

# Processes in the Interpretation of Generics and CP-Laws\*

Bernhard Nickel · Harvard University

September 20, 2007

## 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 taken this feature to introduce the presumption that cp-laws are trivial, one that needs to be countered if we are to appeal to cp-laws in the course of scientific investigation or our philosophical theorizing about it. I argue that the triviality worry is misplaced by pointing out that cp-laws are just a subset of uncontroversially meaningful and contingent expressions of natural language, the generics. I then present an account of these generics that elucidates some of their most puzzling features, especially the ones that suggested the triviality worry in the first place.

## 1 Introduction

So-called *ceteris paribus*-laws (cp-laws) are claims we make in the course of pursuing the non-fundamental sciences, such as biology, psychology, history, and economics. An example is (1).

(1) (*ceteris paribus*), ravens are black.

There are many problems about the analysis of laws generally. But in addition, there is one specific to cp-laws. We qualify (1) with the phrase ‘*ceteris paribus*’ to signal that we are ready to accept (1), even though there are some ravens that are not black. The

---

\*[[Acknowledgments suppressed]]

fact that (1) is compatible with the existence of exceptions in the way that it is—more on this way in a minute—is particularly troublesome.

According to most theorists, we can best sharpen the debate to a useful question thus. There is at least a *prima facie* case to be made that generalizations like (1) are trivial. If trivial, they cannot capture substantive hypotheses about the world, making the disciplines in which results are characteristically stated in terms of cp-laws either charlatanism or at least up to something quite different than giving us ordinary, propositional knowledge of the world. Thus most theorists ask: are cp-laws really trivial?

It is a striking fact that linguists and philosophers of language investigate a rather similar phenomenon under the heading of *generics*. These are claims we make in the course of everyday life rather than scientific inquiry, but they too seem to be generalizations that tolerate exceptions, such as (2).

(2) Ravens are black.

It certainly does not mean that all ravens are black, nor that some are. Just like (1), we are ready to assert it even though we know that some ravens are not black.

Indeed, I shall argue that this is more than a striking coincidence. The sentences we use to state cp-laws are just a subset of the generics. I make this point because it will allow us to recast the debate about cp-laws in the non-fundamental sciences in different terms, and then to make progress on some genuinely puzzling features of both cp-laws and generics. We shouldn't ask whether cp-laws are trivial. We should instead ask directly about what puzzles us. I'll put more meat on this contrast shortly.

In section 2, I'll introduce the standard way of framing the debate about cp-laws, focused on the concern about triviality. I'll argue that this concern is easily dealt with, and that it obscures important questions about cp-laws and generics. In section 3, I present a two-stage analysis of some generics, which I then apply to answer the questions in section 4. I address concerns about the empirical adequacy of my treatment in section 5.

## 2 CP-Laws, Triviality, and the Questions

We may write any proposed law in the form *it is a law that p*. For many laws, it is relatively easy to say what is expressed by *p*, the complement of the nomic operator, and much attention has been paid to analysing that operator. But in the case of cp-laws,

the distinctive problem concerns the meaning of that complement. I want to focus on this problem, and to do so, I want to call sentences like (1) *cp-generalizations*, leaving aside the question whether they are properly called laws.

### *Motivating Triviality*

The problem usually discussed under the heading of *triviality* 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*. So we might try to spell out (1) as (3).

- (3) All ravens are black, unless they are albinos.

But (3) does not capture the force of (1), since ravens that are non-black because of the environmental conditions they experience, rather than the genetic endowment they are born with, would falsify (3) 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.<sup>1</sup> 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.<sup>2</sup> 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

---

<sup>1</sup>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.

<sup>2</sup>Pietroski and Rey (1995, p. 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, p. 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).

trivial. That is the *triviality worry*. Applied to my example, it says that ‘ravens are black’ 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 exception 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, 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 a 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, claiming that cp-generalizations cannot be empirically confirmed or disconfirmed.<sup>3</sup> 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, that anything goes. For that reason, the concern about confirmation stands and falls with the triviality worry, and I won’t discuss it separately here.

### *CP-Generalizations are Contingent*

However, pursuing the triviality worry is not the right way to get at what is really puzzling about cp-generalizations. The worry is easily met, but without addressing the originally puzzling features of cp-generalizations, their open-endedness and the impossibility of genuine, individual exceptions. I’ll first adduce two considerations to show that cp-generalizations definitely aren’t trivial, and then state different questions

---

<sup>3</sup>See for example Earman et al. (2002, p. 293) and Schurz (2002, pp. 360-2).

that are more directly about their puzzling features. These are the questions I'll be concerned to answer in this paper.

CP-generalizations do not, in general, behave at all like trivial claims.<sup>4</sup> If they were trivial, we would predict all of the following to be true.

- (4) a. Ravens are black.
- b. Ravens are polka-dotted.
- c. Ravens are white.
- d. Ravens have the color scheme of Crimson Rosellas.

Uncontroversially, (4a) is a far better thing to say in just about any context than any of (4b)-(4d). Now, 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 being true. However, if (4b) really were true, denying it would yield a falsehood. So if (4b) is true, (5) must be false.

- (5) It is not the case that ravens are polka-dotted.

But (5) is obviously acceptable, and this speaks strongly in favor of its truth, which in turn requires (4b) to be false.<sup>5</sup> Analogous arguments are available for (4c) and (4d).

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 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.

---

<sup>4</sup>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 this claims when I talk about the non-triviality of cp-generalizations

<sup>5</sup>Observe that in (5), the negation has wide scope over the whole sentence. That is, (5) denies something that, according to the theory we're considering is, trivially true. (5) does not just say that ravens have the property of being not-polka-dotted, which is true by anybody's lights.

### *CP-Generalizations are Generics*

Here is another reason to reject the triviality worry. One might think that cp-generalizations make use of a special, technical vocabulary, the expression “*ceteris paribus*.” When we make use of technical vocabulary, we need to show that it is meaningful—technical jargon requires proper constraints on use.

Showing this need not take the form of giving an explicit and reductive definition of the jargon. One could simply exhibit a well-confirmed, predictively and explanatorily powerful theory in which that jargon plays a significant role. This is why, one might think, 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 this way of showing that the technical machinery of cp-generalizations is legitimate is unavailable. We are left with the unpalatable option of trying to provide an explicit definition. Short such a definition, we face the triviality worry.

We can disarm this alternative source of the worry by showing that cp-generalizations do not make use of technical jargon, or special technical senses of otherwise ordinary expressions. Rather, cp-generalizations *just are* expressions of ordinary language.

I agree with Schiffer (1991, p. 10) and Woodward (2002, p. 305), who point out that claims that intuitively express cp-generalizations need not be stated with the explicit *ceteris paribus* qualification. Hence, the question of whether they involve special technical vocabulary reduces to the question whether the constructions practitioners of the non-fundamental sciences use have a special, non-standard sense.

I deny that they do. Paradigmatic cp-generalizations behave *exactly* as a class of natural language expressions do. This class is a subset of so-called *generics*. Genericity in natural language encompasses a range of phenomena that differ both in how they are expressed and in their semantic properties. I’ll focus on sentences with subject terms that are in the plural but are not preceded by a definite article or any explicit quantifier, so-called *bare plurals*, such as ‘ravens are black.’

Among the bare plurals, we again find sentences with very different semantic properties. Some bare plurals seem to make claims directly about kinds—*genera*, hence the name *generics*—such as ‘dodos are extinct.’ Others seem to instead state a generalization about members of the kind, such as ‘dogs bark.’ Following the terminology of Krifka et al. (1995, pp. 2-3), I’ll call these latter *characterizing sentences*. They are the sentences I’ll be concerned with. It is currently a matter of debate where to draw

the line among the bare plurals. My strategy here is to state a sufficient condition for being a characterizing sentence, one that points to the core cases. We can then extend the definition as theoretical considerations warrant.

[CHARACTERIZING SENTENCE] A sentence *S* is a characterizing sentence if

*S* is of the form *As are F*, *S* expresses a generalization, and *S*'s truth is compatible with what would be counterexamples to the corresponding universal generalization.

The first clause focuses us on a particular surface form, the bare plurals. The second excludes a number of types of sentences that share this form but seem like good pretheoretic bets to be different kinds of sentences, including kind-predicating sentences—e.g., ‘dodos are extinct’—as well as straightforward existential claims like ‘dogs are on my lawn.’<sup>6</sup>

Given these stipulations, I claim that cp-generalizations just are characterizing sentences. My reason is that both kinds of sentence share several semantic properties. Aside from the fact that both tolerate exceptions, they are easily paraphrased by inserting ‘normal’ or ‘normally’; they usually cannot be paraphrased by using a progressive form; and they exhibit neither upward nor downward entailment.<sup>7</sup>

Since cp-generalizations are generics, and as such are uncontroversially meaningful, we should reject the triviality worry. However, I certainly do not take this discussion to show that there are no problems about cp-generalizations. The spirit that gave rise to the worry is still very much with us.

Disarming the triviality worry was still worth-while, however, since it set the parameters of the debate. Theorists who took themselves to have to argue against it found themselves trying to give an explicit, reductive definition. As the criticisms of Earman and Roberts (1999) and Earman et al. (2002) show, that goal is almost impossible to accomplish. Once we reject the triviality worry, however, we can make progress on the remaining questions by giving a non-reductive account of characterizing sentences, connecting them to other notions without attempting to eliminate them completely. That is what I aim to do here.

---

<sup>6</sup>So when I talk about a generalization in the definition of characterizing sentences, I should be taken to mean something stronger than just an existential generalization.

<sup>7</sup>For a longer discussion of these features, see Krifka et al. (1995, pp. 8-14).

### *The Remaining Questions*

Specifically, we should like to see answers to three questions.

- What distinguishes the merely apparent exceptions from the remaining instances of the kind?
- Why is the class of merely apparent exceptions to a characterizing sentence so heterogeneous?
- Why do only large-scale variations in the world make for variations in the truth-value of a characterizing sentence?

The first two are about the open-endedness of characterizing sentences, the third about the fact that, though these sentences are contingent, we cannot imagine individual genuine exceptions to them. Ideally, answers to the first two questions will also show how speakers with limited cognitive resources can represent the relevant division, since we obviously do not memorize by rote what counts as a merely apparent exception.

## 3 A Two-Stage Analysis

To answer these questions, I present a two-stage analysis of characterizing sentences. In the first stage, I analyze at least very simple characterizing sentences as restricted universal generalizations. For example, ‘ravens are black’ says roughly that all normal ravens are black. In the second stage, I give a non-reductive analysis of normality in terms of processes. A raven counts as normal just in case it goes through a certain kind of process, what I’ll call a *characteristic* process. This stage of the analysis is non-reductive, in that I will not be able to say what makes a process characteristic in a given case without appeal to normality. But as I’ve argued, that’s alright.

Ultimately, I think that the two stages of this analysis correspond to a difference between the logical form of characterizing sentences and the interpretation of an element in that logical form. That is, I believe that their logical form contains an unpronounced element analogous to the predicate ‘normal.’ I also believe that other features of the logical form conspire to introduce the force of a universal generalization, and I believe that we can further analyze ‘normal’ in terms of processes, although I do not believe that explicit mention of processes enters the logical form.

That my two stages of analysis correspond to what is part of the logical form and what is not is a substantive claim, well beyond the scope of this paper. I mention it

here because I want to indicate how my account can fit into an empirically adequate analysis of characterizing sentences. For the purposes of this paper, we can take the two stages merely as an expository device.

### *Simple Characterizing Sentences*

I will only sketch an analysis of very simple characterizing sentences. I restrict myself in this way because the semantics of very many characterizing sentences exhibit interactions with other, complex and ill-understood constructions. Specifically, I'll limit myself to what I'll call *simple* characterizing sentences.

[SIMPLE CHARACTERIZING SENTENCES]

A characterizing sentence *As are F* is simple iff

- the main verb is in the present tense;
- the subject is not a mass term;
- the predicate is distributive and intersective;
- the predicate is not logically complex (doesn't contain conjunction or disjunction); and
- there is no other true characterizing sentence *As are F and G*.

The first two conditions should be self-explanatory. Let me comment briefly on the remaining ones. The difference between distributive and non-distributive predication only arises for plural nouns. Sometimes, we use them to summarize what a number of people or things did individually, as in 'five students put on gloves.' Sometimes, we use them to say what a number of people or things did together, even though they did not perform the same action individually, as in 'five students formed a circle.' In the former case, the predicate is distributive. It distributes across each of the things picked out by the subject term. In the latter case, the predicate is non-distributive. Generally: a predicate *G* is *distributive* iff the inference from *some/five/all/many Fs are G* to *each F is G* is valid.

I also restrict myself to intersective predicates. Suppose we have two predicates *F* and *G*, and consider the expression *F Gs*. If the extension of the larger expression is the intersection of the extension of the two constituent ones, the constituent predicates are intersective. For example, the predicates 'black' and 'raven' are intersective because the extension of 'black raven' is just the intersection of the black things and the ravens.

Both non-distributive and non-intersective predicates interact in complex ways with the plurality of the subject. That makes it hard to give genuinely compositional semantics for bare plurals containing these trickier expressions. But these are problems that bare plurals share with other plural expressions, such as plural definite descriptions. Consider these examples.

- (6) a. Ravens are widespread.  
b. All normal ravens are (individually) widespread.
- (7) a. Girls do better than boys in grade school.  
b. All normal girls do better than all normal boys (=than any normal boy) in grade school.
- (8) a. The ravens in my study are black.  
b. All the ravens in my study are black.
- (9) a. The books in my study are scattered.  
b. All the books in my study are (individually) scattered.
- (10) a. The girls in the fifth grade of MLK do better than the boys.  
b. All girls in the fifth grade of MLK do better than all the boys (=than any boy).

Clearly, the (b)-paraphrases of the (a)-examples in (6) and (7) are non-starters. But that, by itself, doesn't show that we've got the truth-conditions for simple characterizing sentences wrong, as the analogy with plural definite descriptions suggests. The analysis of (8a) as (8b) seems very promising, even though the (b)-paraphrases of the (a)-examples in (9) and (10) are, just as in the case of the bare plurals, completely hopeless.<sup>8</sup> That suggests that the problem arises from how the non-distributive or non-intersective predicates interact with the plural subjects. An extension of my account to characterizing sentences that are not simple in my sense needs to await a more detailed understanding of these other constructions. Clearly, a final semantics needs to predict the preliminary analysis I am giving here as a special case, otherwise my analysis is untenable. For now, I'll proceed in the hope that this is possible.

---

<sup>8</sup>The analysis of (8a) as (8b) is in line, for example, with the influential view proposed in Sharvy (1980).

Now for the last pair of restrictions. Characterizing sentences with syntactically complex predicates behave differently from quantified sentences with quantifiers that are downward entailing in their second argument. That is to say, while we can always draw the inference from *All As are F* to *All As are F or G* for any predicate *G*, we cannot generally draw the inference from *As are F* to *As are F or G*. The contrast is illustrated in (11) and (11).

- (11) a. All emeralds are green.  
b. All emeralds are green or paisley.
- a. Emeralds are green.  
b. Emeralds are green or paisley.

Further, *As are F and G* does not require any individual *A* to be both *F* and *G*, whereas *All As are F and G* does, as is illustrated in the contrast between (12) and (13).

- (12) Cats are black, white, tabby, and orange.
- (13) All cats are black, white, tabby, and orange.

This last is clearly a feature specific to characterizing sentences, and it merits further investigation. I believe, however, that the truth-conditions for simple generics I present here will be a special case of a more general semantic theory that deals with these complex predicates.<sup>9</sup>

### *The Normality Analysis*

As a first step, I propose that we analyze (14a) as (14b).<sup>10</sup>

- (14) a. Ravens are black.  
b. All normal ravens are black.

---

<sup>9</sup>See my (forthcoming).

<sup>10</sup>For the idea of interpreting generics as having something to do with quantification and normality, see, for example, Asher and Morreau (1995), Dahl (1975), Farkas and Sugioka (1983), Heim (1982), Krifka et al. (1995), and Lawler (1973). However, the particular proposals differ both from each other and from my own, on the role of references to normality.

Even for simple sentences, the final analysis needs to be a lot more complex, but one additional layer of complexity will do for us.<sup>11</sup> Consider the pair of sentences in (15).

- (15) a. Chickens lay eggs.  
b. Chickens are hens.

Clearly, (15a) does not entail (15b) but the simple analysis predicts that it does. By that analysis, (15a) means that all normal chickens lay eggs. Since only hens lay eggs, it follows that all normal chickens are hens. And this is just the meaning the simple analysis predicts for (15b).

There is a popular answer to this problem, one that I endorse, as well. The restriction of the domain of quantification is, in part, a function of the predicate, not just of the subject term. In my case, that takes the form of saying that what is at issue is not normality *per se*, but normality *in a respect*.<sup>12</sup> We should analyze (15a) and (15b) as (16a) and (16b), respectively.

- (16) a. All chickens that are normal with respect to how chickens give birth lay eggs.  
b. All chickens that are normal with respect to the sex of chickens are hens.

(16a) does not entail (16b), and hence we can consistently account for the intuitive truth-value judgments about (15a) and (15b).

More generally, I'll work with these semantics.

[SEMANTICS FOR SIMPLE CHARACTERIZING SENTENCES]

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

Being an *F*-normal *A* is short for being an *A* that is normal with respect to a certain determinable of *F* of *As*. For example, if *F* is being black, then being an *F*-normal *A* is being an *A* that is normal with respect to the color of *As*.<sup>13</sup>

One can have at least two different attitudes towards the appeal to normality in these semantics. Either, we take 'normal' there to have the same meaning as it does in ordinary English. In that case, the analysis is complete, but it is also dependent on

---

<sup>11</sup>For details on more complications, see my (2006).

<sup>12</sup>See for instance the treatment in Cohen (1999). Cohen achieves the same results by different means.

<sup>13</sup>This determinable is determined as a function of the property of being *F*. I'll remain silent on how.

our ordinary understanding of normality. Alternatively, and this is my attitude, we take ‘normal’ as a heuristic place-holder that needs to be further explicated and constrained. That is the job of the second stage of the analysis.

### *Analyzing Normality In Terms of Processes*

Here is my basic hypothesis. Every simple characterizing sentence *As are F* has an associated *characteristic* process. By this process, an *A* comes to have one of the determinates of the determinable involved in being *F*-normal. Going through this process, and having the outcome of the process persist unaltered, is necessary and sufficient for being an *F*-normal *A*.

For example, the characteristic process associated with ‘ravens are black.’ is one by which a raven comes to have a color, since being colored is the determinable at issue in that sentence. This process is one among many actual and possible processes by which ravens come to be some color or other. Going through this process is necessary and sufficient for being a raven that is normal with respect to the color of ravens.

Processes, in this context, are sequentially ordered sets of events. These events need not stand in cause-and-effect relations, though they often will. The events need not only concern the intrinsic features of the *A* going through the process. In most cases, they won’t. For example, many of the events that constitute the characteristic process for ‘ravens are black’ are events of a raven’s eating a certain standard diet. One day’s eating will not usually causally depend on the previous day’s—the raven might have eaten the same way even if it had skipped dinner—and eating a certain kind of food will not usually be an event that involves only intrinsic properties of the raven doing the eating.

Events that constitute a characteristic process are a very broad category. They include very many things that, in ordinary parlance might not be naturally called ‘event,’ such as nothing’s happening to a raven for five minutes, or its being black. Neither of these “events” constitute the raven’s undergoing a change, but that does not disqualify this chunk of the raven’s life from being an event in my sense.<sup>14</sup>

Events can be very flexibly individuated, at times coarser, at times finer. I make this point about event-individuation because I have been speaking about *the* characteristic process associated with a sentence, implying uniqueness. Given that events are so flexible, that isn’t much of a commitment. In principle, there could be a single process

---

<sup>14</sup>Events in my sense are very similar to Lewis’s. See his (1986).

that every raven goes through in the course of its life, so long as the events that constitute the process are individuated with sufficient coarseness: being born in some way or other, developing, dying. Obviously, such a process is not terribly interesting, and it is unclear what use it could possibly serve to theorize about it. But it shows that more substantive constraints are needed to impose a significant division among the As into the relevantly normal ones and the rest.

In discussing the example of the ravens, I already gestured at such constraints. Not just any genetic endowment is such that a raven born with it is set upon the characteristic developmental process for ravens and their color. Likewise, only some dietary regimes will keep ravens going in this process. Others will derail it. When I mentioned these, I spoke of the ordinary or normal endowments and diets, and that is what makes the account non-reductive. At least at this stage of inquiry, I appeal to intuitions about what's normal as far as the color of ravens is concerned to help pick out the associated process, and our ideas about the process to help pick out the relevantly normal ravens.

In order to evaluate this proposal, I suggest that we look towards three questions. Does it allow us to address questions about characterizing sentences that we could not answer without it, or not as well? Are there features of characterizing sentences that rule out the proposal? And: does it fit into a more general account of how inquiry and representation work? That is, we should evaluate the proposal according to its theoretical fruitfulness, its empirical adequacy, and its holistic fit within a larger theory. The next two sections address the first two points, respectively. At the end, I'll point briefly towards the direction in which I think an affirmative answer to the third question can be found.

## 4 Answering the Questions

Let's begin with the theoretical fruitfulness of the proposal, specifically with how I can answer the three questions I proposed at the end of section 2 on the hypothesis that characterizing sentences are interpreted in terms of characteristic processes. Here they are again.

- What distinguishes the merely apparent exceptions from the remaining instances of the kind?
- Why is the class of merely apparent exceptions to a characterizing sentence so heterogeneous?

- Why do only large-scale variations in the world make for variations in the truth-value of a characterizing sentence?

I can answer the first two just on the hypothesis that there are characteristic processes. The third can be answered if we assume that characteristic processes satisfy two further constraints. These constraints have theoretical utility beyond answering the third question, so my appeal to them is not *ad hoc*.

We've already seen the answer to the first question: in order for an *A* to fall within the scope of the characterizing sentence *As are F*, it needs to go through the characteristic process. This idea allows me to immediately explain the heterogeneity of the merely apparent exceptions. An exception to a characterizing sentence *As are F* is merely apparent iff it is an *A* and is neither *F* nor *F*-normal. Hence, a merely apparent exception is an *A* in which the characteristic process either never got started, was derailed along the way, or whose outcome was altered after it had run its course. But there are ever so many ways for this to happen. There is a normal way for a raven to come by its color, and there are extremely many ways to prevent this process from getting started, to derail it once underway, or to alter its outcome once complete. Each of these ways produces a different, merely apparent exception. That is why the class of merely apparent exceptions is so heterogeneous.

An appeal to such a process also allows us to say how a speaker with limited cognitive resources could represent where to draw the line between the merely apparent exceptions to a characterizing sentence and the rest. She simply has to represent the characteristic process, which only requires that she represents the (presumably quite small) sequence of events that constitute the characteristic process. The class of the merely apparent exceptions is then determined as deviations from this characteristic process.

The idea is perhaps best illustrated on analogy with a rather different example. Consider the class of things that are not books, the class of non-books. This, too, is an extremely heterogeneous class, and one might ask how a speaker with limited cognitive resources could manage to represent the extension of this class well enough to tell whether an object belongs to it or not. The answer is obvious. In the first instance, a speaker represents not the class of non-books, but rather the class of books. Once she has done that, she can also represent the complement of the class of books, the non-books. But crucially, she can represent the non-books simply as the complement of the books, rather than needing an explicit, enumerative representation of all the non-books. I'm suggesting that the same thing holds for the merely apparent exceptions,

and indeed, for the non-*F*-normal *As*. A speaker in the first instance grasps what it is to be an *F*-normal *A*, and that allows her to represent the class of non-*F*-normal *As* derivatively.

This way of answering the questions represents a complete reversal of how they are usually approached in the literature on cp-laws. There, the notion analyzed first is that of a merely apparent exception, since theorists usually try to place substantive constraints on what can go in to the unless-clause of an explication like ‘all ravens are black, unless...’ The normal ravens are then defined derivatively as whatever isn’t covered by the unless-clause. By contrast, I first try to analyze what it is to be relevantly normal, and then only derivatively what it is to be a merely apparent exception.

In order to answer the third question—why only large-scale changes suffice for a variation in truth-value—I need to introduce two constraints on characteristic processes.

[SUFFICIENCY] The characteristic process for a given characterizing sentence settles whether something that goes through it is *F*.

[MODAL ROBUSTNESS] If a process is characteristic in a world *w*, then it is also characteristic in worlds close to *w*.

SUFFICIENCY says that characteristic processes are individuated in part by their results. In the case of (2), ‘ravens are black,’ the characteristic process is instantiated by a particular raven only if that raven ends up being black. This is the feature that accounts for the fact that we cannot imagine a non-black raven that is nonetheless normal: given that the characteristic process is individuated in part by way of its endpoint, it’s impossible for a raven to go through this process, thus being relevantly normal, and be non-black. By counterposition, any non-black raven fails to instantiate this process, and therefore isn’t normal.

The deeper point made by SUFFICIENCY is that a characterizing sentence *As are F* isn’t directly about individual instances of the predicate *A*. Rather, it is about whether a certain process is characteristic for the particular *A* and *F* at issue, in this sense of *about*: once we’ve fixed on which process is characteristic in a world, we’ve settled that sentence’s truth-value in that world. There is no further variation of truth-values once the process is fixed. So the only way to vary the truth-value of a sentence is to vary which process is characteristic for the interpretation of the sentence.

SUFFICIENCY by itself is not yet enough to explain why we only see variations in the truth-value of a characterizing sentence when there are large-scale changes. For

all that SUFFICIENCY says, which process is characteristic in a world could vary with very slight changes in the distribution of properties among individuals, for example, slight changes in the distribution of blackness among ravens. Indeed, for all that SUFFICIENCY says, worlds that differ only with respect to whether a single raven is black or not could differ in whether a process whose outcome is a black raven is characteristic. That is, *given* that we accept that *As are F*, we cannot imagine a relevantly normal *A* that is not *F*. But if we only assume SUFFICIENCY, we might still be able to imagine a world that is only slightly different in which it would be false that *As are F*. That is clearly not the case.

Once we add MODAL ROBUSTNESS, this option is foreclosed. On the assumption that a change in whether a single *A* is *F* does not take one out of the neighbourhood of a world, MODAL ROBUSTNESS says that worlds that differ only in whether a single *A* is *F* or not can never make a difference to whether a process is characteristic or not. To get out of such a neighbourhood requires a large-scale change.

I said earlier that the two constrains I need to answer the third question have theoretical uses beyond just answering this question. In the case of SUFFICIENCY, we can immediately predict another feature of characterizing sentences. Consider what happens when we negate a simple characterizing sentence, as in (17).

(17) Ravens are not black.

This sentence is stronger than merely the claim that not all normal ravens are black, since this latter paraphrase is compatible with there being some normal, black ravens. Rather, I think that if we accepted (17), we would be committed to treating all black ravens as exceptions. In this respect, compare (17) with (18).

(18) Ravens are not white.

It seems to me that one of the commitments we incur in asserting (18) is that all white ravens are somehow merely apparent exceptions to the characterizing sentence and thus abnormal.<sup>15</sup> If we assume SUFFICIENCY, we can predict this fact about simple characterizing sentences immediately.

From the semantics I've suggested, it follows that if ravens aren't white, some relevantly normal ravens are not white. And by SUFFICIENCY, if some relevantly normal

---

<sup>15</sup>This fact about simple characterizing sentences has been observed before, among others by von Fintel (1997).

ravens aren't white, then the associated characteristic process is individuated in part by issuing in a raven that is not white. Hence, all such relevantly normal ravens are not white.

While SUFFICIENCY shows its utility in accounting for other semantic features of characterizing sentences, MODAL ROBUSTNESS is motivated by our interest in making these generalizations in the first place. It seems uncontroversially true that characterizing sentences play a role in guiding our actions. We contemplate various courses of action and use the generalizations to shape our expectations about the outcomes of these actions. In order to play this role, it's important that the truth of the generalization is invariant across the possibilities we are choosing between. But that is just to say that very small differences in the way the world is, the kinds of differences we can bring about by acting one way rather than another, should not make for a difference in the truth-value of the generalization. And given the rest of the semantics I have proposed here, that is just to say that MODAL ROBUSTNESS holds.

### *Further Theoretical Fruits*

So far I've focused very narrowly on how I can answer the questions I posed at the end of section 2 by appeal to characteristic processes. Here are some other features of characterizing sentences I can elucidate on the hypothesis that we interpret them in terms of processes.

A process can be interfered with many times, and some of these interferences might offset each other, so that the end result is the same as if the process had run its course without any interferences. On the view I'm proposing, if this happened to an *A*, it would not count as normal, even though it might well conform to the generalization that *As are F*. And this is precisely what we find. Take, for example, an albino raven that has been painted black. That is a black raven, and so conforms to the generalization, but intuitively, it is not relevantly normal.

If it's true that there is a significant distinction among the black ravens into ones the generalization is intuitively about and ones that are, intuitively, irrelevant to it, then this speaks very strongly in favor of a view according to which what makes something relevantly normal is a matter of how it came to be black. That the way to distinguish among the black ravens goes by way of how they came by their color certainly does not *follow* from the fact that we distinguish among them as we do. However, other likely candidates for doing the work of appropriately restricting the quantifier don't work. It is clearly not a matter of a spatial or temporal proximity to the speaker. Nor is it

otherwise a difference in temporal or spatial location, since the “normal” black ravens and the other black ravens might be sprinkled randomly. One might wish to say that some black ravens are important, or salient, while others are not. But salience is not a primitive property, things are always salient to us in virtue of some other feature. And while I am certainly not giving a general theory of salience, I am suggesting that in the very specific context of evaluating characterizing sentences, salience is a matter of the process by which, say, a raven comes by its color.

The next theoretical virtue concerns a class of sentences that intuitively should be grouped among the characterizing sentences, yet fails to conform to the paradigm I set out. These are sentences just like the paradigmatic characterizing sentences I’ve been discussing so far, except that there cannot be any exceptions to the generalization they express, such as the examples (19) and (20).

(19) Ravens are birds.

(20) Lions are mammals.

Considering their similarities, it would be extremely odd if (19) had different semantics from (2), ‘ravens are black.’ Yet there can be exceptions to (2), while there cannot be any to (19).

Theorizing about normality in terms of processes allows me to account for this difference without positing a semantic difference. In both cases, something is a relevantly normal raven only if it goes through the characteristic processes associated with the generalization. The difference between the two is a result of the metaphysics of the world, specifically, the fact that nothing can go through a process that produces a raven without also being a bird. Semantically the two sentences (2) and (19) are on a par: both are interpreted in terms of processes, and both are such that if the process was derailed in such a way as to produce a raven that did not conform to the generalization, the generalization would still be true. Metaphysically, the two differ, however, in that in the former case, but not the latter, the process really can be derailed in this way.

An analogy might make this idea clearer. Consider the two descriptions in (21).

- (21) a. The President of the US is male.
- b. The successor of 2 is odd.

The first description is not rigid: it does not pick out the same object with respect to every possible world at which it is evaluated. By contrast, the second description is rigid. It does pick out the same thing with respect to all possible worlds, viz., the number 3. But we do not infer from this difference that the two descriptions have different semantics. Rather, the difference in rigidity is due to the metaphysics of the world. It's necessarily the case that 3 is the successor of 2, and that's why 'the successor of 2' always picks out the same thing. I am suggesting that the difference between (2) and (19) is similarly due to a metaphysical difference, not a semantic one.

Thus, the introduction of processes allows me to answer all three of the questions I was left with at the end of section 2, and at the same time, to account for further features of characterizing sentences that the processes weren't originally designed to handle. The introduction of processes into the analysis of characterizing sentences is explanatorily powerful and theoretically fruitful. But is it empirically defensible?

## 5 Is the Two-Stage Analysis Empirically Adequate?

I have so far emphasized one of the shifts in the debate about cp-generalizations and characterizing sentences. Once we accept that cp-generalizations are characterizing sentences, we can make progress on questions about their semantics without having to give a reductive account of these semantics. But there is another shift. The scope of our theory becomes much wider. No longer do we need to just give a theory that's adequate to the phenomena in the non-fundamental sciences, we now need to give a theory that also works well for other areas in which we use characterizing sentences. I think that the gravest worries about my account arise when we consider how we can extend this account to other areas that are, at least *prima facie*, inhospitable to the notion of a process, characteristic or otherwise.

### *The Extension to Other Areas of Discourse*

Consider (22) and (23).

(22) Soccer games are played eleven to a side.

(23) Prime numbers are odd.

I want to sketch how my account can work in interpreting such sentences, and to show that the treatment is quite natural.

By my lights, (22) is true iff the characteristic process for creating soccer games ends in a soccer game being played that has eleven players to a team. And this may well be the case. All that is required is that the characteristic process is one in which the game is set up according to the rules that dictate that each team be composed of eleven players. This pattern will hold generally for characterizing sentences that reflect rules of games or other rule-governed activity. The characteristic process is simply one in which the game or activity is performed according to the rule.

One issue left open by this interpretation of (22) is whether it tolerates any exceptions. But this is not a semantic question on my view, but a broadly metaphysical one which turns on whether rules that are constitutive of a practice—such as soccer playing—can be broken without thereby making it the case that the people breaking the rule have stopped engaging in the practice. Less abstractly, if a team suddenly fields only ten players, maybe because only ten showed up for one of the sides, the resulting game of eleven on ten may or may not be soccer, depending on whether this departure from the rule makes the game something other than soccer. And I think it is exactly the right result that the semantics of (22) does not settle this issue.

(23) is an example from mathematics. On my view, this sentence is interpreted as making a claim about relevantly normal primes. In turn, the reference to normality needs to be spelled out in terms of processes. And while the first step in the interpretation may seem odd, the second is surely fatal. Numbers do not come to be via some process.

However, my account can naturally handle these examples. In fact, that numbers do not come to be by some process provides the crucial idea. The relevant process here simple consists of a single event, the existence of a number. If that is correct, then every number is equally normal, and hence characterizing sentences about numbers—or any mathematical subject matter—are equivalent to their corresponding universal generalizations. This accords well with our initial unhappiness at glossing (23) in terms of normal numbers, since all numbers are equally normal. Hence, to say that all *normal* primes are odd is to suggest that there are, or at least could be, *abnormal* primes and since this suggestion is false, asserting the sentence that is naturally heard as making is less than fully excellent. In essence, my account makes characterizing sentences about mathematics equivalent to their corresponding universal generalization. I think that this accords well with our truth-value judgments.

### *Characterizing Sentences and Habituals*

I want to end this section by distinguishing a class of sentences that are often included among the cp-generalizations, but where talk of processes seems out of place. They are exhibited in the examples in (24).

- (24) a. Mary smokes.  
b. Iron bars expand when heated.  
c. Dogs bark.

We can interpret (24c) as expressing a generalization over events that involve Mary, claiming that some appropriate subset of these events are also ones in which Mary smokes.<sup>16</sup> Likewise, (24b) would then be interpreted as saying that events in which iron bars are heated are also events in which these bars expand. (24c), on this view, says that a certain set of events involving dogs are also events in which they bark. What is more, each of these generalizations has to be non-strict. There are events involving Mary that are not smoking events, there are events of iron-bar-heating that are not also iron-bar expansions, and there are events involving dogs in which they do not bark.

But (24a) cannot be interpreted in terms of processes. All of the events quantified over involve a single individual, so whatever the process is by which she comes to be the way she is won't distinguish between those events in which she smokes and those in which she doesn't. Once we see the problem there, we can observe the same one in (24b). Even if we focus just on a single iron bar, we can still see a non-strict generalization: 'This iron bar expands when heated' is about a particular bar, but it is compatible with there being situations in which the bar is heated in a non-standard way, for instance, in which the bar is simultaneously compressed to counteract the effects of heating. Again, since we are dealing with a single bar, characteristic processes don't seem to apply.

But it's important to note that the examples in (24), *habituals* as they are usually called, are semantically distinct from the characterizing sentences. In the latter case, we have a generalization over objects, in the case of habituals, we have a generalization over situations. These two kinds of generalization can even occur side-by-side, as in

---

<sup>16</sup>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 (1973). For a more recent elaboration, see Fara (2001, 2005).

(24c). The bare plural subject ‘dogs’ makes the generalization non-strict with respect to dogs, since some dogs, such as those with defective vocal tracts, do not bark at all. The habitual predicate ‘barks’ makes the generalization non-strict with respect to the situation in which a relevantly normal dog actually manifests its habit or disposition to bark, excluding situations such as ones in which the dog is asleep or muzzled.

Since these two kinds of sentences are semantically distinct, involving generalizations over different kinds of things, and by being associated with different elements in the sentence, we need not assume that a single treatment is required for both characterizing sentences and habituals. In particular, both characterizing sentences and habituals may be given a normality-analysis, but it does not follow that we need to spell out normality along the same lines in both cases.

## 6 Conclusion

I’ve argued that cp-laws are generics, or more precisely, that they are characterizing sentences. I’ve suggested that some of their most puzzling features are captured by the three questions I address in this essay. The core idea of my response is that characterizing sentences are interpreted in terms of their associated characteristic processes.

One issue I have remained silent on, and which remains as a pressing question, is what makes a process characteristic. This is a question about categorization quite generally, since if I am right, that is just a question of why we focus on one class of objects in order to frame a generalization rather than another. As such, this question must be answered by giving a theory of categorization, or perhaps as it is more commonly known, a theory of natural kinds. I cannot present such a theory here, and it is important to me that the theory I’ve presented in terms of processes be able to stand on its own. A theory of natural kinds then has the job of explaining, *inter alia*, how categorization by process fits into our investigative strategies. I myself am drawn to a theory of homeostatic property clusters along the lines proposed by Boyd (1991, 1999) and Kornblith (1993). But a presentation and defense of this view has to await another occasion.

## References

- Asher, N. and M. Morreau. “What Some Generic Sentences Mean”. In G. N. Carlson and F. J. Pelletier, eds., *The Generic Book*, 300–339 (Chicago: University of Chicago Press, 1995).
- Boyd, R. N. “Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds”. *Philosophical Studies*, (61), (1991), 127–148.
- . “Kinds, Complexity and Multiple Realization: Comments on Millikan’s ”Historical Kinds and the Special Sciences””. *Philosophical Studies*, 95, (1999), 67–98.
- Cohen, A. *Think Generic!* (Stanford, CA: CSLI Publications, 1999).
- Dahl, O. “On Generics”. In E. Keenan, ed., *Formal Semantics of Natural Language*, 99–111 (Cambridge: Cambridge UP, 1975).
- Davidson, D. “The Logical Form of Action Sentences”. In *Essays on Actions and Events*, 105–148 (Oxford: Oxford UP, 1980).
- Earman, J. and J. Roberts. “*Ceteris Paribus*, There is no Problem of Provisos”. *Synthese*, (118), (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.
- Farkas, D. F. and Y. Sugioka. “Restrictive *If/When* Clauses”. *Linguistics and Philosophy*, (6), (1983), 225–258.
- 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.
- Heim, I. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst (1982).
- 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).
- Lawler, J. *Studies in English Generics* (Ann Arbor: University of Michigan Press, 1973).
- Lewis, D. K. “Adverbs of Quantification”. In E. L. Keenan, ed., *Formal Semantics of Natural Language*, 3–15 (Cambridge, UK: Cambridge UP, 1973).

- . “Events”. In *Philosophical Papers, Vol. 2*, 241–269 (New York: Oxford UP, 1986).
- Mott, P. “Fodor and Ceteris Paribus Laws”. *Mind*, 101(402), (1992), 335–346.
- Nickel, B. “Semantics for Characterizing Sentences”. In J. Gajewski, V. Hacquard, B. Nickel, and S. Yalcin, eds., *Recent Work on Modality* (Cambridge, MA: MIT Press, 2006).
- . “The Ways of Normality”.
- 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.
- Schiffer, S. “Ceteris Paribus Laws”. *Mind*, 100(1), (1991), 1–17.
- Schurz, G. “Pietroski and Rey on Ceteris Paribus Laws”. *British Journal for the Philosophy of Science*, (52), (2001), 359–370.
- . “Ceteris paribus Laws: Classification and Deconstruction”. *Erkenntnis*, 57(3), (2002), 351–372.
- Sharvy, R. “A More General Theory of Definite Descriptions”. *The Philosophical Review*, 89(4), (1980), 607–624.
- von Fintel, K. “Bare Plurals, Bare Conditionals, and *Only*”. *Journal of Semantics*, 14(1), (1997), 1–56.
- Woodward, J. “There is no such Thing as a *ceteris paribus* Law”. *Erkenntnis*, 57(3), (2002), 303–328.