Physically Contingent Laws and Counterfactual Support

Michael Strevens

Philosopher's Imprint 8(8)

Abstract

The generalizations found in biology, psychology, sociology, and other high-level sciences are typically physically contingent. You might conclude that they play only a limited role in scientific investigation, on the grounds that physically contingent generalizations offer no or only feeble counterfactual support. But the link between contingency and counterfactual support is more complex than is commonly supposed. A certain class of physically contingent generalizations, comprising many, perhaps the vast majority, of those in the high-level sciences, provides strong counterfactual support of just the sort that appears to be scientifically important. This paper explains why.

1. Contingency in the High-Level Sciences

The generalizations that do predictive and explanatory work in the high-level sciences—the laws of biology, psychology, geology, and so on—appear to be physically contingent, often to an extreme: not only are they not entailed by the fundamental laws of physics, they require for their truth an antecedently quite improbable array of initial conditions.

Consider, for example, generalizations attributing a trait to a biological taxon, such as *All normal ravens are black*, or *The HIV virus has a highly mutable outer coat*, or *Adult humans' inferences treat natural kinds as though they have essences*. If these generalizations are true, it is only because the members of the taxa in question—ravens, HIV, humans—have very particular properties, or you might say, a very particular structure. That they have this structure is a physically contingent fact; ravens, for example, might have evolved brown rather than black plumage. Thus the truth of the generalizations is physically contingent.

Even some laws of physics are contingent. According to Newton's law of cooling, a substance that is hotter than its surroundings cools at a rate proportional to the temperature differential. You will not, in your lifetime, see the law violated, but it is nevertheless not a consequence of the fundamental laws. Rather, it is a consequence of the fundamental laws together with the fact that the initial microstate of the kind of system to which it applies is highly likely to be of a certain class, which you might call *entropic*, since Newton's law is a nomological avatar of the second law of thermodynamics. Anti-entropic initial conditions are physically possible; indeed, on the standard understanding of statistical mechanics, if you wait long enough, you will see temperature differentials widen just as often as you see them, as Newton's law says they must, narrow (though entropic conditions predominate in our region of space-time). Newton's law is thus not even close to being a physical necessity. The same can be said for any "high-level"—that is, non-fundamental—law that depends on entropic processes, thus on entropic initial conditions. Among scientific laws, physical contingency is the rule, necessity the exception.

Ought you to be concerned? There is, I think, a certain anxiety about contingency to be found in the philosophy of the higher level sciences, perhaps felt most keenly by the writers who have worked hardest to draw attention to its existence. Beatty (1995), for example, argues that because physical necessity is a hallmark of laws of nature, the contingent generalizations of biology must lack lawhood. Rosenberg (2006) carries the point further: he agrees with Beatty that contingency undermines lawhood, proposes that a generalization must be a law to explain, concludes that contingent generalizations lack explanatory power, and so reaches the view that in any biological domain whose regularities are contingent, all explanation must proceed at the physiochemical level. Weber (2005) largely concurs.

More positively, other writers, moved by the same anxiety, have attempted to restore necessity to biological and other high-level generalizations. Lange (2000, 2004) suggests that every domain of investigation has its own contextually determined facts about necessity, and that within a given domain, scientifically important generalizations will be necessarily true. The same generalization that is contingent from the perspective of fundamental physics, then, will be necessary from the perspective of the domain of inquiry to which it matters. Lange further hints that this necessity plays a central role in undergirding the counterfactual support offered by a domain's distinctive generalizations—or as Lange is happy to call them, once their domain-relative necessity is established—laws.[1]

1. On Lange's view, necessity within a field is relative to the counterfactual suppositions of interest to the field, and lawhood in a field is a matter of stability under such suppositions.

I should add that Lange has a powerful theory of the relation between laws, modality, and inductive inquiry that provides independent motivation, aside from any anxiety about contingency, for his positing of multiple nomological necessities.

Given the topic of this paper, I should also note that, unlike the other authors cited here, he acknowledges explicitly that all contingent states of affairs have some degree of counterfactual robustness; however, he seems to think that this robustness lacks systematicity

Waters (1998) takes a different route to restoring biological generalizations' modal standing, arguing that generalizations that are on the face of things physically contingent—such as *Broadleaf trees tend to grow canopies that cast shadows on rival saplings*—are in fact necessary. How can this be? Suppose that the tenebrous tendency of broadleaf trees is due to most or all such trees having a physical property *P* that causally necessitates canopy production of the appropriate sort. It is contingent that broadleaf trees have *P*, and thus contingent that they grow the canopies in question. How, then, can the canopy generalization be physically necessary? Waters argues that, contrary to appearances, the canopy generalization expresses roughly the proposition *All things with P tend to grow such-and-such canopies*.

All five writers are moved, it seems, by the thought that the apparent contingency of high-level generalizations has ruinous implications for their scientific status. What might the implications be? What scientific or metaphysical roles are exclusively or at least predominantly the domain of necessities? Rosenberg and Weber point to explanation, and Weber explicitly connects explanatory power to counterfactual support (p. 30).

I venture that counterfactual support lies at or near the core not just of Weber's but of everyone's angst about contingency. A generalization that is highly physically contingent, it is supposed, can provide at best only fitful support for counterfactual conditionals—certainly not the kind of robust, across-the-board support that something serving the predictive and

---

and significance: there is nothing *special*, as he puts it, about the range of counterfactual suppositions with respect to which a typical contingency holds (Lange 2000, 15). Further, implicit in this and other passages I detect a reluctance to attribute to contingencies, however robust, an ability to provide counterfactual support. They may themselves have a certain modal persistence, but it is passive rather than active; what *explains* such persistence is never something contingent.

Finally, let me mention Lange's interesting view that physical necessity does not entail biological necessity: just as some biogeographical laws are physical contingencies, so some physical laws are biogeographical contingencies—they are not mandated or presupposed by the biogeographical laws.

explanatory role of a law or near-law should provide.[2] Hence my topic in the present paper: the ability of physically contingent generalizations to support counterfactual conditionals.

It is not a case of all or nothing. On the one hand, it is clear that physical contingency does not entirely obviate a generalization's counterfactual-supporting power: If these two ravens had bred last year, their offspring would have been black. If I had put an ice cube in my mouth five minutes ago, it would have melted at a rate proportional to the temperature difference between ice and tongue. On the other hand, there are some counterfactual conditionals with physically possible antecedents that a physically contingent generalization cannot support, namely, conditionals counterfactually supposing the absence of the contingent grounds of the generalization itself: If the course of raven evolution had gone differently, ravens might very well not have been black.

What, then, is the relation between a generalization's contingency and its ability to offer counterfactual support? You might think that, as contingency increases, the power to support counterfactuals declines in roughly inverse proportion. This proportionality thesis, I suggest, underlies unease about contingency in the high-level sciences: it entails that highly contingent generalizations can offer only the most meager counterfactual support, certainly too little to be of any broad predictive or explanatory use.

But the proportionality thesis is false, as I will show in the course of examining the foundations of contingent generalizations' support for counterfactual conditionals. Further, it is false in the most interesting possible way: it turns out that highly contingent generalizations supply a degree of counterfactual support almost as great as that offered by physical necessities. The scientific cost of dealing in contingent rather than necessary generalizations, then, is not nearly as great as some have supposed—indeed, it is in

2. Weber (2005, 30) writes that a regularity's providing counterfactual support "indicates" its nomic necessity, from which it follows that without necessity, there is no such support.

many cases close to zero.

Although it is essential to this paper to show that, say, *All ravens are black* provides strong and systematic (though not exceptionless) counterfactual support, I do not take the establishment of this conclusion to be the paper's main business. The existence of the support can be demonstrated informally but effectively by running through a few examples; further, I suspect that writers such as Lange and Waters are well aware of the scope of the support. What I have to offer that is new is my explanation for the phenomenon. The generalization *All ravens are black* provides counterfactual support not because, as Lange and Waters suppose, it has or inherits a kind of necessity (biological necessity for Lange; veiled physical necessity for Waters). On the contrary, it is in every respect just as contingent as it seems; its contingency, however, does not seriously undermine its ability to support counterfactual conditionals.

It would be interesting to go further and show that, in virtue of this counterfactual support (or perhaps, in virtue of the facts that underlie the support), contingent generalizations may be attributed considerable explanatory power, or may be counted among the laws of nature. But that would require the establishment and defense of criteria for explanatoriness and lawhood, efforts that are better pursued elsewhere (e.g., Strevens 2009). The reader of the present paper will have to be content with the removal of a powerful prima facie objection to the explanatoriness and lawhood of contingent generalizations: their putative counterfactual impotence.

* * *

A terminological issue: the term *law* is rather contentious. Some writers take the fundamental laws of nature as their paradigms, using *law* to pick out only deep, wide-ranging, unifying generalizations. Other writers apply the term also to those shallower or more phenomenological generalizations whose very names proclaim their lawhood—Boyle's law, Fitts' law, Mendel's laws, and so on. I propose that we call the latter class the *nomes*, in honor of the great linguist. A nome, I stipulate, is any scientifically significant regularity, in the broadest sense of *regularity* and in a sense of *significant* to be determined by the scientists themselves—so that the nomes of biology, for example, are whatever kinds of regularities might be found playing predictive and explanatory roles in the biological literature.

Think of the term *nome*, then, as a shallow but convenient term, rather than as an attempt to pick out a deep, philosophically integrated category. Certainly, nomehood as I have so sketchily characterized it is no replacement for the kinds of categories found in the philosophical literature on lawhood, or for philosophical attempts to abandon or revise the notion of lawhood, such as Mitchell's graded conception of the nomological (Mitchell 2002). My introduction of the word is, in short, an attempt to avoid rather than to comment substantively on the question of the nature of lawhood—except perhaps in one respect. I do assume that the question of counterfactual support can be addressed independently of, and perhaps should be addressed prior to, the question of lawhood.

There is a further ambiguity in the term *law*: it is sometimes used to refer to a feature of the physical world, and sometimes used to refer to a scientific representation of such a feature. A nome is a physical feature; from this point on, I use the term *generalization* to refer to a representation of a nome or other regularity, and not to the regularity itself.

I will restrict my attention to a certain class of nomes, those picked out by generalizations that describe the consequences of a causal mechanism. I call these the *causal nomes*. In the next section, I sketch an account of the physical grounding of causal nomes that should clarify their nature. I am inclined to think that all, or at least almost all, scientifically important high-level "laws" are causal nomes; from this point of view, the paper's restriction to causal nomes is not a significant limitation. You need not sign on to this imperialist program; what matters is that causal nomes raise the problem of contingency and counterfactual support in a perspicuous way.

This paper will ignore entirely one class of high-level laws that is clearly not physically contingent: laws whose obtaining is logically entailed by the functional nature of the properties that they relate, a feature commonly attributed to, for example, the laws (if that is the right term) of folk psychology.

2. Causal Nomes

A statement of a causal nome describes (I stipulate) the consequences of a particular causal mechanism, such as the mechanism by which ravens acquire their black plumage or by which hot objects cool at a rate proportional to the temperature differential. Two things make such a statement true: a certain pattern of initial conditions, and the fundamental laws of physics in virtue of which the initial conditions constitute a mechanism producing the stated consequence.

Consider, for example, the raven nome, *All (normal) ravens are black*. Normal ravens are black because

1. All normal ravens possess a certain constellation of physical properties, properties that constitute a mechanism for producing black plumage,

2. The various background conditions required for the operation of this blackness mechanism usually obtain, and

3. The fundamental laws of nature entail that, provided that the background conditions obtain, the constellation of properties produces black plumage.

Two remarks. First, what I have given is a purely physiological reason for blackness in ravens; there may in addition be further, evolutionary reasons for blackness, in particular, evolutionary considerations that account for the fact that ravens possess the blackness-producing mechanism to begin with. A focus on the physiological foundations of raven blackness will suffice for my purposes in this paper.[3] Second, I will for simplicity's sake largely ignore the importance of the background conditions that are the topic of clause (2); there is contingency enough to be found in the fact of the universality of the mechanism stated in clause (1).

Putting aside background conditions, then, the raven nome holds because

1. All normal ravens possess a certain constellation of physical properties *P*, and

2. The fundamental laws of nature entail that *P* causes blackness.

I call the pattern stated in (1)—the *P*-hood of all (normal) ravens—a *basing pattern*. The raven nome holds, then, because of the combined effect of first, a certain basing pattern, and second, the fundamental laws of nature.

For the sake of tidiness, let me stipulate that the blackness-producing structure *P* includes everything about the raven that contributes at any point to producing blackness, including the mechanisms of reproduction by which the relevant genes (for example, the gene for tyrosinase, an enzyme that catalyzes the production of the pigment melanin) are passed on from parent to offspring. This ensures that every physical property of ravens on which their blackness depends, either directly or indirectly, is integrated into the basing pattern, so that when I talk about the basing pattern, I am talking about all the physically contingent properties of ravens on which the nome depends.

My focus in this paper is on nomes of the form *All Fs are G* that hold for reasons that are similar, schematically speaking, to the reasons that the raven nome holds, that is, nomes of the form *All Fs are G* that hold because

1. All *F*s possess a certain constellation of physical properties *P*, and

3. An evolutionary pressure towards blackness will both decrease the raven nome's degree of contingency and increase the range of counterfactuals that it supports (by removing some exceptions due to early antecedents and late consequents; see section 5).

2. The fundamental laws of nature entail that *P* causes *G*,

or in short, because all *F*s possess a *G*-producing mechanism *P*.

Some more terminology: I will call a statement of a basing pattern, such as (1), a *basing generalization*. And I will call that aspect of the fundamental laws in virtue of which a structure *P* causes a property *G*—the sort of fact described by (2)—a *causal-mechanical law*. (I hope that there will be no confusion between causal nomes, which are often physically contingent, and causal-mechanical laws, which are by definition physically necessary.)

Why define a causal-mechanical law so that it is a function only of the fundamental laws, and not of higher level laws (or causal nomes)? For now, because I want to separate the contingent from the physically necessary constituents of a nome, and later, because the fundamental laws play a special role in evaluating counterfactuals (see note 5).

The raven nome, and other nomes of the form *All Fs are G*, are a special case, I think, of nomes of a more general form that might be captured as follows, where *S* is some class of systems: *Within S, all Fs are G*. Such nomes make their universal claim not about just any *F*s, but about *F*s found in a particular kind of environment. They do not state entirely general consequences of *F*-ness, in other words, but only consequences that emerge in the causal framework supplied by systems of type *S*. Mendel's laws of genetics provide an example: the claims they make about genes are restricted to sexually reproducing life on earth.

A nome relativized to a class of systems *S* holds in virtue of two basing patterns and a causal-mechanical law:

1. All *F*s in *S* possess a constellation of physical properties *P*,

2. All *S*s (or at least, the ones containing *F*s) possess a constellation of physical properties *Q*, and

3. The fundamental laws of nature entail that *P* and *Q* together cause *G*.

The schema for causal nomes can be generalized in other ways. For example, a causal nome's universal claim might be relativized not only to a certain kind of system but also to other background conditions, yielding a nome of the form *In such-and-such conditions and within class S, all Fs are G.* Or a causal nome might assume a functional form stating the dependence of one variable on another. But the very simple cases examined here are sufficient to raise the issues I wish to discuss in this paper; additional complications would enhance realism but would not bring any concomitant philosophical advantages.

It is often supposed that statements of causal nomes must be qualified by a ceteris paribus hedge, or by some other special language that narrows or otherwise weakens the claim made by the statement to allow for the vagaries of high-level processes. (Perhaps the raven nome's restriction to "normal" members of the species exemplifies this sort of qualification.) There are a wide range of proposals for understanding ceteris paribus hedges, as well as arguments for thinking that they are in fact unimportant or dispensable (Earman et al. 2003). I take it that on all relevant proposals, contingent causal nomes will turn out to hold in virtue of some basing generalization or other, and some causal-mechanical law or other, with the form of the basing generalization and the causal-mechanical law determined differently by different accounts of the significance of ceteris paribus hedges. My explanation of causal nomes' counterfactual support depends only on the fact that a causal nome is founded in a contingent basing pattern and a necessary causal-mechanical law, and not on the particular identity or form of the pattern or law; it is compatible, then, with any account of the significance of *ceteris paribus*, and I will accordingly ignore the complications introduced by such hedges in the rest of this paper.

⋆ ⋆ ⋆

Which nomes are causal? Which nomes, that is, satisfy the basing-pattern-plus-causal-mechanical-law schema laid out above? Certainly, all the taxon-trait generalizations, such as *The HIV virus has a highly mutable outer coat* and *Adult humans' inferences treat natural kinds as though they have essences*, are true in virtue of basing patterns and causal-mechanical laws: the basing pattern is in each case the universality or near-universality of a certain mechanism in the normal members of the taxon in question, while the causal-mechanical law encapsulates those aspects of the fundamental laws in virtue of which the mechanism ensures the presence of the trait.

Some more general biological nomes also fit the mold. For example, as I have already remarked, generalizations about inheritance, such as Mendel's laws and the later more sophisticated generalizations of classical genetics are true in virtue of the prevalence of certain kinds of inheritance mechanisms in sexually reproducing life on earth, along with the physics that governs the operation of the mechanisms.[4]

The social sciences too, I would argue, deal in generalizations that have basing patterns and causal-mechanical laws as truthmakers, not least because they depend so much on psychological nomes, which are a species of taxon-trait nome. Some writers would resist this conclusion for reasons given earlier: the psychological nomes, and many other social-scientific nomes, they would argue, hold in virtue of the functional nature of psychological and social kinds such as *belief* or *democracy*. They overestimate the scientific importance of functional properties, I think, but this is not the place to give the reasons why; it is enough, in what follows, that a good proportion of high-level nomes are causal nomes.

4. Waters (1998) provides a more extended defense of the view that many nomes in biology have basing patterns and causal-mechanical laws as their foundation. He prefers to break down the nomes into their constituent parts, understanding biology as principally concerned with two quite different classes of regularities, basing patterns and causal-mechanical laws, but not with a third class of regularities—my causal nomes—that melds the two. (As noted above, he holds that what appear to be statements of causal nomes in fact pick out causal-mechanical laws.) Waters' bifurcation, it will turn out, misses something of great importance: the counterfactual support offered by the raven nome and others like it depends in part on the synergy between the basing pattern and the causal-mechanical law.

High-level nomes are physically contingent because the basing patterns on which they in part depend are physically contingent. Ravens might not have had the blackness-causing constellation of properties *P*; they might have evolved some other coloration mechanism. Humans might not have been quasi-essentialist in their thinking about natural kinds; their minds might have been structured so as to engender a more empiricist psychology. It is fruitful, I suggest, to understand the special sciences as systematic explorations of the causal consequences of the contingent but pervasive presence, in one part of the universe or another, of certain mechanisms, structures, or even (in the case of Newton's law of cooling and other manifestations, precisifications, or localizations of the second law of thermodynamics) initial conditions. The special sciences, then, deal with contingency just as much as they deal with causality; causal nomes, being one part basing pattern and one part causal-mechanical law, delicately fuse the two, and as such are the proper object of special science research.

3. The Nature of Counterfactual Support

For a nome of the *All Fs are G* variety to support counterfactual conditionals is for the following sorts of conditionals to be true:

> *If x had been F, it would have been G.*
>
> *If a new F had come into existence in such-and-such a way, it would have been G.*
>
> *If conditions had been different in such-and-such a way, then x, an F, would still have been G.*

Further, they should be true for the same reasons that the nome holds, that is, not by accident but in virtue of the basing patterns and causal-mechanical laws that provide the nome's foundation. (Of independent interest are negative counterfactuals of the form *If x had not been F, it would not have been G*; they will not, however, be discussed in this paper.)

Only a subset of the counterfactuals conforming to the above schemas are of scientific importance. Morton (1973) quite reasonably remarks that there is not much scientific value in ascertaining the truth or otherwise of *If I were a dry well-made match, I would light if struck* (see also section 5.4). Woodward (2003, 279–285) goes further, developing a criterion to determine which of the counterfactuals supported by a nome contribute to its explanatory potential. And the practical importance of counterfactuals surely depends to some degree on human needs and values. For simplicity's sake, I will rely in this paper on a loose and intuitive feel for what counterfactuals matter to science, rather than any kind of deep philosophical characterization of the scientific significance of counterfactual reasoning.

In order to understand counterfactual support, we need— of course— an account of the truth conditions for counterfactual conditionals, preferably one that is simple and uncontroversial. I will adopt the appealing account developed by Bennett (2003), omitting a few of its more sophisticated elements.

Bennett's account does not claim to give a complete theory of counterfactuals. It rather describes what is close to a consensus as to how the ultimate truth conditions, whatever they may be, function when applied to a class of counterfactuals which include most everyday conditionals as well as, more significantly here, all or almost all the counterfactual conditionals that are of immediate interest to science— such as, for example, *If I had released the martini, it would have fallen to the ground*, and *If the 2005 avian flu had crossed over to humans, it might well have resulted in a pandemic*, and *If Archduke Ferdinand had not been assassinated, war would anyway have broken out not much later*.

The Bennett algorithm functions as follows. To assess a counterfactual conditional of the form *If c had occurred at t, then e would have occurred*, turn your attention to the following sort of possible world. Up to a time shortly before *t*, the possible world is identical to the actual world. At that point, a small deviation from the actual world's history takes place, as a result of which the course of events in the possible world diverges from the actual course of events just enough that the counterfactual antecedent *c* occurs at *t*, without any further unnecessary changes. After *t*, the world strictly obeys the actual laws: the course of events follows the path that the fundamental laws of our world prescribe, given the state of the possible world at time *t*.[5] In an indeterministic possible world, the deviation that brings about *c* may be due to some irreducibly stochastic process's having an outcome different from its outcome in the actual world. Otherwise, the deviation will involve what Lewis calls a miracle: a violation of the actual world's laws.

Call the set of possible worlds best satisfying the previous paragraph's description—that is, the worlds satisfying the description in virtue of the most conservative possible deviation from the actual world before *t*—the counterfactual's *evaluation worlds*. (On the method for quantifying the conservativism of a deviation, I refer you to Bennett's chapter 13 and Lewis's work on counterfactuals (Lewis 1973, 1979). Suffice to say that the criteria take into account to some extent the probability of the deviation (where relevant), the degree to which the fundamental laws are violated (where relevant), the closeness of the time of the deviation to the time of the antecedent *c*'s occurrence, and the side effects of the deviation.)

The counterfactual is true if its consequent *e* occurs in all, or almost all, of the evaluation worlds. The "might" version of the counterfactual—*If c had occurred at t, then e might have occurred*—is true if *e* occurs in a set of

5. Here I follow Lewis in supposing that it is the fundamental laws only that matter after *t*. You might alternatively suppose that some additional, higher level laws also make a contribution; in determining the truth of a biological counterfactual, for example, you might insist that after *t* the world follows not only the fundamental laws but also some "biologically basic" (but physically contingent) laws. You could then go on to give precisely the sort of story offered below about the counterfactual support offered by causal nomes that are not "biologically basic".

evaluation worlds that have non-negligible probability.[6]

For later reference, let me underline the fact that there are three temporal stages to any of a counterfactual's evaluation worlds, each bearing a different relation to the actual world: first, the stage up to the moment when the deviation from the actual course of events begins, second, the stage between the beginning of the deviation and the occurrence of the counterfactual's antecedent, and third, the stage after the occurrence of the antecedent. Call these stages I, II, and III; they are shown in table 1.

	Duration	Relation to actual world
Stage I	Up to shortly before occurrence of antecedent	Identical to actual history
Stage II	Up to occurrence of antecedent	Deviation from actual history sufficient to bring about antecedent, changing as little else as possible
Stage III	From occurrence of antecedent on	What the actual fundamental laws prescribe, given the conditions holding at the end of stage II

Table 1: The three stages of the Bennett evaluation worlds

As noted above, the Bennett algorithm cannot be applied to all counterfactuals. It requires that the antecedent take place at some particular time, and so fails to deal with, for example, the conditional *If gravity had obeyed an inverse cube law, Kepler's second law would still have held.* It also requires that the antecedent be accessible by way of a conservative deviation, and so perhaps fails to handle, for example, *If Caesar had directed the invasion of*

6. The "almost all" and the "non-negligible probability" in the main text point imprecisely to some necessary qualifications of the otherwise refreshingly simple view that a "would" counterfactual requires the consequent to occur in every evaluation world and a "might" counterfactual requires it to occur in at least one. There is much to say here, but let me simply remark that in evaluating counterfactuals, we tend to ignore outlandish possibilities of the sort traditionally perpetrated by typewriting monkeys (Bennett 2003, chap. 16). For a discussion of alternative views of "might" conditionals, see Bennett (2003, §73).

Iraq, there would have been no looting of the archeological museum (though Bennett argues that his full treatment applies to many conditionals of this sort). Finally, it does not deal with what Lewis calls backtracking counterfactuals, conditionals that implicitly or explicitly call for the ramifications of a counterfactual antecedent to be propagated into the past as well as into the future.[7] Some of these cases—the first in particular—are of scientific interest (though the fact about Kepler's second law and others like it can be expressed, I think, without recourse to counterfactual conditionals). But most scientifically significant counterfactuals, including the examples given earlier in this section, are amenable to the Bennett approach; it is, then, a solid foundation for a discussion of counterfactual support.

* * *

A note on the relation between a nome's ability to offer counterfactual support and its counterfactual robustness, that is, its obtaining under various counterfactual suppositions. In general, it is more or less sufficient for the truth of *If x had become F, then it would have become G* and so on that the causal nome *All Fs are G* holds in every evaluation world (*more or less* because the nome might allow exceptions). It is not necessary, however: it is possible that, though the nome does not hold throughout an evaluation world, it holds in the relevant locale, and that its holding locally is sufficient to support the counterfactual.

To see this, consider the counterfactual *If these two ravens were to have bred, their offspring would have been black.* Suppose that the course of raven evolution depends very sensitively on current initial conditions, so that changing even the most insignificant-seeming details in the present has the capacity to divert the course of future raven evolution, perhaps resulting, millennia

7. Additionally, it does not cope with cases in which there is no most conservative deviation bringing about *c* because, for every conservative deviation, there is one that is even more conservative. This and other such niceties are dealt with in Lewis (1973). Bennett, too, treats this issue, and discusses several other classes of counterfactuals that are not covered by the simple version of his algorithm discussed here.

from now, in ravens becoming mottled or dark gray. Then in some evaluation worlds for the counterfactual, the changes required to bring about the counterfactual antecedent, the raven mating, will also bring about the evolution of mottled or gray raven plumage in the world's distant future. The raven nome, because it is temporally unrestricted, does not hold in such worlds. Yet, in these same worlds, the ravens' present-day counterfactual offspring are black, and they are black for the same reason that the raven nome holds (as I will explain in detail below). So although the raven nome fails somewhere in the evaluation world, it provides genuine support to the counterfactual in question because it holds over the local region of the world relevant to the evaluation of the counterfactual—because it holds, as it were, for the present if not for the future of the raven species.

There has been far more work on nomes' counterfactual robustness than on their ability to offer counterfactual support, ranging from research concerned to establish the reasons for high-level nomes' contingency (Beatty 1995; Mitchell 2002), through essentialist arguments for the metaphysical necessity of nomes that appear to be contingent (Ellis 2001; Bird 2005), to Lange's (2000) project of understanding the principal properties of lawhood in terms of laws' invariance under certain sets of domain-relative counterfactual suppositions. None of this work attempts a direct or constitutes an entirely satisfactory indirect explanation of why contingent causal nomes support counterfactuals (though a heterodox interpretation of the raven and other taxon-trait nomes as "inference tickets" is offered by Lange (1995, 440–442)). Robustness and support may be closely related, then, but inquiries into the one question tend to take place more or less independently of inquiries into the other, a tradition that I will for brevity's sake perpetuate here.

4. How the Raven Nome Provides Counterfactual Support

The counterfactual conditional *If these two (normal) ravens were to have bred at time t, their offspring would have been black* is true; further, it is not true by coincidence but rather in virtue of the facts that found the raven nome *All (normal) ravens are black.* This is as straightforward a case as could be of a physically contingent nome's offering counterfactual support. How does it work?[8]

The Bennett algorithm evaluates the raven counterfactual according to the following procedure: Find the most conservative deviations in the actual world's history, starting shortly before t, that result in your two ravens mating at t. For each such deviation, feed the exact resulting state at t to the actual laws. If the laws mandate black offspring in (almost) every case, the counterfactual is true.

Applying the algorithm, then, the counterfactual turns out to be true because

1. The potential raven parents have property P (in the actual world),
2. The more conservative alterations of the actual world that would lead to the parents' mating do not affect the parents' P-hood, and
3. According to the actual causal-mechanical laws, raven parents' P-hood first replicates itself in their offspring (recalling from section 2 that P includes relevant reproductive mechanisms), then causes the offspring to grow black plumage.

By (1), the parents in the evaluation worlds have P until shortly before t; by (2), the counterfactual processes that lead to their mating in the evaluation worlds leave their P-hood intact; by (3), their P-hood then causes them to have offspring that are black.

8. I have specified a particular time to simplify the discussion; I leave it as an exercise for the reader to generalize to the case where a range of times for mating is allowed.

To put things in more general terms, facts (1), (2), and (3) respectively guarantee that everything goes smoothly in the three temporal stages of a counterfactual's evaluation world, which I called in section 3 stages I, II, and III; see table 1 (p. 16). Fact (1) ensures that the right initial conditions obtain at the end of stage I, fact (2) ensures that the deviation from actuality in stage II does not interfere with these conditions, and fact (3) of course ensures that, during stage III, the conditions cause the counterfactual consequent. Let me say something more about the three facts and their roles in the three stages.

Fact (1), the actual *P*-hood of the raven parents is carried over without amendment to the evaluation worlds. This is not due to any special modal dominion that *P* has over the physical possibilia, but for a purely parochial reason: ravens have *P* in the actual world, and the possible worlds that are relevant to the evaluation of counterfactuals are those that are identical to to the actual world up until shortly before the antecedent occurs. In evaluating counterfactuals, then, the actual facts have a certain "modal inertia" by which they are transferred by default to the evaluation worlds, quite regardless of their degree of contingency. This inertia's capacity to set up appropriate initial conditions in evaluation worlds is the single most important element of the explanation of contingent causal nomes' ability to offer counterfactual support.

Fact (2) might seem to hold almost by definition. What makes an alteration conservative is its affecting absolutely as little as possible beyond what is necessary to bring about the counterfactual antecedent in a timely way; conservative deviations, then, will not affect the raven parents' *P*-hood. The availability of deviations that are conservative in this sense is not, however, guaranteed. They exist, when they do, for three reasons.

First, *P*-hood does not depend either causally or metaphysically on the occurrence or otherwise of mating, so the facts about mating may in principle be manipulated without thereby manipulating the facts about *P*-hood (contrast with, say, two traits realized by overlapping aspects of the same physiological structure).

Second, we live in a world in which small localized changes have, in the short term, small localized effects; if it is possible in principle to manipulate one property without affecting another, then, it is normally possible in practice.[9] This locality principle holds in part because of the way the world's fundamental laws work; arguably, it holds also because the world is entropic, that is, because the world conforms to the second law of thermodynamics. However, the case for this claim is complicated, and it depends on the way in which the conservatism of deviations is evaluated, so let me not dwell on the importance of entropic initial conditions here.

Third, *P* has at least a small degree of stability, in the sense that if it is not actively interfered with, it will tend to stay around, if only for the short amount of time that elapses between the deviation from actuality and the consequent occurrence of the counterfactual antecedent. In ravens, there are biological reasons for the durability of *P*; what about other causal nomes? Sufficient durability is in general rather likely, since the relevant property *P* features in a basing pattern—it is present in most or all *F*s (for fanciful exceptions, see section 6). Indeed, the stability normally required to underwrite a basing pattern is far greater than the stability required to prevent *P*'s disintegration in the short term. So although, for example, the basing pattern may be maintained by the kind of active causal process that Boyd (1999) calls "homeostasis", even without such nurturing a typical *P* might well be expected to endure for the brief interval spanned by a typical stage II.

9. More exactly, manipulations have small, localized *macrolevel* effects. They may have more diverse and diffuse microlevel effects. For example, although a deviation that brings about a mating between ravens will not change the fact that the ravens have *P*, it will almost surely change the exact microlevel realization of the *P*-hood, by altering the rate at which various molecules vibrate, and so on. It will also change the way things are at a level that is higher than the microlevel in the technical, statistical-mechanical sense—for example, it will change the positions of macrolevel objects, not least the ravens.

To fact (3), then: the causal-mechanical laws. In a way, there is nothing to say about the importance of the laws that is not obvious. It is a consequence of the Bennett algorithm that the causal-mechanical laws, being facets of the fundamental laws of physics, hold for the duration of stage III in the evaluation worlds for any counterfactual for which the algorithm is valid.[10] They are sure to be there whenever the raven nome, or any other causal nome, needs them. Furthermore, they are essential if the raven nome is to provide counterfactual support, as without them *P*-hood is not converted into blackness.

Nevertheless, if you are impressed by the power of modal inertia, you might wonder if the causal aspect of causal nomes adds anything special to their ability to provide counterfactual support. Let me divide this worry into two parts.

First, does the fact that causal-mechanical laws are causal give them any advantage over non-causal physical necessities, when it comes to explaining counterfactual support? No. Any physical necessity holds through stage III; that the necessities relevant to this paper are causal is merely a side effect of my restricting my attention to causal nomes, nomes whose physically necessary part is by definition causal. (That said, it is not easy to think of examples of interesting contingent generalizations true in virtue of a basing pattern and a non-causal physical necessity.)

Second, does the fact that causal-mechanical laws are physically necessary distinguish them from contingent regularities, when it comes to explaining counterfactual support? Does a causal nome provide more counterfactual support, all other things being equal, than a contingent basing pattern, which after all enjoys the benefits of modal inertia? Yes. Of the three temporal stages relevant to the evaluation of a counterfactual, modal inertia affects the course of events in the first two but not the third. It is unable to shepherd the

10. The causal-mechanical laws may be locally violated outside of stage III by the "small miracle" that sets off the conservative deviation from actuality at the beginning of stage II.

counterfactual process from antecedent to consequent in stage III; here the causal-mechanical laws play their part, in virtue of their physical necessity.

To see this, consider the following imaginary scenario. Color in ravens is caused by one of two mechanisms, *Q* and *R*; any given raven has one or the other, but not both. The properties are transmitted to offspring down the maternal line, so a raven has the same mechanism as its mother. One mechanism, *Q*, always causes its possessors to be black. The other, *R*, sometimes causes them to be black, sometimes white. (In case you need to know, the process is irreducibly stochastic.) Yet although some ravens have *R*, by a strange twist of fate all actual ravens are black.

Now take a pair of ravens; suppose that the female has *R*. Consider the truth values of the following two conditionals:

1. If the ravens mated, their offspring would have *R*.

2. If the ravens mated, their offspring would be black.

By the Bennett algorithm, the first is true and the second—because in some evaluation worlds, the offspring's *R*-hood will cause it to be white—is false. I take this judgment to be in agreement with our intuitions. Thus the fact that all ravens are black, despite its universal scope, holds no sway over the counterfactual offspring. Why not? A world in which an offspring's *R*-hood caused it to be black would arguably be more similar to the actual world than a world where its *R*-hood caused it to be white, because a world in which all ravens are black is more like the actual world than one in which some ravens are white. But this similarity counts for nothing, because it develops in stage III. All that matters in stage III is what the causal-mechanical laws prescribe, given the state of the world at the end of stage II; similarity's role is exhausted once the stage II state is determined. Contrast with the first counterfactual; although the mother's *R*-hood is not only accidental but is not part of some larger pattern (most ravens, let's say, have *Q*), it is carried over to the offspring in stage III by the causal-mechanical laws.

In short: the modal inertia of basing patterns is good only as far as the end of stage II (and even persistence through stage II is not guaranteed, for the reasons given above). But the truth of a counterfactual depends on how things stand at the end of stage III. For this reason, a basing pattern on its own does not provide counterfactual support. It must be supplemented with something that establishes a relationship between the conditions at the end of stage II and those at the end of stage III, and that is modally robust enough to hold in physically possible worlds for reasons that go beyond similarity—something like a causal-mechanical law. Causal nomes support counterfactuals, then, because they unite two complementary kinds of robustness: their basing patterns establish the relevant initial conditions through modal inertia, and their causal-mechanical laws convert initial conditions to consequent events through physical necessity.

My canonical raven counterfactual concerns the color of a counterfactual raven, that is, a raven that does not actually exist. But as noted at the beginning of this section, some conditionals relevant to the question of counterfactual support concern the color of actual ravens in non-actual circumstances: *If this raven had been raised on a vegetarian diet, it would still have been black*, and so on. The raven nome supports these kinds of counterfactuals for the same reasons that it supports counterfactuals about non-actual ravens: by modal inertia alone, the actual P-hood of the raven in question carries over to the evaluation worlds in stage I (note that you no longer need to fall back on the P-hood of the parents); the conservative deviations in the evaluation worlds do not undermine P, which is therefore still present at the end of stage II; by the causal-mechanical laws it then causes blackness in stage III.

On the basis of the examination of counterfactual support offered above, it is possible to see, though perhaps only dimly, how the raven nome might support (as Lange (2000) has urged) even some counterlegals, that is, counterfactuals with physically impossible antecedents, for example: *If these two* *ravens had been teleported instantly to the same love nest, the offspring resulting from their mating would have been black*. The Bennett algorithm very likely must be amended, or at least made more sophisticated, to deal with conditionals of this sort, but even putting the details aside, you can see how the *P*-hood of the actual parents might be carried over, by modal inertia, to a physically impossible world where instantaneous teleportation can be arranged. Then, provided that the relevant causal-mechanical law also survives in the new world— and such a law is typically consistent with a wide range of physical underpinnings— certain raven blackness counterlegals will be supported by the raven nome for much the same reasons that more mundane conditionals enjoy the nome's support.

5. When the Raven Nome Fails to Provide Support

The raven nome does not support every counterfactual that it might. Let me discuss four cases in which support breaks down, in order to better understand the scope and limits of contingent nomes' modal power.

5.1 Mechanism Failure

A part of the structure *P* that constitutes the mechanism by which ravens acquire their color is the gene for tyrosinase. (As already mentioned, tyrosinase catalyzes the production of melanin.) If the gene is damaged, a raven typically cannot produce the pigment; this is the cause of albinism. Thus the following counterfactual, uttered pointing to two ravens with normal tyrosinase genes, is false: *If these two ravens were to have mated, and were something to have gone wrong in the transmission of their tyrosinase genes to their offspring, their offspring would have been black*.

Why does the support fail? You might think that the counterfactual's antecedent, in stipulating that transmission is unsuccessful, supposes that the causal-mechanical laws misfire. This is possible if transmission is a stochastic process: even physically necessary laws do not support counterfactuals whose antecedents specify that some essential, irreducibly stochastic process falls through. I want to focus, however, on the case where the causal-mechanical laws are deterministic, and so where their "misfiring" is impossible.

If transmission fails in a deterministic world, then something has gone wrong with *P*, understood to be the complete set of conditions required for transmission to succeed. Thus a deterministic evaluation world in which transmission fails is one in which, at the end of stage II, some element of *P*, hence *P* itself, is missing. (Here you should understand *P* as a property of the two parents and the environment as a whole, rather than, as previously, a property possessed separately and in entirety by each of the parents—a harmless equivocation, necessary among other reasons because this paper for the most part ignores the role of background conditions.) The cases in which the raven nome will not support a counterfactual because of mechanism failure or "misfiring" are, in other words, cases in which *P* disappears during stage II.

Such cases come in two varieties. First, there are cases in which the counterfactual antecedent entails that the relevant raven does not have *P*: in the evaluation worlds for such a counterfactual, *P* will be present at the end of stage I, but then the deviation that brings about the antecedent of the counterfactual will of necessity undo *P* by the end of stage II. Second, there are cases in which the antecedent, though it does not entail the absence of *P*, is nevertheless brought about most conservatively by deviations that also undermine *P*. Mechanism failure in other causal nomes can be understood in the same way.

You might wonder whether a causal nome, when properly understood, has an obligation to provide support in cases of mechanism failure. I have in mind the following line of argument: as noted in section 2, any real special science nome has certain conditions of application that restrict its scope. The raven nome, for example, does not apply to all ravens but only to normal ravens (or at least, that is one way to understand it). The antecedents of the sorts of counterfactuals under consideration, you might think, violate these conditions of application. The counterfactual scenarios they envisage, then, do not fall within the scope of the nome, and thus the conditionals' falsehood does not constitute a genuine failure on the part of the nome to provide support.

This is, I think, the correct way to understand the falsehood of some such counterfactuals. But not all of them: the only way to protect a nome from every antecedent implying, logically or causally, the absence of *P* is to have the nome's conditions of application entail (perhaps in conjunction with the fundamental laws) the presence of *P*. But then, of course, the nome would not be physically contingent after all. Every physically contingent nome, then, fails to provide complete counterfactual support in at least some cases because its mechanism fails—because the conditions required for the operation of its underlying mechanism are somehow subverted by the process actualizing the counterfactual antecedent.

5.2 Early Antecedents

The contingency of the raven nome is due in large part to the contingency of the events surrounding the appearance of the raven species (or so I will suppose). So as to vividly introduce a second class of counterfactuals unsupported by the raven and other causal nomes, let me indulge in some free speculation about this speciation event. My story is quite likely false, but something like it is true—so biologists believe—of many traits in many species; it is biologically realistic, then, at one remove.

Back in ornithological prehistory, a few members of some ancient corvid species were separated from the rest by a geological incident: perhaps they migrated to a distant island, or perhaps the territory that once united them with their conspecifics became impassable. The species in question had plumage that varied from mid-gray to black; however, in the separated group, black was by chance more common than in the species as a whole. As time went on, the isolated subpopulation changed in response to the local environment; it became raven-like. Black remained the predominant, but not the only, coloring. Then some sudden environmental change came close to driving the new ravens to extinction; as so often happens in such events, a certain amount of variation in the species was lost. In particular, the ravens became exclusively black, not because black coloring offered any advantage in dealing with new challenges, but simply by chance.

Now consider counterfactuals such as the following: *If the early raven population had not gone through the bottleneck, present-day ravens would still have been black* or *If ravens had evolved from a somewhat different founder population, present-day ravens would still have been black.* Such counterfactuals are false, hence not supported by the raven nome. The reason, of course, is that their evaluation by way of the Bennett algorithm requires that the early history of the raven population be rerun with new initial conditions; since it was only a matter of chance that this process led to universal raven blackness in the actual world, there is no guarantee, or even much of a probability, that a counterfactual process will produce the same outcome. More generally, a nome *All Fs have G* will tend not to support counterfactuals having antecedents that occur before any *F*s come into existence, unless the actual conditions at around the time of the antecedent are already set up so as to ensure in advance the existence of the relevant basing pattern.

As you will see, causal nomes' failure to provide support for counterfactuals with "early antecedents" is closely connected to the nomes' contingency; indeed, it is the falsehood of such counterfactuals that we typically use to diagnose nomic contingency. This is the one case where there is a direct connection between contingency and the lack of a certain kind of counterfactual support, a connection that is further discussed at the end of section 6.

5.3 Late Consequents

A further group of counterfactuals not supported by the raven nome and other causal nomes are those in which the interval of time between the occurrence of the counterfactual's antecedent and its consequent is too long. Consider, for example, the following conditional: *If these two ravens were to have bred, their descendants a million years later (should there be any) would have been black*. On the assumption that the course of raven evolution is sensitive to initial conditions (see the end of section 3), the counterfactual is false. The descendants might be black, but then again, they might not be, since a deviation that causes the two ravens, contrary to the actual facts, to mate, might well set off a chain of events that causes the raven species, again contrary to the actual facts, eventually to evolve a mottled color. There are some evaluation worlds for the counterfactual, then, in which ravens evolve to become mottled; because of these worlds, the counterfactual does not hold true.

What has gone wrong? That a mating pair of ravens have the basing property *P* is sufficient to guarantee the blackness of the next generation, and the generation after, and the generation after that . . . but at some point, the power of *P* winds down; thus, ravens' present *P*-hood does not preclude a gradual evolutionary change in their color. Consequently, counterfactuals of the form *If this pair of ravens were to have bred, the* n^{th} *generation of their descendants would have been black* are true for low values of *n*, but false for high values. (For what value of *n* does the counterfactual become false? That depends on what it means for the consequent to hold in "almost all" the evaluation worlds; see note 6.) More generally, whenever a basing pattern is contingent, and so may be eroded with the passage of years, causal nomes founded on the pattern will support counterfactuals less securely as the interval of time between antecedent and consequent lengthens.

Why does *P*'s ability to guarantee the consequent decline, as it were, with age? Why does the *P*-hood of mating ravens guarantee the blackness of their immediate offspring, but not the blackness of their n^{th}-generation offspring when n is large?

The basing property *P*, properly understood, includes many environmental conditions required to guarantee that everything goes smoothly in the transmission of the blackness-producing mechanism from one generation to the next, and equally smoothly in the mechanism's production of blackness itself. I assume these conditions acquit themselves successfully in this respect; it does not follow, however, that they are entirely self-perpetuating. In particular, they may not guarantee that the same environmental conditions will hold for the next generation, that is, for the grandchildren. If not, they give you one but not two generations of blackness—though they may well entail a high probability for second-generation blackness.

Additionally, because of the possibility of random mutations in the raven genome, it may be that *P* does not strictly speaking guarantee the blackness of even the next generation, but only confers extremely high probability on blackness, high enough for the raven counterfactual to qualify as true (again, see note 6). As the number of generations between antecedent and consequent increases, the probability of the successful transmission of blackness would then get whittled away, until (though still high) it dips below the threshold required for counterfactual truth.[11]

You might see a way to rescue causal nomes from this failure to provide counterfactual support. Surely, if *P* were augmented with sufficiently many further conditions, its reach over the generations could be extended as far as you like? And since I am assuming that the raven nome holds in the actual world, could this strengthened version of *P* not be veridically attributed to all ravens? In short, is there not a way of adding conditions to *P* to obtain a new and stronger property *P**, such that, first, all (actual) ravens have *P**,

11. The explanation given in the main text is intended to generalize to most or all causal nomes. I have ignored, or at least relegated to an abstraction, a very important feature of the biological case: in any generation, the ravens may not survive to reproduce—though extinction will of course in any case fossilize rather than falsify the raven nome.

and second, given the P^* of present generations of ravens, the blackness of all future generations is assured? (Again I am equivocating between a P that is a property of individual ravens and a P that is a property of the entire system, ravens plus environment. Again, no philosophical harm is done, no unfair advantage obtained.)

If the world is irreducibly stochastic in relevant ways—for example, in such a way as to make genetic mutation stochastic—the answer to this question is certainly negative: no set of initial conditions can guarantee the blackness of ravens for all time. But suppose that the world is deterministic, and so that such a P^* is available. What then? Take a closer look at the conditions added to P to obtain P^*. It is quite plausible, I think, that some of these conditions will involve very finely specified configurations of properties across a large region of space. As such, they stand to be easily disrupted by the sort of small deviation required for the two raven parents counterfactually to mate. Thus in some evaluation worlds at least, though P^* obtains at the end of stage I, by the end of stage II it no longer holds; it is undermined by the deviation that brings about the counterfactual antecedent. The counterfactual claim therefore falls through. Even in a deterministic world, then, a failure to provide counterfactual support is quite possible for any causal nome when the interval between antecedent and consequent is large.

5.4 Unnaturally Generated Ravenhood

When explaining the blackness of counterfactual ravens in section 4, I restricted myself to ravens produced in the biologically normal way by other ravens. It is because the raven parents are actual, and have P, that modal inertia ensures the presence of P in the counterfactual scenario in which they breed.

What if counterfactual ravenhood is not produced by ravens? Consider (with a nod to Morton) the conditional *If I were magically transformed into* *a raven, I would be black.*[12] Is it true? Arguably not, on the grounds (I am making this up as I go along) that at least some relatively conservative human-to-raven transformations result in a raven of a color partway between black and the original skin tone of the former human. And even if the counterfactual is true, its holding is not due to the facts invoked in section 4; it is due rather to the specifically thaumaturgical fact that the most conservative human-to-raven transformations endow the transformed creature with blackness (and not necessarily by way of *P*).

The kind of counterfactual support explained in section 4, then, applies only to those non-actual ravens that are naturally produced. More generally, a causal nome attributing a trait *G* to a biological taxon *F* will support counterfactuals about non-actual *F*s for the reasons given in section 4 only when the non-actual *F* is naturally produced, that is, produced in the standard way by other *F*s. However, as I remarked at the beginning of section 3, counterfactuals about unnatural *F*s are perhaps of limited scientific significance.

5.5 Summary

I have identified four reasons that a contingent causal nome may fail to provide counterfactual support:

1. Every contingent basing pattern allows the possibility of exceptions, that is, *F*s that do not have *P*. Some antecedents pick out evaluation worlds in which the salient *F*s are just such exceptions, either by specifying explicitly that *P* is absent or by specifying an antecedent that is most conservatively brought about at the expense of *P*.

2. Early on in the actual world, there are (typically) no *F*s. Evaluation worlds for counterfactuals having antecedents occurring at these

12. Put aside the entirely reasonable worry that this is not the sort of counterfactual whose truth is to be evaluated using the Bennett algorithm.

primeval times need match the actual world only for a period of time in which there is no pattern of the *P*-hood of *F*s; if the basing pattern is contingent, some such evaluation worlds will have futures in which few or no *F*s have *P*.

3. The presence, at any one time, of the basing property *P* does not guarantee the world's conforming to the nome for indefinitely long periods. Counterfactuals in which there is an inordinate gap between antecedent and consequent may not, then, enjoy the nome's support.

4. If a counterfactual *F* is produced in an abnormal way, there is no causal process by which the basing pattern carried over to the evaluation worlds in stage i can ensure the *P*-hood of the *F* created during stage ii.

6. Beyond Taxon-Trait Nomes

The treatment of the counterfactual support offered by the raven nome can be extended to any causal nome satisfying the characterization given in section 2. For the most part, the generalization is a matter of replacing the relevant properties in section 4's account—ravenhood, blackness, and the all-powerful *P*—with appropriate placeholders, since almost nothing in the explanation turns on the specific nature of these properties.

With one exception: the raven basing pattern and the basing patterns of other taxon-trait nomes are causally maintained by the mechanisms of biological inheritance, thanks to which *P*-hood is passed down the generations from *F* to *F*. This fact plays an important part in explaining the ability of taxon-trait nomes to support counterfactuals concerning non-actual *F*s: the non-actual *F*s inherit their *P*-hood from their parents, who being actual, have it by modal inertia. Is the ability of a causal nome to support counterfactuals compromised if its basing pattern is not maintained by inheritance? Or more generally, is counterfactual support compromised if the basing pattern is not actively maintained by what Boyd calls a homeostatic mechanism?

Perhaps surprisingly, it is difficult to find useful examples of basing patterns not perpetuated by inheritance, because the causal nomes of so many of the high-level sciences have their roots in biological patterns. I have in mind not only nomes of explicitly biological sciences like genetics and population ecology but also, for instance, the nomes of developmental psychology and sociological nomes concerning suicide or educational underachievement, which have basing patterns that depend on properties of the human mind maintained by inheritance mechanisms, both genetic and cultural.

To see past inheritance-supported basing patterns, it is better to look below the level of biology. Consider, for example, the various geological generalizations true of our planet because its outer layer is composed of tectonic plates: generalizations about sea floor spreading, hot spots, patterns of earthquakes, mountain building, and so on. The geological nomes articulated by these generalizations are founded in part on contingent facts about the earth's structure, including in some cases facts about the particular plates and their relations to one another.

When evaluating a geological counterfactual about our planet, your evaluation worlds will share the actual history of the planet, and thus will attribute to the earth the same tectonic structure in stage I; this structure will then persist in the evaluation worlds' stage II. Nomes that depend on the structure are therefore quite capable of supporting counterfactuals.

That the earth has roughly the same plate structure that it did hundreds of millions of years ago, and that this structure survives typical counterfactual suppositions, is not an accident. But it is not due to anything that could be called inheritance. Rather, the plate structure is simply very robust in the face of the prevailing planetary forces; it has a certain "physical inertia". Such inertia is as good as inheritance in securing counterfactual support.

You might think that the geological basing "pattern" is really a single, one-off event, something to do with the formation of the earth that entails the existence of the tectonic structure at later times. This is, I think, a mistake, just as it would be a mistake to think of the raven nome's basing pattern as being the single event of the founding subpopulation all coming to have *P*. Why a mistake? Because in general, such a one-off event will not quite succeed in entailing the necessary pattern of initial conditions; or more exactly because, for reasons given in section 5.3, in order to secure such an entailment too many highly fragile initial conditions would need to be built into the event, initial conditions that might not survive even the conservative deviations required to bring about the antecedents of typical counterfactual conditionals. As a consequence, a nome founded on the one-off event will not provide the same degree of counterfactual support as a nome founded on the pattern brought about by the event.

Note, by the way, that the geological generalizations are system-relative, in the sense characterized in section 2: their scope is implicitly restricted to the planet earth. In this respect they are like many of the biological generalizations that are not taxon-trait generalizations, which state the consequences of causal mechanisms that are characteristic of life on earth, but which might not be found in life elsewhere.

So much for geology; let me give you a nome of physics with a basing pattern that is not maintained by a mechanism of inheritance, namely, Newton's law of cooling, alongside which may be placed every other nome that depends on entropic initial conditions (which would be pretty much all the nomes of the high-level sciences). Our world has so far enjoyed entropic conditions, starting as it has from a very low entropy state and then proceeding more or less smoothly up the entropy gradient to its present state. The evaluation worlds for counterfactuals share this history. By itself, having such a history does not preclude a world's present-day conditions (more exactly, its state at the end of stage II) turning out to be, for the first time, negentropic. Indeed, for any counterfactual, there will be evaluation worlds with relevantly negentropic initial conditions (that is, conditions that are, at the beginning of stage III, negentropic in the vicinity of the relevant causal processes). But entropic initial conditions are overwhelmingly likely to generate more entropic conditions, thus, there are very few such worlds: the majority of evaluation worlds are entropic, by so great a margin that the negentropic worlds will make no difference to the truth values of counterfactuals, which hold provided that their consequents obtain in "almost all" of the evaluation worlds (note 6). It is for this reason that Newton's law offers counterfactual support.[13]

Newton's law's basing pattern—the entropicity of initial conditions—is not maintained by an inheritance mechanism. But it is actively perpetuated nevertheless, in the sense that there is a causal mechanism or mechanisms that, given entropic conditions at one time, in almost all cases produce entropic initial conditions at a later time.

What the raven basing pattern, the entropic basing pattern, and the geological basing pattern have in common, you will see, is that their persistence can be given a systematic causal explanation (though one that invokes contingent initial conditions); where they differ is the degree to which the relevant causal facts might be said to be positively propagating the pattern. But even a causal nome that depends on a more or less accidental basing pattern may offer some amount of counterfactual support.

How much support? How accidental, how contingent, can a nome's basing pattern be, without destroying its ability to offer counterfactual support? Let me conclude by answering this question.

For simplicity's sake, and I think without any great loss of generality, focus on cases where the system or class of objects that is the subject of the nome comes into existence at some particular point in time by way of a "founding

13. Although for expository reasons I have not explicitly invoked the statistical-mechanical probabilities, I do think that these probabilities are relevant: in evaluating counterfactuals, the probabilistic weighting of possible worlds matters.

event" such as the evolution of the raven species or the formation of the planet earth. If the basing pattern is captured by the canonical form *All Fs in S are P* (where *S* is optional), then before the founding event there were no *F*s or at least no *S*, while after the founding event all *F*s in *S* are *P*.

The contingency of the basing pattern may be divided into two aspects: the contingency of the founding event, and the contingency of the basing pattern conditional on the founding event. The second of these is a kind of conditional contingency, then: to evaluate it, you restrict your attention to (nearby, physically possible) worlds in which the founding event occurs, and ask in what proportion of these the basing pattern holds as a result of the founding event. In the case of the raven basing pattern, for example, the founding event itself is highly contingent, but given the way that the event unfolded, the fact that all ravens have *P* is not accidental at all, because there is a causal mechanism that passes on the *P*-hood of the first ravens to later members of the species (albeit not infallibly).

A nome's founding event may be as contingent and accidental, as astoundingly improbable, as you like without compromising the nome's ability to offer counterfactual support (except to counterfactuals with early antecedents; section 5.2). The contingency of the basing pattern given the founding event has considerably more importance for its counterfactual implications. Let me consider several cases.

First, there may be, as in the case of the ravens, some mechanism that actively maintains the basing pattern, perhaps by repairing the *P*-hood of damaged *F*s or ensuring that new *F*s have *P*. In such cases, counterfactual support is robust.

Second, even if there is no active maintenance of the basing pattern, no Boydean homeostasis, the pattern may be highly resistant to disruption in a passive way. The geological structure of the earth provides an example: due to its large scale and other properties, it has what I earlier called physical inertia. This inertia has two related consequences: that the structure is not so contingent given the founding event—in worlds where the founding event occurs, the geological basing pattern is not accidental—and that the structure is highly likely to withstand the kinds of conservative adjustments made in the Bennett algorithm's stage II, ensuring that geological nomes offer counterfactual support that is as robust as that offered by the raven nome.

Third, consider a case where the basing pattern is highly contingent conditional on the founding event, as in a world where ravens have only a 50% chance of passing the black-producing mechanism *P* to their offspring, but in which by a fabulous fluke, all actual ravens nevertheless have the mechanism. In such a world, the counterpart of the raven nome will tend to support counterfactuals about actual ravens, given the physical inertia of *P* itself, but it will not support counterfactuals about non-actual ravens, such as ravens produced by non-actual matings.

The raven basing pattern can be made more contingent still. Imagine a world in which the blackness-producing mechanism *P* has little or no physical inertia—a world in which *P* is a fragile mix of properties that is apt to fall apart at any time—but in which, now by an even more fabulous fluke, all actual ravens nevertheless have *P*. Here, the raven nome's counterpart has effectively no power to sustain counterfactuals (and is for that reason surely unworthy of the epithet *nome*).

As you can see from these illustrations, the contingency of a nome's founding event makes little or no difference to its ability to support counterfactuals. The conditional contingency of its basing pattern given the founding event, by contrast, makes a great difference, or more exactly, it diagnoses the kind of persistence in the basing property *P* that makes a great difference. Here the proportionality thesis is more or less true; that is, the level of counterfactual support offered by a causal nome is roughly inversely proportional to the conditional contingency of its basing pattern: low conditional contingency (that is, near necessity conditional on the founding event) indicates the persistence that protects and propagates *P* during Bennett's stage II, guaranteeing strong counterfactual support, while high conditional contingency indicates the partial or total counterfactual evanescence of P.

It is no surprise, then, to find that the nomes of the special sciences tend to owe their contingency largely to the contingency of their founding event; conditional on the founding event, they are not that contingent after all. Special science nomes are in other words a kind of "accidentality frozen into nomic universality" (Schaffner 1993, 121).

7. Conclusion

Causal nomes support counterfactuals in a robust way, even when the basing patterns on which they depend, and thus the nomes themselves, are highly contingent or entirely accidental. There are systematic exceptions to the support offered, but none is a serious spoiler.

Why would anyone have thought otherwise? I suspect that it has not generally been appreciated just how exiguous a slice of the space of possible worlds is consulted in evaluating a typical counterfactual conditional. Causal nomes that depend on highly contingent basing generalizations hold in only a small region of the space of physically possible worlds, but this is more than enough modal territory to guarantee, for most scientific and everyday purposes, counterfactual support.

You might wonder what the interest of counterfactual support is, if it turns on such a narrow range of possibilia (a kind of revival of section 1's contingency anxiety that does not rest on the mistaken assumption that high contingency means thin counterfactual support). When you are investigating the fundamental physical laws, you can pride yourself on your contemplating the shared properties of the entire class of physically possible worlds. By comparison, an interest in the sort of counterfactual support that is the subject of this paper might seem modally rather provincial.

Two reasons to care, theoretical and practical. Theoretically, the high-level sciences can be understood, I have suggested, as inquiring into the nature and consequences of certain contingent structures: the genome of the HIV virus, the body of the raven, the brain of the human, the tectonic structure of the earth. A theoretical interest in such structures can hardly be impugned. But it is, of necessity, an interest in a very small part of the space of possible worlds.

Practically—I am thinking now of the role of hypothetical reasoning in decision-making—past events that are contingencies from the celestial perspective are to you virtual necessities. ("Who can turn skies back and begin again?") You must focus, then, on the small set of possibilities that can be reached given your and your world's history. Hence your narrow gaze.[14]

---

14. Thanks to Adam Elga, Justin Fisher, and Ken Waters; to audiences at the White Stallion Ranch, the University of Auckland, and the University of Minnesota; and to the *Imprint* referees.

References

Beatty, J. (1995). The evolutionary contingency thesis. In G. Wolters, J. G. Lennox, and P. McLaughlin (eds.), *Concepts, Theories, and Rationality in the Biological Sciences*, pp. 45–81. University of Pittsburgh Press, Pittsburgh.

Bennett, J. (2003). *A Philosophical Guide to Conditionals*. Oxford University Press, Oxford.

Bird, A. (2005). Unexpected a posteriori necessary laws of nature. *Australasian Journal of Philosophy* 83:533–548.

Boyd, R. (1999). Homeostasis, species, and higher taxa. In R. A. Wilson (ed.), *Species: New Interdisciplinary Essays*, pp. 141–185. MIT Press, Cambridge, MA.

Earman, J., C. Glymour, and S. Mitchell (eds.). (2003). *Ceteris Paribus Laws*. Kluwer, Dordrecht.

Ellis, B. (2001). *Scientific Essentialism*. Cambridge University Press, Cambridge.

Lange, M. (1995). Are there natural laws concerning particular biological species? *Journal of Philosophy* 92:430–451.

———. (2000). *Natural Laws in Scientific Practice*. Oxford University Press, Oxford.

———. (2004). The autonomy of functional biology: A reply to Rosenberg. *Biology and Philosophy* 19:93–109.

Lewis, D. (1973). *Counterfactuals*. Harvard University Press, Cambridge, MA.

———. (1979). Counterfactual dependence and time's arrow. *Noûs* 13:455–476.

Mitchell, S. (2002). Dimensions of scientific law. *Philosophy of Science* 67:242–265.

Morton, A. (1973). If I were a dry well-made match. *Dialogue* 12:322–324.

Rosenberg, A. (2006). *Darwinian Reductionism: Or, How to Stop Worrying and Love Molecular Biology*. University of Chicago Press, Chicago.

Schaffner, K. F. (1993). *Discovery and Explanation in Biology and Medicine*. University of Chicago Press, Chicago.

Strevens, M. (2009). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.

Waters, C. K. (1998). Causal regularities in the biological world of contingent distributions. *Biology and Philosophy* 13:5–36.

Weber, M. (2005). *Philosophy of Experimental Biology*. Cambridge University Press, Cambridge.

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