatment here. Nor shall we attempt to formalize our methodology. It involves a set of *prima facie* reasonable, but non-trivial, epistemic transformations, and even if we confine ourselves to some species of Bayesianism, we might proceed in multiple ways.[9] Our informal account suffices to make a compelling case that law confirmation provides an important and distinctive way of meeting the predictive goals of science, and we don't want to give the mistaken impression that our case is hostage to the details of some particular, inevitably somewhat tendentious, formalization.

[9] Notably, the rejection of background beliefs involved in the consolidation of kinds of PDs demands careful treatment, since you can't downgrade a proposition's probability from 1 by conditionalization. However, since background beliefs commonly get discarded, this is an issue Bayesians must grapple with anyway. Plausibly, they cannot maintain conditionalization as an entirely general policy for updating.

5. Relevant Causal Similarity, Law Confirmation, and Scientific Practice

Does science plausibly proceed in this way? Scientific theorizing largely involves positing entities we suspect might causally explain the observed phenomena—electrons, atoms, molecules, antibodies, hormones, psychological states—and associated causal generalizations and laws. Scientific confirmation largely concerns those generalizations, and they're not stitched together from a patchwork of parochial causal claims: they support explanations that go hand in hand with the broad relevant causal similarity that underpins our confirmation theory. And when scientists pursue their broadest predictive ambitions, they confirm laws. So, the primary objects of confirmation in science are of the right kind for our methodology.

Moreover, our confirmation theory also seems a reasonable fit. A scientist who sought to confirm the law that all ravens are black but didn't recognize the possibility that arctic environments might well *cause* white (hence non-black) ravens as providing a reason to specifically research arctic ravens would be obviously incompetent. Thus, the potential relevant causal dissimilarity that delineates our kinds of PDs clearly matters for scientific confirmation. Further, if we're not confident that all arctic ravens are relevantly causally similar, it's hard to see how we could make a rational induction regarding their blackness based on the observed ones.[10] It's not as if we can *literally* randomly sample the population of arctic ravens, or any population that includes them. It has countless past and future members that have zero probability of being in our sample. Thus, if we're to justify holding the population we can sample as representative of that total population, and hence legitimately make an induction regarding all the arctic ravens being black, we must be confident the total population doesn't include members that *relevantly causally differ* from those in the sampled population; if we're not confident of this, we must recognize our

[10] Both of these considerations plausibly have implications regarding standard random sampling resolutions of the Raven paradox, notwithstanding that such resolutions concern the non-nomic generalization, not the law. See Ward (forthcoming).

sample as potentially biased. So, the judgements of representativeness that are required to legitimately apply a random sampling model—or indeed, any reasonable model of an induction regarding the blackness of the total population of arctic ravens—are parasitic on our confidence that past and future arctic ravens are relevantly causally similar to those we can observe. The confidence in relevant causal similarity that unifies each kind of PD clearly matters for rational induction even in this simple case. Moreover, our inability to *literally* randomly sample the populations of sodium salts and massive objects throughout spacetime—and the attendant need for judgements of relevant causal similarity to justify judgements of representativeness required for inductions that cover those total populations—is even more apparent. So, there's a good *prima facie* case that scientific confirmation commonly involves our judgements of relevant causal similarity and potential relevant causal dissimilarity.

Obviously, this is a cursory assessment. However, it suffices for a *prima facie* comparison with the accounts of Goodman and the non-Humeans—Armstrong, Dretske, and Tooley. For them, law confirmation is a *sui generis* mode of induction, one that involves their respective notions of lawlikeness and non-Humean necessitation relations, neither of which are manifestly implicated in scientific practice.[11] On ours, it's directly integrated with the workaday business of confirming causal claims. Now, it's not required that scientific practice wear its heart on its sleeve, but other things being equal, an account that appeals to manifest features rather than ones that have no *evident* role, is to be preferred. Moreover, barring Tooley (as per footnote 2), they erroneously demand that law confirmation is generally required for induction, whereas we recognize our approach as merely providing one inductive methodology. We don't deny the manifest legitimacy of random sampling as commonly deployed in the sciences—and in Sober's and van Fraassen's examples (cited in footnote 1)—to inductively confirm accidental generalizations.

[11] As Lewis (1983, 366) observed, calling something 'necessitation' doesn't make it genuine nomic necessitation any more than one can have mighty biceps just by being called 'Armstrong'. Similarly, Goodman's notion of lawlikeness should not be casually identified with any notion of law operative in scientific practice.

On the other hand, the very considerable attention paid to non-nomic generalizations as opposed to laws in twentieth and early twenty-first century confirmation theory should give confirmation theorists pause.[12] Laws are a primary focus of much scientific activity and we never had strong reasons to suppose their nomic character was generally incidental to their confirmation. So, their largely peripheral status in the confirmational literature is certainly odd and may well be a source of considerable misunderstanding of scientific practice.

Indeed, it's no part of our account that we should expect to find working scientists explicitly engaging our initial question of how to well confirm non-nomic generalizations, like Reichenbach's. Science is concerned with predictive comprehensiveness, but since scientists primarily concern themselves with nomic claims, its predictive goal will plausibly be addressed by scientists in terms of the laws themselves. Just as our epistemically impoverished position led us to confirm individual laws to ensure predictive comprehensiveness regarding actual sodium salts, scientists seek to confirm that the world conforms to all the nomologically possible kinds of PDs generated by their entire nomic apparatus. In that way they have their guarantee of predictive comprehensiveness, whatever contingencies the vast and varied universe might produce.

Why then begin our discussion with a non-nomic generalization? Primarily, to persuade *philosophers* that law confirmation can make a distinctive, valuable contribution to prediction. Further, the Reichenbach example served as a simple illustration of our primary methodology for well confirming generalizations, whatever their character: considering ways in which disconfirming cases can be caused. As subsequently applied to causal generalizations in particular and conjoined with the use of relevant causal similarity as a basis for induction, that became our methodology for confirming laws.

[12] We don't have quantitative data for the comparative claim, but a perusal of the mainstream confirmational literature will bear it out. To take one telling example, it's surely remarkable that there's no explicit treatment of law confirmation in Howson and Urbach's (2006) fine, and quite comprehensive, text on Bayesian philosophy of science. Indeed, the terms "law", "scientific law", and "natural law" don't even appear in the index.

6. The Prospects for an Analysis

Laws are entailed by the specified causal generalizations and that's all our confirmational story needs. We could, however, attempt an analysis by identifying them. Roughly,

> A law is a unified causal generalization that invariantly holds for all nomologically possible cases.

Where, by 'unified', we mean a causal generalization that subsumes relevantly causally similar cases—the feature caused is caused by the same factor(s) in each case. Such an analysis has *prima facie* problems. It seems *circular* both in its appeal to the set of nomologically possible cases and its demand that the causal generalization *would* hold for those possibilities, implicitly invoking counterfactuals that are commonly taken to be grounded by the laws. Metaphysicians will also find it crucially *incomplete*, since it analyses laws in terms of causation but gives no account of the latter. Is the associated ontology Humean, non-Humean, grounded in causal powers or dispositions? It doesn't say.

James Woodward has an analysis that faces similar objections:

"Laws are generalizations about repeatable relationships that are invariant over variations in initial and other sorts of conditions, at least within an appropriate range of such variations—invariant in the sense that laws will or would continue to hold under such variations." [13]

Both Bird (2007) and Psillos (2004) raise both concerns and we can benefit from Woodward's defense of his own project regarding laws and causation more generally. He characterizes his analysis as a *functional* one: its primary purpose is to illuminate the role that laws play in service of various goals or purposes we have. Such an analysis succeeds insofar as it provides associated insights regarding the kinds of evidence

[13] Woodward (2018, 158). Woodward (2003), (2013), and (2020) provide additional discussion.

required to confirm laws and the norms for assessing law claims.[14] From the functional perspective, circularity and incompleteness are beside the point.

Our analysis fits the functional template quite well. We motivated our approach with respect to the goal of prediction of the actual, and the result was a substantive—although, obviously incomplete—*prima facie* plausible confirmation theory. Indeed, we've some cause to be optimistic about its development *qua* functional analysis. Contemporary research on causal confirmation is both highly sophisticated and flourishing[15], and integrating our account of law confirmation with this work is an interesting prospect.

It also illuminates our norms for distinguishing laws from non-laws. Consider "All solid spheres of U-235 have a diameter of less than 1 mile", an example commonly paired with Reichenbach's generalization to raise problems for naïve regularity analyses of laws.[16] Colossal spheres of U-235 cannot be caused in any nomologically possible case, and for the same reason—aggregates above critical mass would very rapidly explode due to a runaway thermonuclear chain reaction.[17] By contrast, suitable aggregates of gold *would* cause colossal spheres by gravitational accretion. So, no corresponding causal generalization invariantly holds across all nomologically possible cases. Further, the reasons a Jupiter-sized aggregate of hydrogen, say, and a collision of neutron stars would not cause such gold spheres are not subsumed under any *unified* causal generalization: in the former case, the aggregate's size is insufficient to even cause nuclear fusion; in the latter, the violence of the collision causes the gold produced by fusion to be extremely rapidly dispersed.

[14] Woodward (2014) discusses this approach to causation in detail, contrasting it with other species of analysis. Woodward (2018, 166-7, 180) further alludes to it regarding laws in particular.
[15] See, for instance, Spirtes, Glymour and Scheines (2000), Pearl (2009), and a wealth of more recent additions to this research program.
[16] See van Fraassen (1989, 27) and numerous other discussions.
[17] Strictly speaking, the U-235 law should be a probabilistic one that specifies the probability of such colossal spheres as very close to zero. A thermonuclear explosion is not guaranteed, but its probability is exceedingly high for any nomologically possible formation process. However, we shall ignore this complication here.

The analysis doesn't just get the right answers in these cases, but does so for the right reasons, citing the considerations that would be used by anyone versed in the relevant physics. Further, in line with the functional approach, these normative judgements can be ultimately justified by the generalizations' respective fitness and unfitness for achieving our goal of generally reliable prediction of the actual, given our impoverished epistemic state. As it happens, our current beliefs about relevant boundary conditions make it plausible that Reichenbach's generalization is actually true, but an encounter with technologically sophisticated aliens with a perverse interest in very large gold spheres could change that.

What are the prospects for a metaphysical analysis? Woodward (2018, 177-80) strongly suspects the metaphysicians' aspirations are misguided, because the pursuit of truth-makers for laws is an ill-posed problem. We could continue to follow him down the functional road, and it would not be a dead end. Developing the analysis by incorporating something along the lines of his (2003) interventionist analysis of causal explanation is another interesting prospect. At the least, it should provide a foil for his analysis of laws. On the other hand, we might attempt to eliminate the incompleteness and circularity by providing truth-makers for causal relations and counterfactuals that are not grounded in the laws and analyzing the set of nomological possibilities as the set of possibilities that can be caused. However, even a preliminary consideration of such proposals is beyond our current scope.

7. Conclusion

There are familiar reasons to care about causal knowledge. It provides effective strategies for manipulating the world and facilitates understanding. To this we add another: it provides a distinct pathway to broad predictive knowledge. In one respect, that's hardly an original observation. Causal generalizations are central to Hume's epistemology for matters of fact with which we're not directly acquainted. However, we've added three pieces to that story: the causal methodology for well confirming generalizations, its development into the relevant causal similarity methodology for well confirming causal generalizations, and how using this latter methodology to fulfill the sciences' broadest predictive ambitions leads to law confirmation.

Moreover, in showing how law confirmation can be prosecuted using familiar methodologies we apply to causal claims, we've domesticated it, largely reducing it to something familiar and commonsensical. No-one can reasonably deny the role of causal inference in daily scientific and folk practice. Why then deny laws their rightful place in our understanding of confirmation?

References

Armstrong, D.M. (1983). What is a Law of Nature? (Cambridge: Cambridge University Press).

Bird, A. (2007). Nature's Metaphysics: Laws and Properties. (Oxford: Oxford University Press).

Dretske, F.I. (1977). Laws of Nature. Philosophy of Science 44, 248-268.

Goodman, N. (1983*).* Fact, Fiction, & Forecast (4th edition). (Cambridge, MA: Harvard University Press).

Hempel, C.G., (1945). Studies in the Logic of Confirmation (I). Mind 54, 1-26.

Howson, C. and Urbach, P. (2006). Scientific Reasoning: The Bayesian Approach (third edition). (Chicago and LaSalle, Illinois: Open Court).

Lewis, D. (1983). New Work for a Theory of Universals. Australasian Journal of Philosophy 61, 343-77. Reprinted in Lewis 1999, 8-55.

Lewis, D. (1999). Papers in Metaphysics and Epistemology. (Cambridge: Cambridge University Press).

Mill, J.S. (1868). A System of Logic, Ratiocincative and Inductive Vol. I (7th edition). London: Longmans, Green, and Co.

Pearl, J. (2009). Causality: Models, Reasoning and Inference (second edition). (Cambridge: Cambridge University Press).

Psillos, S. (2004). A Glimpse of the Secret Connexion: Harmonizing Mechanisms with Counterfactuals. Perspectives on Science 12, 288-319.

Salmon, W.C. (1989). Four Decades of Scientific Explanation. In W. C. Salmon and P. Kitcher (eds.) Minnesota Studies in the Philosophy of Science 13, (3-219). (Minneapolis: University of Minnesota Press).

Sober, E. (1988). Confirmation and Law-Likeness. Philosophical Review 97, 93-98.

Spirtes, P., Glymour, C., and Scheines, R. (2000). Causation, Prediction, and Search (second edition). (Cambridge, MA: MIT Press).

Thagard, P. (1978). The Best Explanation: Criteria for Theory Choice. The Journal of Philosophy 75: 76-92.

Tooley, M. (1987). Causation: A Realist Approach. (Clarendon: Oxford University Press).

van Fraassen, B.C. (1989). Laws and Symmetry. (New York: Oxford University Press).

Ward, B. (forthcoming). Resolving the Raven Paradox: Simple Random Sampling, Stratified Random Sampling, and Inference to the Best Explanation. Philosophy of Science (forthcoming).

Woodward, J. (2003). Making Things Happen: A Theory of Causal Explanation. (New York: Oxford University Press).

--- (2013). Laws, Causes, and Invariance. In S. Mumford, S. and M. Tugby (eds.) Metaphysics and Science. (Oxford: Oxford University Press).

--- (2014). A Functional Account of Causation. Philosophy of Science 81, 691-713.

--- (2018). Laws: An Invariance Based Account. In W. Ott and L. Patton (eds.) Laws of Nature. (New York: Oxford University Press).

--- (2020). Physical Modality, Laws, and Counterfactuals. Synthese 197, 1907-1929.