Ceteris Paribus Laws: Genericity and Natural Kinds

Bernhard Nickel · Harvard University

January 31, 2009

Abstract

Ceteris Paribus (cp-)laws may be said to hold only "other things equal," signaling that their truth is compatible with a range of exceptions. Several theorists have suggested that because of this feature, these generalizations are semantically defective. Against these concerns, I argue that cp-laws are often stated using a special kind of construction found in natural language, *generics*, which, though puzzling, is obviously semantically in order. I then argue that we can understand many otherwise puzzling features of cp-laws as an interaction between the semantics of generics and the structure of natural kinds in the special sciences.

1 Introduction

Practitioners of the special sciences, such as psychology, biology, or geography, articulate generalizations that seem to differ in important respects from the laws that physicists present. The former, but not the latter, tolerate exceptions, and we can mark this fact by saying that they are generalizations that hold other things equal, or *ceteris paribus*. These are the generalizations that are traditionally called cp-laws, though I'll argue below that this label is misleading.

CP-laws have played a prominent role in several debates in the philosophy of science, usually because there is unclarity, and hence disagreement, about what they say. Let me make this more precise. We can regiment statements of laws into the form *it is a law that p*. Initial discussions of laws of nature focused on analyzing the nomic operator *it is a law that*, trying to determine which claims, if true, are (or express) laws of nature.[1] To take one well-known example, one might want to know why *all uranium*

[1] See, e.g., Armstrong (1983); Hempel (1965); Lewis (1973b); Nagel (1979); Van Fraassen (1980, 1989).

spheres are less than 100,000km in diameter expresses a law, while the similar *all gold spheres are less than 100,000km in diameter* does not. In these examples, the content of the claim that is said to be a law is quite transparent.

In the case of cp-laws, that has been anything but transparent. When *p* in the schema *it is a law that p* is replaced by a sentence containing the locution *ceteris paribus*, it is not at all clear what proposition is said to be a law. This has shown up in the debates that cp-laws have played a role in: they all turn in one form or another on that unclarity. Two examples concern explanation and the nature of theories.

In the case of explanation, opponents of the DN-account have argued that cp-laws are unsuited to figure in deductively subsuming explanations, but that the special sciences nonetheless are capable of offering serious explanations of the phenomena they study.[2] In the case of theories, opponents of the view that theories are, or are properly modeled by, deductively closed axiomatic systems have argued that cp-laws are unsuited to enter into the deductive relations that this view would require of them.[3] In both cases, it is a problem about the deductive relations cp-laws enter into, and thus about the proposition said to be a law, that animates the debate.

For this reason, much of the debate about cp-laws has focused on the question how we should interpret the relevant generalization. The main burden of this paper is to make a two-fold contribution to this literature, one methodological, the other more directly substantive. I will argue that in stating the relevant generalizations, practitioners of the special sciences make use of a linguistic resource familiar from ordinary language, specifically, genericity. We can thus make progress on understanding cp-laws by appealing to the semantics of natural language at several crucial junctures. That is the methodological aim. More substantively, I will focus on one kind of genericity we can use in stating *some* cp-laws, what I'll call *characterizing sentences*, and give truth-conditions for these sentences. On my proposal, the cp-laws we state using characterizing sentences are very closely connected to the natural kinds of the discipline in which they are articulated. I'll argue that on this proposal, we can account for several important aspects of these cp-laws, some familiar, some new.

I'll begin by discussing the current focus of the literature on the so-called triviality problem and explain why I think that this focus is mistaken (§2). I then turn to the debate between Pietroski and Rey (1995) on the one hand and Woodward (2002) on the other. Though this debate is officially concerned with the issue of triviality, it will

[2]See, e.g., Dray (1957) and Rosenberg (2001a,b).

[3]See Giere (1988a,b, 1999) and Cartwright (1983, 1989, 1995, 1999).

serve as a useful jumping off point for my own account (§3). §4 is devoted to making the connection between cp-laws and natural language more precise, as well as to focus on the more specific target of analysis for the rest of the paper. §5 then introduces and motivates the view of natural kinds I will appeal to and explains the relationship between characterizing sentences and natural kinds.

Throughout, I won't take a stand on whether we should take cp-laws to be genuine laws. My concern is exclusively with getting clear on the broadly semantic issues. To mark this fact in my terminology, I'll discuss what I'll call *cp-generalizations*. A cp-generalization is a bit of language that can take the place of *p* in the regimentation *it is a law that p*, and when it does, the instance of the schema expresses a purported cp-law.

2 CP-Generalizations and Triviality

The problem usually discussed under the heading of triviality in connection with cp-generalizations arises because it is very hard to say which exceptions would refute a given cp-generalization and which would not. Let's call the former *genuine* exceptions, the latter *merely apparent*. A natural way of trying to say what the merely apparent exceptions are appends a clause headed by 'unless' to a universal generalization derived from the cp-generalization. An example will help to make the point. We might begin with the cp-generalization (1).

(1) *Ceteris paribus*, all ravens are black.

(1) does not have the same force as the corresponding universal generalization *all ravens are black*, since (1) is true, while the simple universal is false, as witnessed for example by albinos. A more promising attempt at saying what (1) amounts to without helping ourselves to the *ceteris paribus* locution is (2).

(2) All ravens are black, unless they are albinos.

However, (2) does not capture the force of (1) either, since ravens that are non-black because of the environmental conditions they experience, rather than the genetic endowment they are born with, would falsify (2) without intuitively leading us to reject (1). And it seems as if, for any way of adding more qualifications to the unless-clause, we can come up with more merely apparent exceptions we have not yet captured. Let's call such an unless-clause *open-ended*, and let's call the cp-generalization that gives rise to such an unless-clause *open-ended*, as well.[4] Of course, there are some ways of spelling out which exceptions are merely apparent that does not result in an open-ended unless-clause. We could say that ravens are black unless they are abnormal. But if we do not help ourselves to these expressions, the list of merely apparent exceptions is open-ended.

Many theorists argue that the open-endedness of cp-generalizations threatens them with triviality.[5] A generalization is non-trivial if there could be genuine exceptions, i.e., if there are circumstances that would falsify it. But the list of merely apparent exceptions to a cp-generalization is open-ended, so anything could be covered by that list. In that case, there couldn't be any genuine exceptions, making the generalization trivial. That is the *triviality worry*. Applied to my example, it says that (1) amounts to no more than *all ravens are black, unless they aren't.*

We can make the case for the triviality worry even stronger than it is usually presented in the literature. Given that we accept (1), we are unable to imagine genuine exceptions to the generalization. Every non-black raven intuitively strikes us as a merely apparent exception. This is different from the situation in which I am certain of a contingent universal generalization. In that case also, I am unwilling to accept anything that looks like a counterexample as really being one. Since I am certain that the universal generalization is true, any apparent counterexample must be merely apparent. But in this case, I can easily imagine what an isolated counter-example would be, whereas for a cp-generalization I accept, I am unable to do that.

If we put together the open-endedness of cp-generalizations with our inability to imagine genuine counterexamples to ones we accept, we can see why there might be a presumption of triviality. The fact that we cannot imagine genuine counterexamples makes them seem like conceptual or necessary truths of some sort. But the subject matter, together with the deep opacity of what we are saying when we assert a cp-generalization, make it implausible that the truths are substantive. They must be trivial.

The triviality worry is by far the most important point of debate in the literature on cp-generalizations. A subsidiary one focuses on epistemic considerations, concerning the claim that cp-generalizations cannot be empirically confirmed or disconfirmed.[6]

[4]The problem concerns the *potential* exceptions to the generalization. Even if at some point only black ravens exist, (1) still poses the same problem, since it *would not* be falsified by certain non-black ravens.

[5]Pietroski and Rey (1995, 87) give a prominent statement. Woodward, who argues that there is a significant problem with cp-generalizations, cites this passage as stating the problem (Woodward, 2002, 308). Other authors who also take triviality to be the main problem include: Earman and Roberts (1999), Earman et al. (2002), Fodor (1991), Mott (1992), Schiffer (1991), Schurz (2001, 2002), and Silverberg (1996).

[6]See for example Earman et al. (2002, 293) and Schurz (2002, 360-2).

Any observed exception, no matter what, could potentially be counted among the merely apparent ones, because there are no significant constraints on them. That makes an "honest test" impossible. But this is really just the triviality worry again. Both turn on the idea that there are no constraints on the merely apparent exceptions. For that reason, the concern about confirmation stands and falls with the triviality worry, and I won't discuss it separately here.

However, pursuing the triviality worry is not the right way to get at what is really puzzling about cp-generalizations. We can see this by observing that the worry is easily met, but without putting us in a position to explain their open-endedness or the inconceivability of an isolated, genuine exception. For notice that cp-generalizations do not, in general, behave at all like trivial claims.[7] If they were trivial, we would predict all of the following to be true.

(3) a. (*Ceteris paribus*), all ravens are black.

b. (*Ceteris paribus*), all ravens are polka-dotted.

c. (*Ceteris paribus*), all ravens are white.

d. (*Ceteris paribus*), all ravens have the color scheme of Crimson Rosellas.

Uncontroversially, (3a) is a far better thing to say in just about any context than any of (3b)-(3d). This difference in felicity does not by itself show that there is a difference in truth-value, since the unacceptable sentences might be unacceptable for reasons compatible with their being true. However, if (3b) really were true, denying it would yield a falsehood. So if (3b) is true, (4) must be false.

(4) It is not the case that (*ceteris paribus*), all ravens are polka-dotted.

But (4) is obviously acceptable, and this speaks strongly in favor of its truth, which in turn requires (3b) to be false.[8] Analogous arguments show (3c) and (3d) to be false, as well.

Here is another respect in which cp-generalizations do not behave like trivialities. We can quite easily imagine states of the world that are ruled out by the truth of a cp-generalization like (1), such as a world in which ravens evolved to have the color

[7] I say "in general," because some generalizations with a similar surface form to (1) obviously are trivial, such as 'ravens are ravens.' From now on, I will systematically disregard such claims when I talk about the non-triviality of cp-generalizations

[8] Observe that in (4), the negation has wide scope over the whole sentence. That is, (4) denies something that, according to the theory we're considering is, trivially true. (4) does not just say that ravens have the property of being not-polka-dotted, which is true by anybody's lights.

scheme of Crimson Rosellas. The generalization is therefore contingent, not trivial. Note that this is compatible with our earlier observation that we cannot imagine a single individual that would be a genuine exception to (1). It just turns out that we *can* imagine worlds that incorporate large-scale changes and with respect to which (1) is false. I'll mention one other reason to think that the focus on the triviality worry is misplaced in section 4.1, below.

We have so far replaced the triviality worry with two other striking and *prima facie* puzzling features: the fact that cp-generalizations are open-ended, and the fact that while we cannot imagine individual genuine exceptions to a true cp-generalization, we can easily imagine scenarios that differ widely from the actual world and in which that generalization is false. A theory of cp-generalizations should yield an account of these features.

3 Pietroski & Rey vs. Woodward

Pietroski and Rey explicitly address the triviality worry in their (1995), but much of what they say is a substantive theory of cp-laws that is of interest quite independently of whether we want to use it to respond to the triviality worry. In discussing their view, I have two aims, one positive, one negative. On the positive side, I want to highlight a basic motivating thought that I agree with, as well: cp-laws reflect the needs of theorizing in sciences that investigate complex phenomena. I'll also show that their view goes some way towards illuminating the striking features I just mentioned, since it offers a ready account of the open-endedness of cp-generalizations. However, I'll argue on the negative side that their view faces some important shortcomings. First, it does not offer an account of why we can only imagine large-scale changes that falsify a true cp-generalization. And second, on the account of Pietroski and Rey (1995), we predict that various cp-generalizations are logically compatible, whereas such compatibility simply isn't attested. These problems will set the stage for the rest of the paper.

3.1 The View

Pietroski and Rey introduce their view thus.

> [O]ur view is that scientists state cp-laws in an attempt to focus on particular factors (e.g., natural selection) and thereby 'carve' complex phenomena

(e.g. the evolution of populations) in a theoretically important way.[9]

This is an idea that I want to pursue, as well. A theoretically important carving of the phenomena a scientific discipline investigates is at the heart of the practice of using cp-laws, and is therefore also the key to understanding that practice. However, this programmatic idea needs to be spelled out, and here is the next step that Pietroski and Rey take.

> Our own view is motivated by the following general consideration: the emergence of any theoretically interesting science requires considerable abstraction and idealization. The actual world is too complex to study all at once, so one proceeds by ignoring some aspects of the world in order to understand others. We idealize away from friction, electric charge, and nuclear forces, for example, when we seek to understand the effect of gravity on the motion of bodies. However, such abstraction guarantees a loss of descriptive adequacy in any generalization we lay down, since actual bodies are always affected by, e.g., friction, at least a little.[10]

However, this way of spelling out the connection between cp-generalizations and the needs of theorizing leads to a fundamental problem. To that end, I'll introduce Pietroski and Rey's specific proposal and then explain how it naturally grows out of the general consideration I just cited.

In the first instance, Pietroski and Rey do not wish to give truth-conditions for statements of cp-laws. Rather, they want to state a condition that, if satisfied by a cp-law, ensures that it is non-trivial. However, the two tasks are not really separate, since a cp-law is non-trivial just in case its truth-conditions cannot be satisfied trivially. And that means that, in stating a non-vacuity condition, Pietroski and Rey state a necessary condition on the truth of a cp-law. Indeed, they go on to make some suggestions about the sufficient conditions, as well.

As the quotes make clear, Pietroski and Rey couch a lot of their discussion in terms of examples from simple physics, such as ideal gases. I'll make use of a similar example, that of pendula, but the points generalize. The relevant cp-law about pendula is that the period of a pendulum depends on the length of the string and the angle of displacement according to a simple formula, captured in this generalization.

(5) (*Ceteris paribus*), all pendula have a period determined by $T = 2\pi\sqrt{\frac{l}{g}}$.

[9] Pietroski and Rey (1995, 92).
[10] Pietroski and Rey (1995, 89).

There are, of course, conditions under which (5) fails to accurately describe the period of a given pendulum: the bob might be subject to friction-forces, or it might, if made of iron, be subject to a magnet placed in its vicinity to either increase or decrease its period. Nonetheless, one might think, the generalization (5) gets at something important, and it certainly seems acceptable.

At any rate, whatever qualms one may have about calling (5) a *law*, it does not seem to be a triviality. Here, according to Pietroski and Rey, is the relevant feature that accounts for its non-vacuity, and the non-vacuity of other acceptable cp-generalizations. Certain cases covered by the law, i.e., certain pendula, conform to that law. Other cases covered by the law, i.e., other pendula, do not. In the latter cases, however, we can explain why any given failure to conform to (5) is a failure by citing one or more factors that do independent explanatory work. To put this in terms of the examples I've already mentioned, if the period of a pendulum fails to conform to (5) because the bob is made of iron and subject to a magnetic field, then we've explained (or at least can explain) why that particular pendulum fails to conform to (5) by citing the magnetic field. Appeals to that magnetic field, at least potentially, do independent explanatory work for us, such as explaining why a compass-needle in the vicinity turns in a particular direction.

The general idea, then, is this. A cp-law is non-vacuous if the following is a necessary condition for its truth: all exceptions to the law can be explained away citing independent causal factors.[11] And as Pietroski and Rey suggest, this can also serve as a sufficient condition.[12] Thus, the truth-conditions are exhausted by the non-vacuity condition. To see how this condition works to exclude a putative law whose acceptability quite clearly has been bought at the price of trivialization, consider (6) about ESP.

(6) (*Ceteris paribus*), on all occasions on which Jones tries to divine the future, she is successful.

Faced with an apparent exception, i.e., a failure to see the future coming, the proponent of (6) might cite a factor to explain the failure away that itself does no independent work, such as an appeal to psychic static. If she countenances any such kinds of interferences she countenances, (6) is trivialized. Then again, she might cite factors that do independent explanatory work, such as sunspots or migraines. In that case, the putative cp-generalization may not be trivialized, but simply false if it turns out that sunspots or migraines don't actually explain the failure.

[11] See Pietroski and Rey (1995, 92) for their more elaborate statement of the view.
[12] See Pietroski and Rey (1995, 92).

Pietroski and Rey do a lot to give an account of what independent explanatory work amounts to, but I will simply grant them the notion, since the points I will make now do not rely on any controversial way of interpreting that idea.

3.2 Interference, Causation, and Open-Endedness

I'll begin by showing that implicit in the view of Pietroski and Rey is a very interesting and promising explanation for why cp-generalizations are open-ended in the sense of §2. To make this more obvious, it will help to think about the proposal in slightly different but equivalent terms.

I suggest that we interpret their view as saying that a cp-generalization like (5) describes the outcome of the operation of a (possibly singleton) set of causal factors, though these causal factors by no means exhaust all of the causal factors that could possibly affect an object of the kind the cp-law is about, such as a pendulum. For example, (5) describes the working of gravity on a pendulum setup, though it ignores some of the causal factors that could potentially operate on such a setup, such as the presence of friction-inducing media or magnetic fields.

This presentation of Pietroski and Rey's account is equivalent to their own official one. On theirs, a cp-generalization is true iff all exceptions to the generalization can be explained away by citing only factors that do independent explanatory work. On mine, a cp-generalization is true iff there is a (possibly singleton) set of causal factors that bring about the state of affairs described by the generalization in at least some of the cases covered by the generalization. To establish the equivalence, suppose first that a cp-generalization satisfies the conditions set out by Pietroski and Rey. Then the cp-generalization describes a kind of phenomenon that is subject to causal influences—if it did not, deviations form the generalization couldn't be explained away by citing independent causal factors. Hence, there is also a set of causal factors that brings about the state of affairs described by the generalization, and thus my alternative truth-conditions are satisfied. Now the other direction: if there is such a set of causal factors in virtue of which my alternative truth-conditions are satisfied, then any deviation from the generalization must be due to the operation of some other causal factors, and since any causal factor can always do independent explanatory work, the original truth-conditions are satisfied, as well.

More abstractly, the shift concerns what we think of as fundamental. Pietroski and Rey couch their discussion of non-vacuity in terms of a feature of the exceptions to a putative cp-generalization. Instead, we should think of features of the cases that conform to it. The argument of the previous paragraph shows that Pietroski and Rey's way of thinking about their theory and mine come to the same thing. However, my way of thinking about it has the benefit of making the connection to their motivating claims about idealizations easier to see, and of showing that they immediately present an explanation of the open-endedness of cp-generalizations.

In very many cases, an idealization of a situation we're investigating is a simplified version of it.[13] Most importantly, the idealization will contain far fewer causal factors than the original situation. In the idealization, these factors are the only ones involved in bringing about an effect. On Pietroski and Rey's account, a cp-generalization about a situation reflects the operation of such a restricted range of causal factors in those cases where they are the only significant ones. We might put the connection between cp-generalizations and idealizations like this. A cp-generalization about a kind of thing—pendula, for instance—is true iff there is an idealization that takes into account some of the factors potentially acting on things of that kind, and in the idealization those factors bring about the state of affairs described by the generalization, and in some members of the kind, these causal factors exhaust the causally relevant ones.

Recall now that cp-generalizations are open-ended if the list of merely apparent exceptions is heterogeneous and open-ended. If Pietroski and Rey are right, then this is unsurprising. A merely apparent exception, i.e., one that does not serve to refute the cp-generalization, arises because a causal factor operates on a member of the kind mentioned in the law that goes beyond those countenanced in the underlying idealization and thus forces that object to behave in a way that doesn't conform to the generalization. But the form that such causal influences can take are legion, and they need not have anything in common except that they make the object deviate from the course predicted by the underlying idealization. Thus, thinking of cp-generalizations as about causal processes that issue in one possible behavior among several offers a natural and simple explanation of the open-endedness of these generalizations.

3.3 Woodward's Charges

However, given this way of thinking about the truth-conditions of cp-generalizations, we can also see a potential problem. In order for a cp-generalization to be true, it is sufficient that *there is* a suitable set of causal factors. CP-generalizations are there-

[13]Obviously, this is not supposed to do justice to the wide range of things one might call an idealization in the sciences, let alone the closely related notion of a model. But I trust that the description in the text is true of an important subclass of idealizations. For at least some other things that are reasonably counted among idealizations or models, see Mäki (2002); Morrison and Morgan (1999)

fore essentially existential claims. And in general, existential claims are compatible with other existential claims about the same things. Hence, we should expect that, for any such suitable set of causal factors, there is a corresponding true cp-generalization. However, that prediction is simply not borne out. Return to the example of pendula. In *some* pendula, the relevant causal and explanatory factors are exhausted by the length of the string and the force of gravity, and in those cases, the period is given by the formula in (5), which is why (5) is true. However, in some *other* pendula, the relevant causal and explanatory factors include those two and others besides, such as the presence of a magnetic field with a particular strength and orientation. In those cases, a different formula applies, call it T'. We should thus expect that another cp-generalization of the form (7) is also true.

(7) (*Ceteris paribus*), all pendula have a period determined by T'.

However, the intuitive judgment here is two-fold: for one, (7) is false. But more importantly for my purposes, (5) and (7) are incompatible. At most one of them can be true. In this case at least, we would have to restate (7) by qualifying which pendula the generalization applies to, as in (8).

(8) (*Ceteris paribus*), all pendula with iron bobs in such-and-such magnetic fields have a period determined by T'.

The pattern of incompatible cp-generalizations that can be saved by appropriate qualification is quite general, as (9) and (10) illustrate.

(9)
a. (*Ceteris paribus*), all ravens are black. (true)
b. (*Ceteris paribus*), all ravens are white. (false)
c. (*Ceteris paribus*), all albino ravens are white. (true)

(10)
a. (*Ceteris paribus*), if the price of a good falls, demand for that good increases. (true)
b. (*Ceteris paribus*), if the price of a good falls, demand for that good falls. (false)
c. (*Ceteris paribus*), if the price of a good falls and the price of a substitute good falls even more, demand for that (initial) good falls. (true)

I think that this observation is also what fundamentally makes Woodward's influential example about charged particles work. He considers (11).

(11) (*Ceteris paribus*), all charged objects accelerate at $10m/s^2$.

And he goes on to say:

> For every charged object, there is an additional condition K (having to do with the application of an electromagnetic field of appropriate strength to the object) that in conjunction with the object's being charged is nomically sufficient for its accelerating at $10m/s^2$. [...] for those charged objects that do not accelerate at $10m/s^2$, there is always an explanation that appeals to some other factor K′ for why this is so—K′ will presumably have to do with the fact that the object in question has been subjected to an electromagnetic field (or some other force) of the wrong magnitude to produce this acceleration. In addition, since classical electromagnetism is a powerful, non ad hoc theory, K′ will figure in the explanation of many other facts. [...] Even more alarmingly, parallel reasoning can be used to show that "All charged particles accelerate at $n\ m/s^2$" is a *ceteris paribus* law for all other values of n.[14]

I especially want to draw attention to the end of this quotation. It is the fact that, as far as Pietroski and Rey's account enables us to see, each of the generalizations of the form (12) is true that is the real concern.

(12) (*Ceteris paribus*), all charged particles accelerate at $n \cdot m/s^2$.

That is to say, what I think is really problematic about Pietroski and Rey's account is the fact that it predicts that all sentences of the form (12) are logically consistent, whereas that is clearly not the case. This way of diagnosing the problem with the account has some significant advantages over Woodward's original presentation. Woodward wants to simply conclude that the sentences of the form (12) cannot all be cp-laws. But this conclusion doesn't allow us to distinguish two potential sources of the difficulty. The problem could either be that the generalization said to be a law (*ceteris paribus or otherwise*) is false, or the problem could be that the generalization, though true, fails to be a law. Given that both of these options are live, it's not clear where to lay the blame for the failure of any instance of (12) to be a cp-law.

By contrast, if the problem is about the consistency of various statements, we know where to lay the blame. The generalizations said to be laws cannot all be true together,

[14] Woodward (2002, 310).

so the problem is quite independent of any issues regarding the nomic operator. This way of diagnosing the problem also casts a different light on a discussion by Pietroski and Rey in a similar context. They consider the concern that, by their lights, it looks as if any true singular causal claim entails a corresponding cp-law.

> Every singular causal claim of the form A caused B—e.g., Alice's favorite event caused Betty's most hated event—is a candidate for an interferable cp-law, since the quantifier might range over all the conditions that prevent As from bringing about Bs in all other cases. So far as anything we have said, there *might* be a cp-law 'cp(A⇒B)' to the effect that, cp, Alice's favorites cause Betty's hateds; it is just that *cetera* have been *paria* only *once* in the history of the world, and, moreover, are not easily made so. However, we are not committed to regarding 'cp(A⇒B)' as a bona fide law, but only to claiming that, given that the singular claim is true, the corresponding CP law cannot be criticized for being *vacuous*.[15]

I take it that similarly, Pietroski and Rey would say that the problem with Woodward's schema (12) is that these claims, though acceptable generalizations, fail to be laws, i.e., that it is the job of the nomic operator to explain why the instances of (12) are unacceptable as cp-laws. But if I am right in saying that the problem is at bottom one about consistency, not the nomic status of the generalizations, then this reply is insufficient. We're missing something fundamental about the semantics of the generalizations said to be laws by being unable to capture obvious inconsistencies among them.[16]

The upshot of the discussion is therefore mixed. On the one hand, we've seen that thinking about cp-generalizations in terms of causal processes or mechanisms is very appealing, because it allows us to give a convincing account of the open-endedness of cp-generalizations—plausibly one of their features that animates a lot of the debate currently conducted as a debate about triviality. On the other hand, we don't have the right way of spelling out that connection, because we cannot capture consistencies and inconsistencies. And finally, we don't have an account of why we can imagine large-scale changes that falsify a true cp-generalization, though we are unable to imagine genuine individual exceptions. In the case of Pietroski and Rey, the problem takes the form of not being able to explain why even large-scale changes in the way the world

[15]Pietroski and Rey (1995, 98-9).

[16]Incidentally, in the linguistic literature, a very similar account to that of Pietroski and Rey has been developed by Cavedon and Glasbey (1994), drawing on work in Barwise (1993) and Barwise and Seligman (1994). Their account suffers from identical drawbacks.

is should falsify a cp-generalization. As I've argued, on their view, cp-generalizations are in the first instance claims about the existence of certain idealizations, and whether such idealizations exist is completely independent of what the world is like.

Notice that especially this last problem is specific to the way Pietroski and Rey spell out their idea that cp-generalizations reflect theoretically important carvings, rather than that general idea itself. What counts as a theoretically important carving may well depend on what the world is like. What we need is some alternative way of spelling out the general idea that doesn't run into the problems I've presented here. That is the task I turn to now.

4 Targeting the Analysis

I have so far gone along with the practice of speaking relatively indiscriminately about cp-laws and cp-generalizations as if these formed a unified class amenable to a unified treatment. However, I think that Woodward is clearly right when he says that there is a

> great diversity and heterogeneity [to] the generalizations that philosophers propose to analyze in terms of the category *ceteris paribus* laws.[17]

That means that any theory of generalizations in this area needs to have a more sharply defined range of phenomena it wants to account for. A natural first thought is that what is special about cp-generalizations is the presence of a certain locution in their statement, to wit, *ceteris paribus*. However, as Schiffer and Woodward both point out, that phrase is hardly ever employed outside of economics.[18] We thus need some other means of singling out the target phenomena and carve them up into unified classes.

4.1 CP-Laws and Generics

My suggestion is that we turn to natural language to do that job. Let me introduce a class of linguistic phenomena that linguists and philosophers of language study under the heading of *generics*. The term derives from the intuition that very often, we speak about kinds—genera—and say something about them. In some cases, we seem to speak of a kind as a whole, as when we say that *quartz is widespread* or *dodos are extinct*. However, some generics behave exactly like statements of cp-laws in the sciences, such as *ravens are black* or *turtles are long-lived*. These sentences are compatible with what

[17] Woodward (2002, 305). Hall (2005) makes a similar observation.

[18] See Schiffer (1991, 10) and Woodward (2002, 305).

would be counter-examples to the corresponding universal generalizations, just as cp-generalizations are. They also have what linguists often call a law-like flavor, which is to say that their semantics have modal import. We can bring this out by observing that their truth is compatible with at least some situations in which the corresponding existential generalization would be false. *Ravens are black*, for example, is true even in a situation in which all ravens have been painted.[19] That means that the truth-value of such a generic doesn't just depend on the state of the world of evaluation at the time of evaluation, but on what is true at other worlds and/or times, i.e., on modal facts. They also exhibit the same open-endedness and contrast between large- and small-scale changes as the cp-generalizations paradigmatically discussed in the literature.

To a good first approximation, I thus want to claim that such generics are used to state very many cp-laws. More specifically, if we continue to regiment the statement of cp-laws into the form *it is a law that (ceteris paribus) p*, many cp-laws can be, and are most naturally, stated by removing the *ceteris paribus* locution along with any explicit quantifiers and replacing the schematic variable *p* with a generic sentence. (13) illustrates this connection by way of some examples.

(13) a. It is a law that (*ceteris paribus*), all slow rivers meander.
b. It is a law that (*ceteris paribus*), all grass is green.
c. It is a law that (*ceteris paribus*), if the price of a good falls, demand always rises.
d. It is a law that (*ceteris paribus*), iron bars always expand when heated.

Removing (*ceteris paribus)* and the explicit quantifiers leaves us with generalizations that have the same features.

(14) a. Slow rivers meander.
b. Grass is green.
c. If the price of a good falls, demands rises.
d. Iron bars expand when heated.

Just as importantly given the concerns I've raised in §3.3, we see the same pattern of incompatibilities, as (15) and (16) illustrate.

[19]Perhaps one would like to respond that, *in the relevant sense*, the ravens are still black when they're painted white, taking inspiration from Austin (1975) and, following him, Travis (1985, 2000). However, the point still holds: *lions have four legs* can be true even in a situation in which all lions have lost one of their legs in accidents.

(15) a. Ravens are black. (true)

b. Ravens are white. (false)

c. Albino ravens are white. (true)

(16) a. If the price of a good falls, demand for that good increases. (true)

b. If the price of a good falls, demand for that good falls. (false)

c. If the price of a good falls and the price of a substitute good falls even more, demand for that (initial) good falls. (true)

So if we want to understand cp-generalizations, we need to focus not on the locution *ceteris paribus*, but on the rest of that statement.

This connection between cp-generalizations and generics in natural language is one more reason to think that the triviality worry is not the right way to pursue a deeper understanding of cp-laws. In fact, recognizing the connection between cp-laws and generics allows us to undermine the most interesting motivation for that worry.

In general, so that motivation goes, we need to show that any technical jargon that we employ is coherent and contentful. If what is distinctive of cp-laws is the use of the locution *ceteris paribus*, then this bit of technical jargon must be shown to be appropriately constrained. Doing so can take two forms. We can give a reductive definition of the technical jargon at issue, or we can simply exhibit a well-confirmed, predictively and explanatorily powerful theory in which that jargon plays a significant role. Because this second option is available, nobody is concerned about the meaningfulness of the technical vocabulary of physics. But cp-generalizations usually appear in arenas of inquiry that cannot boast of predictive success, and whose claims to explanatory success are controversial. So if we have to exhibit the bona fides of the *ceteris paribus* locution, we have to do so by way of a reductive definition. Given that this has proved very difficult indeed—as shown by Earman and Roberts (1999); Earman et al. (2002)—we should worry that it cannot be done, and hence that generalizations that essentially depend on this locution are theoretically defective.

However, this way of motivating the triviality worry goes wrong at the first step when we suppose that the distinctive feature of cp-laws is the use of the locution *ceteris paribus* in their statement. What is distinctive, instead, is the use of generics in their statement. And generics are not some technical vocabulary whose bona fides could be in doubt. Thus, the triviality worry does not get going this way, either. It also explains why reductive analyses have failed. It is simply an instance of the general fact that reductive analyses of expressions of natural language are almost never available. Thus, the correct moral to draw from the continued failure of such reductive analyses is that we're dealing with theoretically interesting natural language, not theoretically defective technical jargon.

4.2 Characterizing Sentences

The class of generics is very broad, and it is simply impossible to do justice to that class as a whole. For that reason, I'll focus my inquiry more narrowly on a class of sentences that I'll call *characterizing sentences*.[20] In the first instance, I'll focus on generics that express a non-strict generalization of some kind and that have the modal component I pointed to.[21] Among these, there are at least two broad classes. One consists of generalizations over objects, and which require for their truth that a suitable subset of the members of the kind at issue conform to the generalization, such as *ravens are black*. These are the sentences I want to focus on. The other class is illustrated by the examples in (17).

(17) a. Mary smokes.

b. Dogs bark.

(17a), for example, says more than that Mary has, on at least one occasion, smoked, but it also does not say that she smokes at all times. Rather, we can get at what (17a) is after by interpreting it as a generalization over events that involve Mary, claiming that some appropriate subset of these events are ones in which Mary smokes.[22] In contrast to characterizing sentences, I'll call sentences like (17a) that we can interpret as a generalization over events *habituals*. (17b) mixes generalizations over objects and events and is therefore non-strict twice over. It is both a characterizing sentence and a habitual: a characterizing sentence because it does not apply to all dogs but only dogs that have normal vocal tracts, and a habitual because it does not say that all events involving these dogs are barking events. I'll only be concerned with characterizing sentences in this paper. I want to emphasize that this is a non-trivial restriction, since

[20] Here, I follow Krifka et al. (1995), which has done a lot to standardize the terminology in the field.

[21] I thus set aside sentences that predicate a property of a kind as a whole, such as *quartz is widespread*, and merely existential sentences such as *ravens are sitting on the wires outside my house*.

[22] This is an idea that traces back to Davidson's discussion of action sentences in Davidson (1980), and from there has become a central plank in what has become known as *neo-Davidsonian* semantics, the view that all sentences, not just action sentences, contain quantification over events in their logical form. One of the early applications of this idea to ascriptions of dispositions is Lewis (1973a). For a more recent elaboration, see Fara (2001, 2005).

it excludes ascriptions of dispositions (ascriptions of properties such as being fragile), and it excludes many statements involving explicit mention of the verb *cause* because these ascriptions are usually couched in terms of habituals.

The kind of proposition that is expressed by a sentence such as *ravens are black* can be expressed in several different ways, aside from the bare plural I just used. One could also use a singular definite or singular indefinite article, i.e., *the* and *a*, as in *the raven is black* and *a raven is black*. Nonetheless, I'll focus on sentences with bare plural subjects.[23] In some cases, a sentence is ambiguous and can express two different generic propositions, such as the famous *typhoons arise in this part of the Pacific*, which can be used to express either that this part of the Pacific is where typhoons arise—they only exceptionally arise elsewhere—or that this part of the Pacific is regularly subjected to typhoons, although typhoons also arise elsewhere.[24] The first of these propositions is the kind I'm after. Since none of the examples I'll discuss exhibit this ambiguity, I'll ignore it from hereon.[25]

4.3 Semantics for Characterizing Sentences

The strategy now is to present enough of empirically motivated semantics for characterizing sentences to allow me account for the puzzling features of cp-generalizations I've collected so far, as well as some further facts that will come to light in the subsequent discussion. The semantics I'll present will, at least initially, be formulated in terms of a primitive notion of normality. I'll then try to show how the special sciences naturally provide the resources to relate that notion to the demands theories in these disciplines face, especially demands that arise from the complexity of the phenomena they investigate. In this way, my view represents an alternative development of the guiding principle that cp-generalizations reflect the needs of theorizing about complex

[23]One reason for this preference is that plurality seems to be a basic ingredient in genericity, and apparently singular subjects are nonetheless interpreted as having some features of plurality. We can see this by noting that certain predicates that can usually only be applied to pluralities can appear in generic sentences with singular subjects. For example, ordinarily we can predicate *form a circle* or *surround* only of pluralities, as the contrast between *the children formed a circle* and **Mary formed a circle* shows. However, many speaker find *the buffalo forms a protective circle* completely acceptable.

[24]The example is due to Milsark (1974).

[25]Woodward suggests that sometimes, we can use such bare plural sentences to express nothing more than statistical generalizations, citing such as examples as *drivers in England drive on the left* (Woodward, 2002, 311). Officially, I can be agnostic on whether it's possible to use bare plurals to express such a pure regularity, saying only that bare plurals, when so used, fall outside the purview of my theory. But as a matter of fact, I disagree with Woodward's characterization of the data. I think what's really at issue is whether this kind of generalization is, or can be, part of a systematic theory and hence deserves the title of *law*, as Woodward seems to acknowledge at (Woodward, 2002, 311).

phenomena.[26]

Following a lot of work in this area, I will assume that the intuitive sense that characterizing sentences express generalizations over the members of the kind mentioned is basically accurate. The question is what kind of generalization characterizing sentences express.[27] A first pass at an answer is this.

[First Pass Truth-Conditions]

As are F is true iff all normal As are F.

The point of this first step is to take into account the fact that characterizing sentences tolerate exceptions: tolerable exceptions are not normal, and hence do not falsify the characterizing sentence.[28]

The next pass at the truth-conditions can be motivated by considering the pair of sentences in (17).

(17) a. Chickens lay eggs.

[26]In providing truth-conditional semantics, I'll take on commitments large and small that I can only acknowledge here. First is the commitment that giving truth-conditions is the right way to give the semantics of natural language in general, or generics in particular. Opponents have cited generics in arguing for alternative semantic frameworks. In the literature on linguistics and philosophy of language, see Chomsky (1975), Leslie (2007, 2008), Schubert and Pelletier (1989), and Veltman (1996). In the literature on philosophy of science, some theorists have suggested that we shouldn't interpret cp-law statements as expressing propositions, but in some other way, perhaps as inference rules, as in Lange (2000). For another alternative, see Glymour (2002). One of the reasons theorists often give for rejecting ordinary truth-conditional semantics is that while it's possible to give somewhat plausible semantics for simple cases, it is extremely hard to show how they can be extended to more complex examples in a theoretically motivated and compositional way. This will be true of the semantics I present here, as well. This is, in the first place, a concern about the viability of a research program, the program of fitting semantics for characterizing sentences into the overall truth-conditional framework. The right way to pursue this point is to see whether the research program yields interesting results in some core cases in order to determine whether it's worthwhile to try to develop it further.

[27]For some reasons to prefer such a quantificational approach to the truth-conditions of characterizing sentences, see Cohen (1999), Krifka et al. (1995). I have also argued for it in my Nickel (2008a,b,c). One reason to adopt the quantificational approach is that we can account for the availability of two readings of *typhoons arise in this part of the Pacific*.

[28]One might worry about potential counter-examples at this stage.

(A) a. Dutchmen are good sailors.

b. Mosquitos carry plasmodia (the organisms that cause malaria in humans).

(Aa) is not well paraphrased as saying that all normal Dutchmen are good sailors, and (Ab) doesn't seem to be as strong as the claim that all normal mosquitos carry plasmodia. After all, (Ab) is true in the actual world, even though only a minority of mosquitos actually carries the organisms, and that minority doesn't seem particularly normal. Suffice it to say that there are responses available. Cohen (1999) and Krifka et al. (1995) suggest that (Aa) involves a different reading of the bare plural. I also offer a treatment of (Aa) that eschews such an ambiguity in my (2008a). And the acceptability of (Ab) may well be due to the fact that the sentence is ambiguous. The acceptable reading might be the one that is a characterizing sentence about plasmodia, to the effect that all normal plasmodia are carried by mosquitos (as part of their normal life-cycle). This isn't the end of the debate, obviously.

b. Chickens are hens.

If we interpreted (17a) as saying that all normal chickens lay eggs, and (17b) as saying that all normal chickens are hens, then we would predict that (17a) entails (17b), since all chickens that lay eggs are hens. However, that entailment clearly does not hold. The most plausible way to block it is to say that when we interpret generics, we never interpret them just by asking about what is normal *per se*. Rather, we are always concerned with what is normal in this or that respect, and the respect of normality is determined by the predicate in the generic. On this strategy, (17a) says not that all normal chickens lay eggs, but rather, that all chickens that are normal with respect to how they extrude offspring lay eggs. By contrast, (17b) says that all chickens that are normal with respect to their sex are hens. These paraphrases are such that we no longer predict (17a) to entail (17b). In fact, we intuitively predict the correct truth-values, since on this paraphrase, the latter is false.

More generally, we should revise the truth-conditions thus.

[SECOND PASS TRUTH-CONDITIONS]

As are F is true iff all F-normal As are F.

Here, *F-normal* is short for *normal in a respect determined by the predicate F*.

However, we need to complicate the semantics further to account for the examples in (18) and (19).

(18) a. Bears live in North America.
b. Bears live in South America.
c. Bears live in Europe.
d. Bears live in Asia.

(19) a. Elephants live in Africa.
b. Elephants live in Asia.

In these cases, the respect of normality is the same—normal in respect of habitat. Applying the current semantics would predict that each bear that is normal in this respect lives on four continents, which the examples in (18) obviously do not entail. The best way to deal with examples like this is to introduce ways of being normal and allow the truth-conditions to explicitly quantify over them. Thus, (18) might be paraphrased as in (20).[29]

[29] I've argued in detail that introducing such ways of being normal and quantifying over them brings significant empirical benefits. See my (2008b; 2008c).

(20) a. There is a way *w* of being a bear that is normal with respect to its habitat, and all bears that are normal in way *w* live in North America.

b. There is a way *w* of being a bear that is normal with respect to its habitat, and all bears that are normal in way *w* live in South America.

c. There is a way *w* of being a bear that is normal with respect to its habitat, and all bears that are normal in way *w* live in Europe.

d. There is a way *w* of being a bear that is normal with respect to its habitat, and all bears that are normal in way *w* live in Asia.

As these paraphrases make clear, there is no reason to think that any of the examples in (20) is incompatible with any other, since being a bear that's normal in one of the ways does not entail anything one way or another in regards to being a bear that's normal in some other way. More generally, then, the truth-conditions for characterizing sentences that I want to work with are these.[30]

[SEMANTICS FOR CHARACTERIZING SENTENCES]

> *As are F* is true iff there is a way *w* of being an *F*-normal *A* such that all *A*s that are normal in way *w* are *F*.

These semantics do not yet incorporate the modal element that accounts for the law-like flavor of characterizing sentences. We can add it simply by introducing a counterfactual element to the truth-conditions, as illustrated for a particular example in (21).

(21) a. Ravens are black.

b. There is a way of being a normally colored raven such that, if there was a raven that was colored in that way, then all ravens that would be colored in that way would be black.

In those cases in which there is a relevantly normal raven, the truth-conditions (21b) just collapse into the semantics I've highlighted, so I'll ignore this counterfactual element.

[30] Incidentally, here the locution *ceteris paribus* seems to make a genuine semantic contribution to sentences in which it appears. When one prefaces a characterizing sentence with *ceteris paribus*, that sentence is true only if it mentions all of the ways of being normal in the respect at issue in interpreting that sentence. Thus, (Ba) is true while (Bb) is false.

(B) a. *Ceteris paribus*, elephants live in Africa and Asia.

b. *Ceteris paribus*, elephants live in Africa.

But aside from this quirk, *ceteris paribus* does not seem to have any semantic impact on the interpretation of characterizing sentences.

As the statement of the semantics make clear, the only so for unexplained primitive notion is that of normality. More specifically, we need to complete the biconditional in (22).

(22) x is an A that is an F-normal A in some way w iff …

The task of the next section is to develop the resources to do just that.

5 Complex Phenomena: Causal Homogeneity and Unification

Different scientific disciplines appeal to different properties in order to formulate their respective theories. Given that these different disciplines investigate different phenomena, this is unsurprising. Different categorizing schemes will serve different investigative aims better or worse. The question I want to address is this: what makes a scheme of classification the (or a) right one for a given discipline? This is one question we can ask by asking what the natural kinds for a given range of phenomena are.

That question is distinct from various other metaphysical and linguistic issues that arise in the context of natural kinds and natural kind terms. I am not, at least in the first instance, concerned with the question whether there are any objective divisions among phenomena, where groupings that correspond to such objective divisions form a natural kind while gerrymandered groupings do not. Nor am I concerned with the question whether, assuming that there are natural kinds, the naturalness of natural kinds is a basic feature of the world, or whether that naturalness can be reduced to something else.[31] I will not ask about the connection between natural kinds and essential properties, either, for example, whether it is true that if an object belongs to a natural kind, that is an essential property of that object. Finally, I will not address the semantics of natural kind terms. I will remain silent, for example, on whether descriptivism is true as a metasemantic theory of such terms, and on whether natural kind terms retain a constant meaning or a constant reference across even substantial changes in theory.[32] These are all important questions, but my discussion can hopefully proceed while remaining neutral on all of them. I need to register only one *caveat*: I will appeal to causation in making theoretical claims. To the extent that such appeals carry commitments on the issues I just mentioned, I won't remain neutral on them. However, the

[31] For citations about these issues, see the updated bibliography in Bird and Tobin (Winter 2008).

[32] That is, I'll be silent on whether Kripke (1972) and Putnam (1975a,b) are right about natural kind terms.

notion of causation seems sufficiently basic to think that any theory of science has to be compatible with broad appeals to it. I thus hope that my account won't suffer from such reliance.

The question then is this: what are some of the desiderata that a property should meet if it is to be useful for couching theories of a particular domain? In broad outline, the answer to that question is clear: practitioners of a discipline want to appeal to those properties that allow them to formulate systematic and powerful theories.[33] The more specific claim I want to make here is that there are at least two such desiderata, which I'll call causal homogeneity and unifying power. I'll then argue that, in any science that investigates complex phenomena, these two demands are in tension, and the resolution of the tension furnishes us with the resources to explicate the notion of normality I appealed to in my semantics.[34]

To impose some order, I'll assume that each discipline, at least at a given time, has a range of phenomena it seeks to investigate and theorize about. Geography, for example, investigates the shape of the Earth's surface, the distribution and shape of rivers, and so on. Evolutionary biology investigates the myriad ways in which organisms fit into their environment and how the species to which these organisms belong evolve. I'll call these the discipline's *target phenomena*. At least one important aim of theorizing is to formulate generalizations about the shape these target phenomena actually take, for instance, generalizations about what shape rivers are, and what colors various species are. The general structure is well captured in terms of determinable and determinate relations. We can think of target phenomena in terms of properties such as the property of having some shape or other and the property of having some coloration or other, and we can think about the generalizations as telling us which determinate of that determinable is actually instantiated, such as which shape rivers actually have, and which color scheme members of a kind actually have.

[33]Many theorists who endorse this general principle in fact hold a more specific version, one that is particularly keyed to *explanatory* power. Thus, for example, Kitcher: "Natural kinds are the sets that one picks out in giving explanations. They are the sets corresponding to predicates that figure in our explanatory scheme." (Kitcher, 1984, 315n11) Similar remarks can be found in Platts (1983, 134), LaPorte (2004, 19), Rieppel and Kearney (2007, 97), and Root (2000, S629). Wilson (1996, 307), a review of Dupré (1993), makes the point that the issue of pluralism about species turns on whether there are different disciplines that are equally legitimate but require different classificatory schemes in order to achieve their aims, thus endorsing my claim, as well. Boyd (1991, 1999a,b) and philosophers following him, such as Kornblith (1993) even go so far as to suggest that a theory of natural kinds just is a theory that addresses the question I'm asking.

[34]In this discussion, I do not want to presuppose that any scheme currently employed actually meets these desiderata. Indeed, changes in such categorizing schemes are the results of investigation as much as anything else, and thus represent genuine progress. But that doesn't mean that we cannot ask about the desiderata that practitioners try to meet more and more successfully as they adapt their categorizing schemes to the phenomena they study.

5.1 Causal Homogeneity

I can now turn to my first constraint. It grows out of the informal suggestion that generalizations in scientific disciplines should map out the causal structure of the domain under investigation. There are several reasons theorists have given for endorsing this claim. One is broadly epistemological. On this strategy, we appeal to causation to explain what makes a hypothesis projectible. Roughly, the reason that *all emeralds are green* is projectible while *all emeralds are grue* is not is that the former reflects the causal structure of the world in a sense to be made more precise. [35] A second strategy begins with broadly explanatory considerations. If the generalizations of a discipline are to play a role in causal explanation,they have to correspond to the causal structure of the world.[36] The third strategy focuses on practical considerations. If we want generalizations to guide our interventions in the phenomena the sciences investigate, these generalizations need to furnish causal information about the world.[37]

The informal notion of a generalization's corresponding to or mapping the causal structure of the world needs to be made more precise in order to be useful in theorizing. I suggest that we do so by requiring that generalizations in the sciences be causally sustained. A generalization of the form *all As are Bs* is causally sustained iff all *A*s are *B* and in all of the cases in which the properties of being an *A* and being a *B* are coinstantiated, that coinstantiation is the result of the same causal mechanism. In this sense, the generalization *all gold melts at* $1948°F$. is causally sustained.

What makes the notion of a causally sustained generalization a good candidate for explicating what it means for a generalization to correspond to or map the causal structure of the world is the requirement that the mechanism that accounts for the coinstantiation of the two properties be the same mechanism in all of the cases in which that coinstantiation occurs. That is what crucially distinguishes causally sustained from merely accidental generalizations. In the accidental ones, it is true that for every case in which the properties are coinstantiated, there is a causal mechanism that accounts for it, but it is not the same in all cases.

However, one might have two concerns about the thesis that all generalizations in the sciences should be causally sustained. First, one might worry that the notion of a single causal mechanism isn't well-defined, because we can individuate causal mechanisms any way we want. For that reason, any given case in which an *A* is a

[35] See, e.g., the papers by Boyd and Kornblith cited in note 33.
[36] See, e.g., Salmon (1984, 1989) and Strevens (2004).
[37] See, e.g., Woodward (2003).

B instantiates indefinitely many mechanisms if it instantiates any, and depending on which of these mechanisms is at issue, a generalization either will or will not count as causally sustained.

The concern is reasonable as far as it goes, and it would be debilitating if there were no further constraints on how we should individuate processes. I want to contend, however, that there are such further constraint, because whether a generalization is causally sustained should be evaluated within the context of a particular discipline. This is important, because distinct disciplines have proprietary ways of individuating mechanisms. We can see this quite clearly in discussions of supervenience. It is commonplace to describe one manifestation of multiple realizability by saying that one and the same mental process can be realized by several, very different neural, biological, or chemical processes. That description presupposes that these different disciplines have their own ways of individuating types of mechanisms. This response to the concern has the consequence that a generalization may count as causally sustained by the lights of one discipline but not by the lights of another. I take this to be a favorable result, since it seems to harmonize with one prominent way of understanding what the autonomy of higher-level disciplines amounts to: being able to recognize commonalities among phenomena invisible to lower-level disciplines.[38]

Even granting that the notion of a single causal mechanism is well-defined within the context of a discipline, one might worry that the requirement that there be a single mechanism is too strong. Perhaps it is enough that there be a small number of processes that jointly sustain a generalization. However, the requirement that the generalization be sustained by a single process is more plausible so long as we keep in mind that the process is individuated in the context of a particular discipline. Let me briefly discuss an evolutionary example.[39] Populations of *E. coli* bacteria evolve very rapidly. Subjected to an environment that is low on a particular nutrient, such as glucose, such populations evolve to become more efficient at using the reduced nutrients. They can accomplish this by various different genetic paths, a fact that shows up once the evolved populations are subjected to other nutrient-deprived environments. Some of the populations that are good at using glucose are not very efficient at using fructose, while others are. The structure of the example is thus that at one level of description,

[38]This response goes only part of the way towards a resolution of the worry since it is a substantive question how to individuate disciplines. For one, simply identifying disciplines as finely (or coarsely) as university departments—biology, chemistry, etc.—is too coarse. A discipline is rather marked out by the coherence between the target phenomena, research strategies, and evidential and explanatory standards. My account trades on the existence of discplines thus picked out.

[39]This example is discussed at length in Travisano et al. (1995).

there are several causal mechanisms that underwrite the efficiency of *e. coli* at utilizing glucose. At another level of description, there is just one. And quite plausibly, the context in which we want to formulate generalizations about the adaptability of *e. coli* to nutrient-deprived environments are ones in which the particular genetic basis of the change doesn't matter.

Given that we want *generalizations* to be causally sustained, we can also draw a conclusion about a desideratum for the properties we mention in these generalizations. If a generalization of the form *all As are Bs* is to be causally sustained, then the predicate *A* in such a generalization needs to have in its extension only objects that are involved in the same causal mechanism with regards to whether they are in the extension of *B* or not. When that is the case, we can say that the *A*s are causally homogeneous with respect to whether they are *B* or not. More explicitly:

[CAUSAL HOMOGENEITY] The property of being an *A* is causally homogeneous with respect to whether *all As are B* iff all *As* are involved in the same causal mechanism in determining whether they are *B*.

For example, the property of being gold is causally homogeneous with respect to the generalization *all gold melts at* $1948°F$.

Thus, we have good theoretical grounds for saying that a property is a natural kind for a discipline only if that property is causally homogeneous with respect to the universal generalizations that discipline seeks to articulate. The example of natural kinds in the fundamental sciences certainly satisfy this condition.[40] However, this condition fits awkwardly with the special sciences, and indeed it is awkward twice over. In the first instance, the special sciences do not seek to articulate straightforward universal generalizations, though I've suggested that in fact, the cp-generalizations they do articulate are restricted universal ones. More importantly, it does not seem as if it is a goal to formulate predicates that denote causally homogeneous classes, though that would clearly be possible. For example, if we focus not simply on all of the ravens, but on ones that go through a particular developmental process *D*, we would have a property that is causally homogeneous with respect to *all ravens that go through D are black*. But we do not see a predicate that picks out such a homogeneous class. In this respect, the situation is precisely the reverse of the one encountered in the fundamental sciences. There, practitioners adjust the extension of their predicates so that they can use

[40]The condition is also in line with Putnam's suggestion that natural kinds are explanatory kinds cf PUTNAM REFERENCE AND EXPLANATION.

them to formulate straightforward, causally sustained generalizations. In the case of the special sciences, practitioners retain the predicates and give up on the straightforward universal generalizations.

This pattern isn't restricted to biology by any stretch of the imagination. We see the same thing in geography. Consider being a river. The true generalizations about all rivers whatsoever seem to mostly mark constitutive connections, such as carrying water. In order to reach a causally homogeneous property, we need to focus on a more narrowly defined one. Being a river that runs through such-and-such soil and subject to such-and-such fluctuations in water level might be causally homogeneous with respect to *all such rivers meander*, but again, we do not see simple predicates that denote such a causally homogeneous property. To understand what's going on, I want to introduce the next desideratum for a natural kind property.

5.2 Unifying Power

One of the aspects of a theory that makes it systematic is its unifying power. Intuitively, a theory has more unifying power if it contains generalizations that connect many targets of inquiry. The crucial question is how to spell out the metaphor of connecting different phenomena via a collection of generalizations. One way to do so is to find a number of true universal generalizations that all agree on their scope but differ in their predicates. Newtonian mechanics paradigmatically exhibits this kind of connection. Bodies with a certain mass have many properties in common. Moreover, the demand that the natural kinds of that discipline by causally homogeneous harmonizes well with this way of cashing out the demand for unification. Its practitioners can attempt—and are at times successful—at satisfying both of these demands.

However, in the case of the special sciences, these two demands are fundamentally in tension. Suppose that we fix on a class of ravens R_1 that is causally homogeneous with respect to *all ravens in* R_1 *are black*. R_1 will in turn not be causally homogeneous with respect to any other generalization: not only is *all ravens in* R_1 *have two wings* not causally sustained, it's false. This is a non-accidental feature of the phenomenon under investigation, specifically of the mechanism that could causally sustain a generalization. Most of these mechanisms can operate independently of each other. To see this point in an example, suppose that in addition to R_1 we focus on a class of ravens R_2 that is causally homogeneous with respect to *all ravens in* R_2 *have two wings*. The two classes R_1 and R_2 don't coincide, since there are winged albinos and black ravens that have lost a wing.

Let me now argue that this tension cannot be analyzed away in order to motivate my preferred resolution, introducing a different way of cashing out the metaphor of connecting phenomena. One might try to say that we can make the property denoted by the predicate *raven* causally homogeneous with respect to many different generalizations by intersecting all of the classes that are causally homogeneous with respect to at least one such generalization. For example, we might simply intersect R_1 and R_2, yielding a class that is causally homogeneous with respect to both generalizations. One should then perform this operation for all targets of inquiry that we can formulate an acceptable characterization about.

There's something intuitively odd about this approach, since it would exclude many ravens from the extension of the predicate *raven*, but the intuitive oddity can be countered by thinking of *raven* as picking out the typical or paradigmatic ravens. We can object to the approach more sharply by pointing out that it is subject to a dilemma. Either we intersect *all* homogeneous subclasses, or we do not. If we do, we will at least sometimes end up with an empty extension. That's because we sometimes find sets of characterizing sentences that are all true but predicate incompatible properties, such as (23).

(23) a. Lions have manes.

b. Lions give birth to live young.

The class of lions that are causally homogeneous with respect to their sexually selected head-dress and that can causally sustain (23a) excludes females; the corresponding class for (23b) includes only females. Hence, the intersection of these two classes is empty. On the other horn of the dilemma, we intersect only some causally homogeneous subclasses. In that case, we've added a bit of unifying power, since the relevant properties are causally homogeneous with respect to more generalizations, but it is still puzzling why there should be a single predicate that denotes different properties in different generalizations. We therefore cannot analyze away the tension between the demands of causal homogeneity and unifying power. Instead, we should reconceive what unifying power comes to in the special sciences.

The key innovation is two-fold, changing both the relata and the posited relation between them. Rather than focus on the predicates that appear in the different generalizations (or the properties denoted by them), we should focus on the mechanisms that causally sustain them. And rather than require that the relation be one of identity, we should focus on the existence of a common explanation. Thus, a set of generalizations is unified if the mechanisms that causally sustain them have a common explanation. In his sense, the generalization *all ravens in* R_1 *are black* and *all ravens in* R_2 *have two wings* are unified, since there is a common explanation for the presence of these mechanisms in the population of ravens. Both are the result of selection. This account also fits naturally into non-biological cases. Focus on two subsets of slow rivers. R_3 is causally homogeneous with respect to *all slow rivers in* R_3 *meander*, R_4 is causally homogeneous with respect to *all slow rivers freeze in the winter*. The generalizations are unified because both mechanisms operate in the class of slow rivers for the same reason, roughly the interaction of slowness with the environment.

This account doesn't yet explain why we see predicates like *raven* in the generalizations of the special sciences. To see that, I need to return to the semantics of characterizing sentences.

5.3 Normality

Our task is to complete the biconditional (22).

(22) *x* is an *A* that is an *F*-normal *A* in some way *w* iff . . .

I'll divide that task into three questions. Why should there be a restriction to a subset of *A*s at all—why restrict ourselves to normal members of the kind in the first place? Why, given that there is such a restriction in the first place, does that restriction depend on the predicate—why do we need normality *in a respect*? And why are there sometimes several ways of being normal in that respect?

We need to restrict ourselves to a particular subset of the *A*s in formulating a generalization about *A*s because such a generalization should be causally sustained, and hence the class of objects the generalization is about should be causally homogeneous with respect to that generalization. As we saw in the previous section, for most choices of *A* and most generalizations, the whole class of *A*s won't be relevantly causally homogeneous. Hence, we need to have a notion of normality.

Moreover, we've also seen that due to the nature of the mechanisms that can causally sustain generalizations, these mechanisms can usually operate independently of each other. That is to say, the subsets of *A*s that are homogeneous with respect to two different generalizations usually do not coincide. That in turn means that the subset needs to depend on the generalization we are considering. That is just to say that what counts as relevantly normal needs to depend on the predicate, which is why what is at issue is not just being a normal *A*, but being one that is normal in a respect determined by the predicate. Notice that here, the notion of a target of inquiry is extremely useful: the respect determined by the predicate *just is* the target of inquiry of which that predicate denotes a determinate. For example, if the predicate is *black*, then the target of inquiry is coloration, and that is just the respect of normality determined by the predicate.

Finally, we can also account for the fact that at times, there are multiple ways of being normal, even once we fix on a particular respect. Suppose that we're thinking about a particular target of inquiry *T*, such as habitat. What makes an *A* normal with respect to *T* is the fact that it goes through a causal mechanism that determines which determinate of *T* this *A* has, together with the fact that this causal mechanism bears a theoretically important relation to other mechanisms for other targets. In the case of habitat, that may be that the presence of the mechanism is the result of an adaptation. But there may be several mechanisms that determine a value for *T* and that all bear the same relation, e.g., that are all adaptations. And that just means that there are multiple ways of being normal in the respect of *T*.

Let me summarize by completing the biconditional.

[NORMALITY] Assume that we are considering the characterizing sentence *As are F* with target of inquiry T_F. Then:

x is an *A* that is an *F*-normal *A* in some way *w* iff

(i) *x* is an *A*, and

(ii) *x* is involved in a mechanism *m* that determines that *x* has one of the determinates of T_F, and

(iii) *m* stands in a theoretically important relation to mechanisms that determines that *x* or other *A*s have determinates of other targets of inquiry $T_1, \ldots, T_n$.

Put in less complex terms, what makes an *A* relevantly normal is that it goes through a theoretically important mechanism for the purposes of a particular target of inquiry. And what makes the mechanism theoretically important is a relational feature of that process, not an intrinsic one. We can now put this theory to some explanatory work.

For example, we can now explain why practitioners of the special sciences retain predicates like *raven*, rather then refining them in order to be able to use them in formulating straightforward universal generalizations that are causally sustained. If they did that, they would be stuck with a very large number of predicates that would appear to be unrelated and hence would not mark out that we are in fact capturing a significant connection among phenomena.

We can also explain why we see the asymmetries I drew attention to in my discussion of Pietroski and Rey (§3.3). There, I noted that *ravens are black* is true and *ravens are white* false, even though there are mechanisms that underwrite appropriately restricted universal versions of each. The difference between these two characterizing sentences cannot be explained by saying that there only is a causal mechanism to underwrite the former, since there are stable mechanisms for both. Rather, the difference consists in the relational features of these mechanisms. Only the one that sustains *ravens are black* is appropriately related to other mechanisms instantiated in the population of ravens, while the latter is not.

This is also why qualifying the subject yields the true *albino ravens are white*. Once we move away from the original subject *ravens*, the mechanism that sustains the generalization about albino ravens only has to bear theoretically important relations to other mechanisms operating in albinos.

Drawing on the same resources, I now want to give an account of the open-endedness of characterizing sentences as well as their sensitivity to large scale changes in the way the world is, but not to small scale ones. The account of open-endedness is essentially the same as the one I presented in discussing Pietroski and Rey in §3.2. I, too, interpret characterizing sentences in terms of the operation of certain mechanisms, and these mechanisms can be derailed in a large number of very different ways. But unlike Pietroski and Rey, I can derive this aspect from the basic assumption that generalizations in the special sciences should be causally sustained.

The account of the contrast between small scale and large scale change depends on a point about causal homogeneity I haven't drawn attention to so far. The causal mechanism or mechanisms that underwrite being an *F*-normal *A* are individuated, in part, by their endpoints. Thus, in the example of *ravens are black*, the process that makes a raven a normally colored raven is one that a raven cannot go through without being black. This observation has the consequence that if it's true that *As are F*, then the process that makes an *A* an *F*-normal *A* (in some way) is such that all *A*s that go through it are *F*. That in turn means that once we've fixed on the mechanism that makes an *A* *F*-normal, we've fixed the truth-value of *As are F*. There is no further variation once the mechanism has been fixed. This is the crucial point for understanding why a change in the truth-value of a characterizing sentence always requires large-scale changes.

In order for a characterizing sentence *As are F* to differ in truth-value when it's evaluated with respect to two different worlds, the mechanism that makes an *A F*-normal has to differ between the worlds. As I just argued, what makes a mechanism the one that certifies an *A* that is involved in it as *F*-normal is the fact that it bears a theoretically important relation to other mechanisms for other targets of inquiry. And whether a mechanism bears that relation can only change if there are large scale changes. For example, if we are dealing with a coloration mechanism that is normal in virtue of being the result of an adaptation, then it can only cease to be the normal coloration mechanism if it is no longer the result of an adaptation. And that requires a large-scale change.

On the theory I've presented, whether a mechanism that sustains a generalization can qualify the objects involved in it as appropriately normal depends on whether that mechanism bears theoretically important relations to other mechanisms sustaining other generalizations. But which relations are theoretically important depends on the theory and that suggests that when there are different theories that investigate the same target phenomena, we can see that different mechanisms count as bearing the relevant relations. Here is an example that shows this possibility to be actual.

Dobermans are born with floppy ears. In some countries, including the US, breeders then cut off parts of their ears and temporarily insert posts into them in order to make the ears grow in a pointy shape. With that information in mind, consider the following texts.

[BIOLOGY] Some breeds of dogs have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, and as that fact might lead you to believe, *Dobermans do not have pointy ears*. Dobermans have floppy ears.

[DOG-SHOW] Welcome to this year's meeting of the Westminster Kennel Club. Some of our breeds have a more relaxed, homely appearance, especially those with floppy ears. Dobermans, however, are regal. Dobermans do not have floppy ears. *Dobermans have pointy ears*.

In the context of their respective texts, each of the italicized claims is true. That shows that the interpretation of these claims depends on the context in which they are produced. On my account, there is a natural explanation of what is going on. In BIOLOGY, the natural kinds are the kinds of evolution, and any mechanism that involves the actions of breeders fails to be related to other mechanisms in theoretically important ways. By contrast, in DOG-SHOW, we are precisely interested in dog-breeding, and hence the mechanisms that involve (at least some of) the actions of breeders are the ones that contribute to making the property of being a doberman a natural kind property for that theoretical enterprise.

Another example offers further support to this analysis. One widely-remarked upon feature of characterizing sentences is that they can be true, even though the minority of members of the kind at issue conform to the generalization, and even if it's not an accident in any sense of the term. A famous example is (24).

(24) Sea-turtles are long-lived.

Sea-turtles reproduce by laying their eggs on the beach. They all hatch at the same time, and the newborn turtles have to race to the ocean in order to develop further. On that race to the ocean, most turtles are eaten by predators who are there, awaiting the turtles' appearance. Thus, most turtles die young, and this is a completely stable fact. Indeed, turtles "plan" for this culling by laying very many eggs.

The example illustrates several important points. First, one might have responded to the doberman-example by saying that what accounts for the difference in truth-value is simply the salience of dobermans with the relevant shape of ears. Somehow, because we picture a doberman with floppy ears, we find *dobermans don't have pointy ears* to be acceptable in the context of BIOLOGY. But such psychological or perceptual salience cannot account for what happens in (24), since even standing at the beach looking out at the turtles being eaten, we still find (24) acceptable, and *turtles die young* to be false.

Second, the example illustrates that an alternative theory of what makes a causal mechanism eligible to underwrite a causally sustained characterizing sentence is mistaken. One might have said that, in the context of a particular discipline, a causal mechanism is eligible to so underwrite a generalization if it is the kind of mechanism studied in the discipline, and/or it has important effects on the targets of inquiry in that discipline. However, the predator/prey interaction that causes so many young turtles to die satisfies both of these conditions, and still isn't eligible. Clearly, predator/prey interactions are part of what is studied in evolutionary biology, so it's the right kind of mechanism. And such predator/prey interactions also have important effects—for example, they influence how many eggs are laid. By contrast, my account deals with this example neatly, because the slow metabolism of turtles that causally sustains the generalization (24) is related to other mechanisms that underwrite other characterizing sentences by being an adaptation.

6 Conclusions

I have offered a new account of how cp-laws fit into the theoretical aims and demands of the special sciences. To find a class of cp-laws that are properly treated together, I've appealed to broadly linguistic considerations. I've argued that some cp-laws are most naturally stated by the use of characterizing sentences, and that these sentences should be interpreted in terms of a notion of normality. I then argued that such a notion of normality can be defined, albeit non-reductively, by developing resources via consideration of natural kinds in the special sciences. One point I want to draw attention to is that the claims I've made about natural kinds in the special sciences are very weak. Thus, my account should be acceptable to theorists with quite different commitments.

Let me finish by raising two questions. The first is how to extend this treatment to other cp-generalizations, especially ones that are stated by using habituals, among which there are many ascriptions of dispositions. The second is how this treatment of characterizing sentences can be extended to non-scientific contexts. We obviously make use of characterizing sentences outside of formal scientific inquiry. My semantics suggest that, even in these contexts, we should somehow have the resources to interpret a notion of normality with the structure I identified. It is an open question what psychological and epistemological mechanisms underlie this possibility in the informal context of everyday life.

References

Armstrong, D. M. *What is a Law of Nature?* (Cambridge, UK: Cambridge UP, 1983).

Austin, J. L. *How To Do Things With Words* (Cambridge, MA: Harvard UP, 1975).

Barwise, J. "Constrains, Channels, and the Flow of Information". In P. Aczel, D. Israel, Y. Katagiri, and S. Peters, eds., *Situation Theory and Its Applications, Vol. 3*, 3–29 (Stanford, CA: CSLI, 1993).

Barwise, J. and J. Seligman. "The Rights and Wrongs of Natural Regularity". In J. E. Tomberlin, ed., *Philosophical Perspectives 8: Logic and Language*, 331–364 (Atascadero, CA: Ridgeview, 1994).

Bird, A. and E. Tobin. "Natural Kinds". In E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Winter 2008).
URL: *http://plato.stanford.edu/archives/win2008/entries/natural-kinds/*

Boyd, R. N. "Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds". *Philosophical Studies*, 61(1-2), (1991), 127–148.

—. "Homeostasis, Species, and Higher Taxa". In R. A. Wilson, ed., *Species: New Interdisciplinary Essays*, 141–185 (Cambridge, MA: MIT Press, 1999a).

—. "Kinds, Complexity and Multiple Realization: Comments on Millikan's *Historical Kinds and the Special Sciences*". *Philosophical Studies*, 95(1-2), (1999b), 67–98.

Cartwright, N. *How the Laws of Physics Lie* (Oxford: Clarendon Press, 1983).

—. *Nature's Capacities and Their Measurement* (Oxford: Clarendon Press, 1989).

—. "*Ceteris Paribus* Laws and Socio-Economic Machines". *Monist*, 78, (1995), 276–297.

—. "Models and the Limits of Theory: Quantum Hamiltonians and the BCS Model of Superconductivity". In M. Morrison and M. S. Morgan, eds., *Models as Mediators*, 241–281 (Cambridge, UK: Cambridge UP, 1999).

Cavedon, L. and S. Glasbey. "Outline of an Information-Flow Model of Generics". *Acta Linguistica Hungarica*, 42(3-4), (1994), 227–245.

Chomsky, N. "Questions of Form and Interpretation". In R. Austerlitz, ed., *The Scope of American Linguistics*, 159–196 (Lisse: Peter de Ridder Press, 1975).

Cohen, A. *Think Generic!* (Stanford, CA: CSLI Publications, 1999).

Davidson, D. "The Logical Form of Action Sentences". In *Essays on Actions and Events*, 105–148 (Oxford: Oxford UP, 1980).

Dray, W. *Laws and Explanations in History* (Oxford: Oxford UP, 1957).

Dupré, J. *The Disorder of Things* (Cambridge, MA: Harvard UP, 1993).

Earman, J. and J. Roberts. "*Ceteris Paribus*, There is no Problem of Provisos". *Synthese*, 118(3), (1999), 439–478.

Earman, J., J. Roberts, and S. Smith. "*Ceteris paribus* Lost". *Erkenntnis*, 57(3), (2002), 281–301.

Fara, M. *Dispositions and Their Ascriptions*. Ph.D. thesis, Princeton University (2001).

—. "Dispositions and Habituals". *Noûs*, 39(1), (2005), 43–82.

Fodor, J. A. "You can fool some of the People All of the Time, Everything Else Being Equal; Hedged Laws and Psychological Explanations". *Mind*, 100(1), (1991), 19–34.

Giere, R. N. *Explaining Science* (Chicago: University of Chicago Press, 1988a).

—. "Laws, Theories, and Generalizations". In A. Grünbaum and W. Salmon, eds., *The Limits of Deductivism*, 37–46 (Berkeley, CA: University of California Press, 1988b).

—. "Science without Laws of Nature". In *Science without Laws*, 84–96 (Chicago: University of Chicago Press, 1999).

Glymour, C. "A Semantics and Methodology for Ceteris Paribus Hypotheses". *Erkenntnis*, 57(3), (2002), 395–405.

Hall, N. "Causation and Ceteris Paribus Laws". *The Harvard Review of Philosophy*, 13(1), (2005), 80–99.

Hempel, C. G. *Aspects of Scientific Explanation* (New York: The Free Press, 1965).

Kitcher, P. "Species". *Philosophy of Science*, 51(2), (1984), 308–333.

Kornblith, H. *Inductive Inference and Its Natural Ground* (Cambridge, MA: MIT Press, 1993).

Krifka, M., F. J. Pelletier, G. N. Carlson, A. ter Meulen, G. Chierchia, and G. Link. "Genericity: An Introduction". In G. N. Carlson and F. J. Pelletier, eds., *The Generic Book*, 1–124 (Chicago: University of Chicago Press, 1995).

Kripke, S. A. *Naming and Necessity* (Cambridge, MA: Harvard UP, 1972).

Lange, M. *Natural Laws in Scientific Practice* (Oxford: Oxford UP, 2000).

LaPorte, J. *Natural Kinds and Conceptual Change* (Cambridge, UK: Cambridge UP, 2004).

Leslie, S.-J. "Generics and the Structure of the Mind". In J. Hawthorne, ed., *Philosophical Perspectives 21, Philosophy of Mind*, 375–403 (Malden, MA: Blackwell, 2007).

—. "Generics: Cognition and Acquisition". *Philosophical Review*, 117(1), (2008), 1–47.

Lewis, D. K. “Adverbs of Quantification”. In E. L. Keenan, ed., *Formal Semantics of Natural Language*, 3–15 (Cambridge, UK: Cambridge UP, 1973a).

—. *Counterfactuals* (Cambridge, MA: Harvard UP, 1973b).

Mäki, U. *Fact and Fiction in Economics: Models, Realism, and Social Construction* (Cambridge, UK: Cambridge UP, 2002).

Milsark, G. *Existential Sentences in English*. Ph.D. thesis, MIT (1974).

Morrison, M. and M. S. Morgan. *Models as Mediators* (Cambridge, UK: Cambridge UP, 1999).

Mott, P. “Fodor and Ceteris Paribus Laws”. *Mind*, 101(402), (1992), 335–346.

Nagel, E. *The Structure of Science* (Indianapolis: Hackett, 1979).

Nickel, B. “*Dutchmen are Good Sailors* and Other Generic Comparisons” (2008a). Ms., Harvard University.

—. “Free Choice and Generic Predication” (2008b). Ms., Harvard University.

—. “The Ways of Normality” (2008c). Forthcoming in *Linguistics and Philosophy*.

Pietroski, P. M. and G. Rey. “When Other Things Aren’t Equal: Saving Ceteris Paribus Laws from Vacuity”. *The British Journal for the Philosophy of Science*, 46(1), (1995), 81–110.

Platts, M. “Explanatory Kinds”. *The British Journal for the Philosophy of Science*, 34(2), (1983), 133–148.

Putnam, H. “Explanation and Reference”. In *Mind, Language, and Reality: Philosophical Papers Vol. 2*, 196–214 (Cambridge: Cambridge UP, 1975a).

—. “The Meaning of ‘meaning”’. In *Mind, Language, and Reality: Philosophical Papers Vol. 2*, 215–271 (Cambridge: Cambridge UP, 1975b).

Rieppel, O. and M. Kearney. “The Poverty of Taxonomic Characters”. *Biology and Philosophy*, 22(1), (2007), 95–113.

Root, M. “How We Divide the World”. *Philosophy of Science*, 67 (Supplement)(II), (2000), S628–S639.

Rosenberg, A. “How is Biological Explanation Possible?” *British Philosophy of Science*, 52, (2001a), 735–760.

—. “Reductionism in a Historical Science”. *Philosophy of Science*, 68, (2001b), 135–163.

Salmon, W. C. *Scientific Explanation and the Causal Structure of the World* (Princeton, New Jersey: Princeton UP, 1984).

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

Schiffer, S. "Ceteris Paribus Laws". *Mind*, 100(1), (1991), 1–17.

Schubert, L. K. and F. J. Pelletier. "Generically Speaking, or, Using Discourse Representation Theory to Interpret Generics". In G. Chierchia, B. H. Partee, and R. Turner, eds., *Properties, Types, and Meaning, Vol. II*, 193–268 (Dordrecht: Kluwer Academic Publishers, 1989).

Schurz, G. "Pietroski and Rey on Ceteris Paribus Laws". *British Journal for the Philosophy of Science*, 52(2), (2001), 359–370.

—. "*Ceteris paribus* Laws: Classification and Deconstruction". *Erkenntnis*, 57(3), (2002), 351–372.

Silverberg, A. "Psychological Laws and Non-Monotonic Logic". *Erkenntnis*, 44(2), (1996), 199–224.

Strevens, M. "The Causal and Unification Approaches to Explanation Unified—Causally". *Noûs*, 38(1), (2004), 154–176.

Travis, C. "On What is Strictly Speaking True". *Canadian Journal of Philosophy*, 15(2), (1985), 187–229.

—. *Unshadowed Thought* (Cambridge, MA: Harvard UP, 2000).

Travisano, M., F. Vasi, and R. E. Lenski. "Long-Term Experimental Evolution in Escherichia coli. III. Variation Among Replicate Populations in Correlated Responses to Novel Environments". *Evolution*, 49(1), (1995), 189–200.

Van Fraassen, B. *The Scientific Image* (Oxford: Oxford UP, 1980).

—. *Laws and Symmetries* (Oxford: Oxford UP, 1989).

Veltman, F. "Defaults in Update Semantics". *Journal of Philosophical Logic*, 25(3), (1996), 221–261.

Wilson, R. A. "Promiscuous Realism". *Philosophy of Science*, 47(2), (1996), 303–316.

Woodward, J. "There is no such Thing as a *ceteris paribus* Law". *Erkenntnis*, 57(3), (2002), 303–328.

—. *Making Things Happen* (Oxford: Oxford UP, 2003).