On Semantics for Characterizing Sentences*

Bernhard Nickel

July 30, 2006

Abstract

The paper presents semantics for a subset of generics, so-called "characterizing sentences". It is argued that claims about the relationship between the truth of characterizing sentences and claims about the distribution of properties among individuals can be viewed independently of considerations about logical form. Some extant approaches are presented and criticized, and a positive analysis of characterizing sentences in terms of normality is introduced and defended. The main innovation is that a notion of normality enters into the analysis in two separate but connected places, not just one as competing accounts suggest.

1 Introduction

A wide range of sentences can plausibly be characterized as generalizations that tolerate exceptions. To say that they tolerate exceptions is to say that some counterexamples to the corresponding universal generalization would not be sufficient to shake our robust intuitions that the generalization is true. These sentences are usually discussed under the heading of 'generics' in the literature. In this paper, I shall focus on a subclass of generics, what I will call *characterizing generalizations*.[1] Specifically, I will discuss various proposals about the truth-conditions of these sentences, before presenting my own analysis in terms of normality. The main conclusion I shall defend is that a notion of normality enters twice over into such an analysis, not just once as most authors would have it.

* [[Acknowledgements suppressed]]

[1] The terminology is taken from Krifka et al. (1995).

I begin by delimiting the class of sentences I focus on (section 2). I then motivate the search for an answer to the question: what is the relation between the truth of a given characterizing generalization and facts about its instances (sections 3-4)? By criticizing various extant answers to this question (section 5-7), I motivate my own (section 8).

2 Characterizing Generalizations

Generics fall broadly into two natural classes. The first consists of sentences that apparently make direct reference to kinds—*genera*, hence the term 'generics'. In these, the predicate attributes a property that particular instances of the kind cannot have. Examples are (1a)-(1c).

(1) a. Diamonds are rare.

b. Bacteria are widespread.

c. Dodos are extinct.

These sentences cannot be true in virtue of the fact that the same predicate applies to particular instances of the kind: (1a) is not true in virtue of the fact that a number of diamonds, each individually, are rare.[2]

The other class of generics are so-called characterizing sentences. These are sentences that seem to have the force of generalizations over particular objects, where the predicate can sensibly apply to instances of the generalization. I'll say more about the qualification that they *seem* to have this force below when I discuss Carlson's account. For now, let me characterize the notion of a characterizing sentence. We already encounter problems here, because there is no easily agreed upon criterion that singles out all and only the characterizing sentences.[3] Instead, we should pick out a core class of cases, which we may then extend on theoretical grounds. Here is my proposal for doing so.

[Exceptions] A sentence S is a characterizing sentence if

[2] These examples are quite central to the discussion in Carlson (1977). One of the great advantages of Carlson's account is that he can give a unified account of both reference to kinds and characterizing sentences, introduced next. Before Carlson, Quine (1964) already drew attention to the fact that explicit reference to kinds cannot be understood as a straightforward generalization over its instances.

[3] For some of the problems, see the classification in Krifka et al. (1995).

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

Intuitively, EXCEPTIONS says that the characterizing sentences are the generalizations that tolerate exceptions. I choose the more cumbersome formulation, because I don't want to rule out from the get-go that characterizing sentences express what are in fact exceptionless generalizations. Indeed, many of the proposals we will see, including my own, make just this claim. On these views, a characterizing sentence expresses a suitably restricted universal generalization.

We can illustrate the generalizations at issue by considering the examples (2) and contrasting them with (3).

(2) a. Ravens are black.
  b. Snow is white.
  c. Grass is green.

(3) a. Mary smokes.
  b. Sugar dissolves in water.
  c. Fire burns human skin.

I have separated these examples into two groups for a theoretical (and hence a defeasible) reason. For the examples in (2), it is easy to say what kind of generalization is involved. (2a) involves a generalization over ravens, (2b) over bits of snow, (2c) over blades of grass.[4] Their instances are particular ravens, bits of snow, and blades of grass, respectively.

It may be somewhat less obvious what kind of generalization the examples in (3) involve. *Prima facie*, (3a) ascribes a property to a person, the property of smoking. Likewise, (3b) seems to ascribe the same property to all bits of sugar, the property of dissolving in water. The same holds, *mutatis mutandis*, for (3c). On the face of it, they do not involve exception-tolerating generalizations. However, we can hypothesize that (3a) in fact expresses a generalization. In this case, it expresses a generalization over situations or events, situations in

[4]In speaking at this level of abstraction, I am gliding over what may be a crucial issue. I have nothing to say about the connection between the mass noun 'snow' and various bits of snow, or how large the bits have to be to count as a bit, and so on.

which Mary may or may not smoke. And (3a) says that in certain situations, she smokes. Likewise for (3b) and (3c).[5]

We can call the sentences in (2) *characteristic generalizations*, or CGs for short, and the proposition expressed by a CG a characterizing proposition or CP. By contrast, we can call the sentences in (3) ascriptions of dispositions or habits, or *habituals* for short. So characterizing generalizations are a subset of the characterizing sentences. In this paper, I focus primarily on CGs. Perhaps the proposals I discuss here can be extended in a natural way to habituals, but that is an issue on which I will remain silent.

One final remark before we begin. It seems as if characterizing sentences are context-dependent in the following respect. Roughly, they depend on what explanatory or predictive aims we have. For example, if we are doing taxonomy, (4) seems acceptable.

(4) Dobermans have floppy ears.

But when we are engaged in dog-breeding, (4) seems bad. In other work, I give some theoretical reasons for taking these intuitions at face-value as tracking semantic facts.[6] For now, I just want to mark the indisputable context-dependence of our acceptability judgments for characterizing sentences. It'll be important to keep this dependence in mind, because many of my arguments turn on acceptability judgments, and these arguments thus require keeping the predictive or explanatory background the same. As I go along, I will remind the reader of these facts.

3 Supervenience

It seems clear that the truth of a given CG rules out certain distributions of properties among its instances. For example, the present distribution of blackness among ravens, perhaps throughout history and across all possible worlds, makes it so that 'ravens are black' is true and 'ravens are green' false. That

[5]This is an idea that traces back to Davidson's discussion of action sentences, Davidson (1980), and from there has become a central plank in Davidsonian semantics. One of the early applications of this idea to generics is Lewis (1973). It is also discussed, with copious references, in Krifka et al. (1995).

[6]See my "CP-Laws, Genericity, and Processes."

suggests a thesis: once we have fixed all the facts about how properties are distributed among individuals, throughout history and across possible worlds, we have fixed all the facts about characterizing sentences, as well. In slightly technical terminology, the characterizing facts *supervene* on the individual facts.

What I have said here is relatively weak. For all this broad supervenience claim says, changing one molecule in outer space could make it so that it is no longer the case that ravens are black. That is unnecessarily weak. The facts relevant to fixing the truth of a given CP form a narrower class. In this paper I ask just what that class is. To introduce some terminology, call the class of facts that fix the truth of a given characterizing proposition its *supervenience base*. So I will ask what the supervenience base of a given CP is. The various proposals I will discuss here are all addressed to this question.

Answers to this question are independent of answers to various questions in compositional semantics, particularly questions about the LF of characterizing sentences. To say what I mean by that, let me briefly discuss Carlson's influential proposal about how to interpret characterizing sentences. I'll point out that even if Carlson's proposal (or something like his proposal) is right, the question I am interested in can still be fruitfully asked and, hopefully, answered.

3.1 Carlson

Carlson (1977) takes seriously the idea that with predicates such as 'extinct', 'widespread', or 'common', we say something directly of kinds. For this reason, he introduces kinds into the ontology. Aside from kinds, there are objects, such as the usual tables, chairs, and trees.[7] Bare plurals, such as 'ravens' are proper names of kinds, so a sentence like 'ravens are widespread' means that the kind raven has the property of being widespread.

If we try to flatfootedly extend this idea to other predicates, such as 'black', we run into a problem. The easiest extension has it that 'ravens are black' means that the kind raven has the property of being black. But that is a bad interpretation. A kind is not the sort of thing that is black—it's not the sort of thing to reflect light, for one. To deal with this problem, Carlson asks us to distinguish between two predicates, 'black' as it applies to kinds and 'black' as it applies to individual objects. Call them 'black_K' and 'black_I', respectively.[8]

[7] See Carlson (1977, pp. 116ff), as well as Carlson (1982).

[8] Of course, the relationship between 'black_K' and 'black_I' is not completely random, as is

On his view, 'ravens are black' should be interpreted not quantificationally, but as predicating the property black_K of the kind raven.

It is crucial to be clear about what Carlson is doing at this point. Some of the facts linguistic semantics seeks to account for concern the availability of distinct readings of a sentence or the possibility of anaphoric relations. The strategy is to posit a representation—a *logical form* or LF—that serves as the input to an interpretation function. Constraints on possible readings and anaphoric relations are then put in terms of these logical forms.[9]

For example, VP-ellipsis is sensitive to facts about the logical form of a sentence. Thus, the sentence (6) is only two ways ambiguous.

(6) I saw the man with the binoculars, and Joan did, too.

Either the man Joan and I saw had binoculars, or both Joan and I used binoculars to see the man, but the sentence cannot mean that I used binoculars, and that Joan saw the man who had binoculars. The fact that the sentence cannot mean this is to be explained by constraints the language puts on interpretation. After all, it's not as if a situation in which I use binoculars and Joan sees a man who has binoculars is impossible. Instead, there are constraints on the LFs we may or may not associate with (6).[10]

Similarly, we appeal to logical form to explain the presence of different readings of a sentence, as well as what the different readings mean. Thus, we explain why

(7) Mary wants to meet a policeman.

---

the case between the different meanings of ambiguous terms like 'bank', where knowing one of the meanings does not help one predict the other meaning. In the case of predicates that apply either to kinds or individuals, it seems as if we can predict one of the meanings given the other. To mark this fact, Carlson introduces an operator G. The operator applies to a predicate that we can only felicitously apply to objects, such as 'black', and yields a predicate we can felicitously apply only to kinds. So on Carlson's view, the logical form of (5a) is (5b):

(5) a. Ravens are black.
b.
ravens
G black

I pass over the operator G in the text, since nothing I say depends on it.

[9] For one implementation of this idea, see Heim and Kratzer (1998).

[10] For a theory that implements this idea, see Sag (1976).

is ambiguous between a reading on which there is some particular policeman she wants to meet, and one on which there is not, by positing two distinct logical forms that this sentence can have. It is then up to the rules that interpret the logical form to deliver the intuitively present readings, and only those.

As a final example, the interpretive possibilities for some pronouns depend on features of the logical form. The sentence

(8) Mary wants to meet a policeman, and John wants to meet him, too.

has only one reading, the one on which there is some policeman that Mary and John both want to meet. The sentence cannot mean that Mary wants to meet some policeman or other, it does not matter which, and that John wants to meet some policeman, where again it does not matter which.

In making his proposal about how to interpret 'ravens are black', Carlson is offering a proposal about the LF of the sentence, specifically that the LF of (9a) is (9b).

(9) a. Ravens are black.
    b.
        ravens black_K

Positing this logical form allows him to account for a number of facts. For example, consider the contrast between (10) and (11).

(10) Mary wants to meet every policeman.

(11) Mary wants to meet policemen.

Intuitively, (10) has two readings. We capture this fact by associating two LFs with (10), one where the quantifier appears in the same place as it does in the sentence, another where the quantifier is moved to a higher position in the tree. But (11) does not have two readings. If 'policemen' were a quantified phrase in the LF, that would be inexplicable (given various other assumptions he takes to be common ground in the debate). On the other hand, if ‘policemen’ just denotes an object, we do not predict any scope effects.[11]

Carlson’s proposal is thus, in the first place, about the LF of characterizing sentences, and he suggests that such an LF only contains a reference to the kind and a reference to the property predicated of that kind. At this stage, one might ask whether this LF ought to be interpreted further, that is, what sort of output our interpretation function should give. In particular, one might ask what the connection is between true ascriptions of ‘black_K’ and ‘black_I’.

In response, Carlson says that semantically, sentences attributing blackness_I to particular ravens and those attributing blackness_K to kinds are “equally basic.”[12] He draws a parallel with predicates such as ‘is a spy.’ Someone is a spy if and only if she has certain other features, as well. But it is not part of interpreting a logical form containing ‘spy’ to say what those features are. The logical form is fully interpreted once we have that ‘Lucy is a spy’ is true iff Lucy is a spy. Giving necessary and sufficient conditions for something to be a spy is not part of giving the semantics of ‘is a spy.’[13]

I do not want to enter into a dispute about what is and what is not a proper concern for semantics. I just want to note that, even if Carlson is right, we can sensibly ask what it takes for something to be a spy. And likewise, we can sensibly ask what it takes for a given characterizing generalization to be true. Thus, we can address the question of the supervenience base even if Carlson is right about the logical form of characterizing sentences, and even if he is right about what is the concern of semantics proper. If he is wrong about either of these, the question becomes pressing for semanticists. If, for example, he is wrong about the logical form of characterizing sentences, and they contain some quantificational element that captures their characterizing force, it is crucial that we know what the interpretation function needs to produce as output when given a CG as input.[14] But either way, the question of the supervenience base is one we can address without prejudging any issues regarding the logical form of characterizing generalizations.

[11] Carlson (1977, chp. 2) is devoted to listing the data and showing that they cannot be accounted for by positing a quantifier in the logical form. Carlson (1977, chp. 4-5) is devoted to showing how these data are accounted for on his analysis.

[12] Carlson (1977, p. 109).

[13] See Carlson (1977, pp. 107ff).

[14] For reasons to think that Carlson is wrong on this point, see Diesing (1992, chp. 2) and Krifka et al. (1995, pp. 116-27).

4 Truth-Conditions for CGs?

One might challenge the idea that the truth of CPs supervenes on individual facts on other grounds. So far, I have simply assumed that characterizing generalizations are true or false. Otherwise, there could be no question of which distribution of properties the truth or falsity of characterizing propositions supervenes on. Though this is controversial, I'll argue now that the main alternative to this view fails. The alternative is to propose an analysis of the meaning of generics in terms of the conceptual role they play in a thinker's mental life. One way of implementing this idea is in terms of update semantics. Instead of identifying the meaning of a sentence with its truth-conditions, one identifies it with its context-change potential: how an agent has to update her mental state if she accepts it.[15] More specifically, the view is this.

[INFERENCE TICKET (IT)] A generic is acceptable to a thinker if and only if she is willing to use that generic in default reasoning.

The kind of default reasoning at issue is exemplified in inferences of this form.

$$\frac{\begin{array}{l} A\text{s are } B\text{s.} \\ \text{This is an } A. \end{array}}{\therefore \text{ This is a } B.}$$

An instance of this reasoning is: Ravens are black; Ray is a raven; therefore, Ray is black. According to IT, a speaker accepts a generic just in case she is willing to deploy it in this kind of inference as the first premise. I call this kind of inference default reasoning, since the inference clearly is not valid. The fact that there are exceptions to the generalization allows that sometimes, the premises are true but the conclusion false. Nonetheless, in the absence of any particular information about the specific instance referred to in the second premise (Ray), the inference is reasonable.

IT is clearly falsified by generics that are true, even though most of their instances are exceptions to the generalization. Thus (12):

(12) Mosquitos carry plasmodia.[16]

[15] For examples of analyses of generics as updates, see Geurts (1985) and Veltman (1996). Another idea is presented in Lange (1995) and Lange (2004). All of these views are committed to the claim IT I argue against in the text.

[16] Plasmodia are the parasites responsible for malaria.

As a matter of fact, most mosquitos do not carry plasmodia, and that means that in general, one does not rely on (12) in default reasoning. Nonetheless, (12) is acceptable.[17] So one direction of IT fails: a generic may be acceptable to a speaker, even though the speaker is unwilling to draw the relevant default inference. The same example shows that the other direction fails as well. Since most mosquitos do not carry plasmodia, we tend to infer that Mo the mosquito does not carry them, given only the information that Mo is a mosquito. Nonetheless, 'Mosquitos don't carry plasmodia' is not acceptable.[18]

This is not to say that no non-truth-conditional account could succeed. But the argument hopefully shows that assuming truth-conditions for CGs is well-motivated, since the main rival to such a treatment does not seem very promising.

5 Some Initial Accounts

As a first hypothesis about the supervenience base of a CP, we might consider just facts about the distribution of the property at issue among individuals in the actual world. Here is an initial proposal. A characterizing sentence is true iff all of its instances have the property predicated. Thus, (13a) has the truth-conditions (13b).

(13) a. Ravens are black.
     b. All ravens are black.

This proposal fails, because (13a) is true in the actual world, while (13b) is false. Indeed, it fails in both directions. (13) shows that a CG does not entail its corresponding universal generalization. A universal generalization does not entail its corresponding CG, either. Suppose all ravens were painted white. Even in these circumstances, it would not be true to say that ravens are white.[19]

[17] The example appears, for example, in Carlson (1977).

[18] Indeed, a truth-conditional account of the meaning of generics is probably the most widely held view. For some other reasons to prefer such a view to IT, see also Asher and Morreau (1995). They urge that a non-truth-conditional account of generics does poorly at explaining the fact that generics can be embedded in larger constructions.

[19] This is one of the places where our explanatory aims matter. I am assuming that we are in a broadly biological context. If we consider instead a situation in which we need to recognize ravens in circumscribed circumstances, we might truly say that ravens are white, such as when we're hunting ravens in a neighbourhood where ravens are habitually and meticulously painted white.

One might try the opposite extreme: a characterizing sentence is true iff at least one of its instances has the property predicated. On this view, characterizing sentences are just existential generalizations, and (14a) has the truth-conditions (14b).

(14) a. Ravens are black.

b. Some raven is black.

While (14a) and (14b) are both true, this proposal fails, as well. It predicts that (15) is true, since some ravens are white.

(15) Ravens are white.

So this proposal will not do. Indeed, it fails in both directions. (15) shows that an existential claim does not entail its corresponding CG. Likewise, a CG does not even entail the corresponding existential claim. (14a) does not entail (14b). Suppose that all ravens have been painted white. In that case, we would think that (14a) is true, while (14b) fails.

I discuss these two proposals not because they are *prima facie* compelling. Rather, they serve to motivate the hypothesis that the supervenience base of a characterizing sentence does not just consist of facts about the distribution of the property among its instances in the actual world at the time of evaluation. The argument is straightforward. A given CG neither entails nor is entailed by its corresponding universal generalization, and a given CG neither entails nor is entailed by its corresponding existential generalization. But all quantifiers that are restricted to the distribution of the property in question in the actual world will entail or be entailed by at least one of the existential or the universal generalization. Hence, the entailment facts predicted by any such quantifier will not be the same as the entailment facts actually observed for a CG. Hence, more than just the distribution of the property in question among objects in the actual world must be in the supervenience base.

It is this observation that forges the link between the interpretation of characterizing sentences and modality. Nonetheless, there is one phenomenon that we need to account for, even once we complicate our semantics by adding a

Ravens are black.
Ray is a raven.
∴ Ray is black.

modal element. There is clearly something right about the idea that CGs can serve as premises in arguments like this, what I earlier called default inferences. If the CG 'Ravens are black' expresses some kind of restricted universal quantification, this argument is enthymematic: the suppressed premise states that Ray belongs to the relevant subclass of ravens. Thus filled out, the argument is standardly valid. Intuitively, something like this suppressed premise is what we assume (reason: suppose that assumption fails. In that case, we no longer trust the inference, as the case of the mosquitos shows). To ensure that this holds, the quantificational force of the CG needs to be downward entailing in its restrictor, and the universal quantifier is a good candidate for that. So we have three elements that need to be combined. Universal quantification, in order to account for the compellingness of the default inference; restriction of the universal quantification in order to account for the toleration of exceptions; and some kind of modal element. The question is how to combine these elements.

6 Normal Worlds Quantification

A proposal that has been defended by several theorists has it that characterizing generalizations quantify over relevantly normal worlds. When a CG is biological, the worlds at issue are biologically normal worlds; when the CG is meteorological ('snow is white'), the worlds are meteorologically normal; and so on.[20]

On such a view, (16a) should be interpreted as (16b).

(16) a. Ravens are black.

b. In all the most normal biological worlds, all ravens are black.

This account seems unsatisfactory in two respects. First, there is the often noted problem that a world in which all ravens are black does not seem to be a biologically normal world. Mutation and albinism are perfectly natural

[20]See for example the discussion in Dahl (1975) as well as Krifka et al. (1995, pp. 49-57). A recent defender of this position in the philosophy of science literature is Silverberg (1996).

phenomena, for example. To make this point even more forcefully, we would like to say that (17a) is true, which would be glossed as (17b).

(17) a. Turtles are long-lived.

b. In all the most normal biological worlds, all turtles are long-lived.

But all turtles are long-lived only where none fall prey to predators or accident. And a world without predators is not biologically normal, since a balance between prey and predator populations seems to be more normal, not just more usual.

Second, it is unclear how to accommodate the acceptability of the inference about Ray the raven. For one, it is quite possible that the individuals that exist in the most normal biological worlds are not the same as exist in the actual world, so that Ray may not even fall within the scope of the set of ravens quantified over. For two, it makes it mysterious why in general it is acceptable to rely on the enthymematic argument, since the argument only goes through if the actual world is biologically normal. And considering how odd the biologically normal worlds are, the assumption that the actual world is biologically normal seems to never be warranted.

7 Cohen: Probability

7.1 Introduction and Preliminaries

Cohen (1999) suggests that we analyze CGs as asserting a claim about probabilities.[21] On Cohen's view, *As are Bs* asserts that for a randomly picked object, the probability that it is B given that it is A is better than even chance.[22] So for example,

(18) 'ravens are black' is true iff the probability that x is black given that x is a raven is greater than .5.

[21] See also Cohen (2004).

[22] Cohen puts the probability in terms of properties, but that strictly speaking doesn't make sense. Probabilities are defined over propositions, not properties.

As Cohen says, his proposal stands and falls with how we interpret the claim on the right hand side, i.e., what it takes for the relevant probability to be high. But as I argue now, nothing turns on his appeal to probabilities. The first thing to do is eliminate talk of them.

Most abstractly, Cohen is a *frequentist* about probabilities. He identifies the probability of some event with the frequency of that event among some suitable reference class. The idea behind this proposal is simple. When we begin to think about probabilities, we think about games of chance, such as coin tosses. And when we intuitively evaluate claims about probability, we "count possibilities."

On a given coin toss, there are two possible outcomes, heads or tails. Call them H and T. When we want to evaluate the probability of a proposition, we count the number of outcomes that would make it true among all possible ones. So to evaluate the proposition that heads comes up on the next toss, we count the possibilities that would verify it. That is one out of the two possible ones, and hence the probability is .5. If we want to evaluate the proposition that on two sequential tosses, heads comes up at least once, we count the total possible outcomes (4: *HT, HH, TH, TT*), and count the outcomes that would verify the proposition (3: *HT, HH, TH*), and hence assign probability .75 to that proposition.

I have so far made an assumption, and making the assumption explicit shows a shortcoming of the model. I have so far assumed that all of the possibilities are equally likely. Suppose that's not the case. For definiteness, suppose that the coin we are considering is heavily biased towards heads, so that heads is twice as likely as tails. In that case, there are still only two outcomes, H and T, but we won't get the right probabilities just by counting favorable outcomes among all the possible ones. Rather, the thought is that the different outcomes have to be represented in our set of possibilities according to their weights. Out of any three tosses, two will come up heads, yielding the right probabilities. We are still counting favorable outcomes among the possible ones, but now we are considering a set of outcomes that reflects the weighted probabilities.

Frequentists say that *that* is all there is to probability: the probability of a proposition is to be identified with the frequency of favorable outcomes among some suitable set of outcomes. Turning back to Cohen's proposal, (18) is thus equivalent to

(19) ‘ravens are black’ is true iff the frequency of black ravens among some class of ravens R is greater than .5.

Or in other words, ‘ravens are black’ is true iff most ravens in some suitable domain are. As Cohen acknowledges, he faces what in the philosophical literature has come to be known as the problem of the reference class. To see the problem, return to the example of the coin toss.

We would make a hash of things if we identified the probability of heads with the frequency of heads among any set of three tosses, such as the three most recent ones. For one, it’s true that the biased coin has a $\frac{2}{3}$ chance of coming up heads even before it has been tossed. But what is worse, sometimes three tosses in a row will land tails. Sometimes, three tosses in a row land heads. And in general, any combination of heads and tails is possible, but only some of them reflect the “true” probabilities. The problem of the reference class is to say how to determine the proper reference class, without taking for granted what the right probabilities are.

To put the same point in terms of ravens, we have already seen that a universal generalization does not entail its corresponding characterizing generalization. Recall a situation in which all the ravens are painted white. If we took that as our reference class R, Cohen would predict that ‘ravens are white’ is true, and that ‘ravens are black’ is false. But that is false, and hence the set of ravens that exist at the world at that time cannot be the proper reference class R in order to evaluate ‘ravens are black’. That means that some classes cannot be the proper reference class for a given CG. Put in Cohen’s terminology, the question about the supervenience base I am addressing is just the question of what the class R is that we need to use to evaluate a given CP. Probabilities have nothing to do with it.

7.2 Criticism

Cohen’s main substantive proposal is that the reference class be drawn from a suitable stretch of the history of the world with respect to which we are evaluating the CG. Given the generalization *As are Bs*, we should consider all the As that exist during the relevant stretch. If most of them are Bs, the generalization is true.

The proposal fails from the outset. To see why, let me look at its motivation.

Consider again ‘ravens are black.’ Situations in which most ravens are not seem like aberrations, and these aberrations will be swamped if we take a sufficiently long view of the matter. In the long run, most ravens are black, even if there are intermittent hiccups. But that pretty immediately suggests a systematic class of counterexamples to Cohen’s proposal: generalizations where the fact that most instances do not conform is not an aberration. We have already seen one example, the claim that mosquitos carry Plasmodia. Through no stretch of history is it the case that most mosquitos do, but the generalization is true nonetheless. Consider also (20).

(20) Turtles are long-lived.

Most turtles die at the hands of predators shortly after they hatch, so most turtles are not long-lived at all. And the presence of predators is a stable fact about turtle environments, so that even in the long run, most turtles aren’t long-lived, either. Once we see the pattern, the examples are easy to multiply.

8 The Preferred Account

I’ll now turn to my own proposal. As I said, there are three pieces that need to be put together: universal quantification, a restriction on the scope of that quantification, and a model element. I’ll argue that we should rely on the notion of normality. That in itself is not new.[23] Where I part company with other theorists is in how these parts interact. The proposal I present here says: suppose we’re given a CG *As are Bs*. The domain at issue is one that is relevantly representative for this CG. And further, the CG is true just in case all *A*s that are normal in the relevant respect are *B*s. And crucially, the representativeness of the domain needs to be spelled out using the same notion of normality as is used in the restrictor of the quantifier.

[23]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), Lawler (1973). However, the particular proposals differ on the role of references to normality.

8.1 Preliminary

I will assume that characterizing generalizations are presented in the following form: *As are Bs* is true iff

(21) $[\textsc{gen}(x) : A(x)](B(x))$.

(21) says that a CG is true iff some kind of quantification over As is satisfied. From here on in, I will assume that a CG is presented in this form of restricted quantification. There are two reasons for this.

First, the class of characterizing sentences is syntactically complex: it includes some but not all sentences containing bare plurals, some but not all conditionals, and perhaps others. For example, on one reading,

(22) Dogs are on my lawn.

is not a characterizing sentence. It simply expresses that there are some dogs on my lawn.

What is more, it is sometimes a nontrivial task to say how to represent a sentence of natural language. Consider (23).[24]

(23) Typhoons arise in this part of the Pacific.

This can mean two things. Either that it is characteristic of this part of the Pacific that typhoons arise there (although they arise elsewhere, as well), or that it is characteristic of typhoons that they arise in this part of the Pacific (this is where typhoons come from). The two readings might be represented as (24a) and (24b), respectively.

(24) a. $\Big[\textsc{gen}(x) : \text{ is this part of the Pacific}(x)\Big]\Big((\exists y)(\text{Typhoons}(y) \wedge y \text{ arise in } x)\Big)$.

b. $\Big[\textsc{gen}(x) : \text{ Typhoons}(x)\Big]\Big((\exists y)(\text{is this part of the Pacific}(y) \wedge x \text{ arise in } y)\Big)$.

In this paper, I want to avoid questions about how to map sentences in English into the quasi-formal representation (21). As I said, I am not taking any stands

[24]The example is taken from Krifka et al. (1995, pp. 24ff), who credit it to Milsark (1974).

on the LF of characterizing generalizations, and so I just assume that some suitable hypothesis can be found that describes a function from natural language sentences (and perhaps other features of the context, such as intonation or background knowledge) to this quasi-formal representation.

I'll now motivate the two aspects of the proposal in reverse order: the particular form of restricted universal quantification I favor, then the reference to the suitable domain.

8.2 Universal Quantification over Normal Instances

The restriction on universal quantification is a restriction to instances that are normal in some respect. That is to say, sentences whose truth-conditions are initially represented as (25a) should be interpreted as (25b).

(25) a. $[\text{GEN}(x) : A(x)](B(x))$.

b. $[(\forall x) : x$ is a normal A with respect to the property ϕ of As$](B(x))$.

There are at least three reasons for adopting this view. The first is that it accounts for the entailment relations characterizing sentences enter into, or rather, fail to enter into. Making this point allows me to show that we need to take the restriction to be normality in some respect. We need this in order to maintain that (26b) can be false even when (26a) is true.

(26) a. Chickens lay eggs.

b. Chickens are hens.

If we interpreted (26a) as meaning simply that all normal chickens lay eggs, and (26b) that all normal chickens are hens, the former would entail the latter. Since only hens lay eggs, the truth of (26a) entails that all normal chickens are hens, and hence (26b) must be true as well. These kinds of examples motivate quantifying not just over all normal instances, but over all instances that are normal with respect to ϕ. Since the respects of normality differ in the two sentences, they do not entail each other.

The other fact about entailment relations I can capture is the failure of downward entailment.[25] Even if the Cs are a subset of the As, the inference from *As are Bs* to *Cs are Bs* is generally not valid. Thus

(27) a. Snow is white.
$\not\Rightarrow$ Dirty snow is white.

b. English verbs form their simple past tense by adding '-ed'.
$\not\Rightarrow$ Irregular English verbs form their simple past tense by adding '-ed'.

Snow that is normal with respect to its color is white. But dirty snow that is normal with respect to its color is not white, even though all dirty snow is snow. So we can account for the failure of downward entailment. Hence, my proposal accounts for the entailment relations.

The second reason to accept this proposal is that it makes intelligible why the meaning of a characterizing sentence is changed only slightly by prefixing 'normally' to it. However, I want to emphasize that the occurrence of 'normal' in the specification of the truth-conditions of characterizing sentences does not completely coincide with our everyday usage of the word, which is why the meaning is changed at all. One way to see this is that our ordinary understanding of normality is not happily wed to the context-dependence of characterizing sentences. We don't think, or at least are unhappy to say, that in a sense dobermans with floppy ears are normal, and in a sense they aren't. It also means that we need to do more work in spelling out just what normality amounts to for various characterizing sentences.

Third, it accounts for facts of outright incompatibility. It seems as if (28a) and (29a) are incompatible, which are represented as (28b) and (29b), respectively.

(28) a. Ravens are black.

b. All ravens that are normal with respect to the color of ravens are black.

[25] We capture the fact that characterizing sentences are not upward entailing just by noting that they are universal generalizations, which in general are not upward entailing.

(29) a. Ravens are green.

b. All ravens that are normal with respect to the color of ravens are green.

The key idea is that what is a raven that is normal with respect to the color of ravens is independent of the predicate in (28a) and (29a). Thus, the restrictor of the universal quantifier is the same in the interpretation of these two sentences. And thus we straightforwardly predict that if all of the things in the restrictor are black, it cannot be the case that all of the things in that restrictor are green, as well.

Let me make two clarificatory points about this stage of the analysis. In the statement (25b), ϕ is a determinable of the property predicated by 'B.' A determinable of B is some property C such that necessarily, all Bs are Cs, but possibly, some Cs are not Bs. Red is a determinable of scarlet, for example. If 'B' predicates a color, then we are concerned with normality with respect to being colored. If 'B' predicates long life, we are concerned with normality with respect to how long one lives. I will not speculate on how ϕ is determined for a particular sentence. Note that it will not do to simply require ϕ to be a determinable of B. ϕ has to be more tightly constrained than just being a determinable, since for any given B, there are multiple determinables, and what counts as normal with respect to one of them may well differ from what does for another.

8.3 The Representative Domain

In the first instance, representativeness arises as an issue because of what would otherwise be vacuous quantification. For example, we take (30) to be true, even when the only lions left have three legs.

(30) Lions have four legs.

The problem for my kind of account arises from the observation that when no lions that are normal with respect to their number of legs exist in the world, all sentences of the form (31a) are true.

(31) a. Lions have n legs.

b. All lions that are normal with respect to the number of legs of lions have n legs.

That is because, for all I have said so far, (31a) is interpreted as (31b), which in the present circumstances would be a case of vacuous quantification. Hence, all sentences of the form (31b) are true, and thus also (31a).

More needs to be said, and a flat-footed proposal would have it that we add a simple counterfactual element. For example, we could interpret (30) as (32).

(32) If there was a lion that was normal with respect to the number of legs for lions, then all lions that are normal with respect to the number of legs for lions would have 4 legs.

Put more simply, (30) says more or less that in the closest world in which there are normal lions, the normal ones have four legs. That works well for cases in which there is just one way of being normal. But it fails when we consider examples for which there is more than one way of being normal.

Take cats. Certain colors and patterns are normal for cats, while others are not. To simplify, suppose that the only colors normal for cats are white and black. Then (33) is true.

(33) Cats are black or white.

Now suppose that at a time, we're missing white cats. There are still black ones. In this case, the flat-footed proposal makes the prediction that (34) is true.

(34) Cats are black (and nothing else).

The flat-footed proposal makes this prediction, because there are in fact normal cats—the black ones—and all normal cats are black. But that is clearly the wrong prediction.[26]

The diagnosis of the failure of the flat-footed proposal is that there are lots of ways of being normal, and we need to make sure that the domain over which

[26] Again, take care to remain in a biological context, rather than thinking about what you'd tell animal control.

we are quantifying has at least one exemplar for each way of being normal. So a better proposal for the truth-conditions of (33) is (35).

(35) If the following were true: for each way of being a normal cat with respect to the color of cats, there is a cat that is normal in that way, then this would also be true: all cats that are normal with respect to the color of cats are black or white.

This proposal does a better job of capturing the requirement that the domain over which we are quantifying is suitably representative. The proposal also does well with another major datum.

I said earlier that the following inference is intuitively good, but enthymematic. Ravens are black; Ray is a raven; therefore, Ray is black. Once we add the premise that Ray is normal (with respect to the color of ravens), the inference is valid. Assume that Ray is a raven. Presumably, that fact remains true under the counterfactual supposition that there are exemplars for every way of having a normal color. So in the world we are considering counterfactually, it is still the case that Ray is a raven. Given that he is also normally colored, he falls within the scope of the universal generalization. And the truth of the universal generalization ensures that, in that case, Ray is also black.

8.4 Asher and Morreau: Counterfactuals

I'll end with a discussion of another proposal, due to Asher and Morreau (1995), that combines the three elements (universal quantification, restriction on the quantification, and modality) differently from mine. The main point I want to make is that it fails because it does not deal properly with the problem of the representative domain. I'll begin by presenting their account, which they introduce with (36).

(36) 'Potatoes contain vitamin C' is true iff

any object would contain vitamin C, if it were a potato and all other things were to hold which would normally hold if it were a potato.[27]

[27] Asher and Morreau (1995, p. 310)

It is actually not clear to me how these truth-conditions work. Consider a potato that does not contain any vitamin C, perhaps because it's been boiled for a few hours. Now instantiate the truth-conditions just quoted to this case: this potato would contain vitamin C if it were a potato and all other things were to hold which would normally hold if it were a potato. It seems to me as if the following holds. Since the boiled potato is a potato, the actual world is the closest world in which the potato is a potato. Indeed, the actual world is the closest world that satisfies the antecedent of the counterfactual *if it were a potato and all other things were to hold which would normally hold if it were a potato.* But the consequent of the counterfactual is false in the actual world, since this potato does not actually contain any vitamin C. So the counterfactual is false. So the boiled potato falsifies the generic sentence 'potatoes contain vitamin C' according to these truth-conditions. That is obviously the wrong result. Boiled potatoes are paradigms of the kind of exception tolerated by 'potatoes contain vitamin C.'

One might object that 'potatoes contain vitamin C' is not a CG, but rather an ascription of a disposition or a habitual. So perhaps the account of Asher and Morreau (1995) works for CGs as stated in (36). However, parallel considerations apply when we consider our example 'ravens are black'. I take it that their truth-conditions would be (37).

(37) 'ravens are black' is true iff

any object would be black, if it were a raven and all other things were to hold which would normally hold if it were a raven.

Consider a raven painted white. Since that's a raven, the closest world that satisfies the antecedent is the actual world, but in the actual world, the raven is white, not black. Hence, the counterfactual is false, and hence 'ravens are black' is predicted false by (37), which is the wrong result.

With this in mind, I'm going to slightly reinterpret Asher and Morreau. I will interpret their proposal as (38).

(38) 'ravens are black' is true iff

any object would be black, if it were a *normal* raven and all other things were to hold which would normally hold if it were a *normal* raven.

(38) differs from (37) only in the additional occurrences of ‘normal’. Taken this way, the raven painted white no longer threatens the truth of ‘ravens are black’. Here’s why. The painted raven is not a normal raven. The closest world that satisfies the antecedent of the counterfactual is a world that is just like the actual world except that the raven is not painted, and thus black. So that world satisfies the consequent, and hence the counterfactual is true. So the raven painted white does not threaten the truth of ‘ravens are black’, as desired.

So (38) is the proposal I shall criticize.

8.5 Criticism

This proposal falls prey to the example about cats I discussed in the previous section. Recall that there are different ways for cats to have a normal color. Sometimes, not all of these ways are exemplified. In such a situation, the proposal of Asher and Morreau makes the wrong prediction. Here’s why.

The sentence that will be predicted true when it is intuitively false is (34), repeated here.

(34) Cats are black (and nothing else).

Suppose again, just for concreteness, that the only cats around are black. On the proposal at issue, (34) is interpreted as (39).

(39) Any object would be black if it were a normal cat and all other things were to hold which would normally hold if it were a normal cat.

And in the situation imagined, (39) is true. The only things that even could be normal cats are cats. And it is true of all the cats around that they would be black, since they in fact are. So (39) is true. And that is the wrong result.

Quite generally, even if there are abnormal instances of a generalization that could be normal, there is no reason to believe that were they normal, there would be at least one exemplar for every way of being normal in the relevant respect.

9 And what about Supervenience?

The account I have given here says that normality enters twice over into the claim a given CG makes. Not only does it serve to restrict a universal quantifier, it is also required in order to determine the domain of that quantifier. So the crucial question for a theory of the supervenience base for characterizing sentences is what determines normality for a given choice of A and ϕ.

Our discussion has shown that normality is not a statistical notion, not even over a sufficiently lengthy course of history. It is not a notion that attaches to worlds. And perhaps the fact that at least some characterizing sentences are sensitive to our explanatory and predictive aims offers the hint that normality is determined by these aims. But it is unclear how much more we can say at this level of abstraction. At some point, a theory of characterizing sentences will turn into a theory of whatever the intellectual home of the relevant sentences may be.[28]

[28] I try to make some progress on these issues in my "CP-Laws, Genericity, and Processes."

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).

Carlson, G. N. *Reference to Kinds in English*. Ph.D. thesis, University of Massachusetts, Amherst (1977).

—. "Generic Terms and Generic Sentences". *Journal of Philosophical Logic*, 11, (1982), 145–181.

Cohen, A. "Generics, Frequency Adverbs, and Probability". *Linguistics and Philosophy*, 22, (1999), 221–253.

—. "Generics and Mental Representation". *Linguistics and Philosophy*, 27(5), (2004), 529–556.

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).

Diesing, M. *Indefinites* (Cambridge: MIT Press, 1992).

Farkas, D. F. and Y. Sugioka. "Restrictive *If/When* Clauses". *Linguistics and Philosophy*, (6), (1983), 225–258.

Geurts, B. "Generics". *Journal of Semantics*, 4, (1985), 247–255.

Heim, I. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst (1982).

Heim, I. and A. Kratzer. *Semantics in Generative Grammar* (Malden, MA: Blackwell, 1998).

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).

Lange, M. "Are there Natural Laws Concerning Particular Biological Species". *The Journal of Philosophy*, 92(8), (1995), 430–451.

—. “The Autonomy of Functional Biology: A Reply to Rosenberg”. *Biology and Philosophy*, 19, (2004), 93–109.

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).

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

Quine, W. *Word and Object* (Cambridge, MA: MIT Press, 1964).

Sag, I. A. *Deletion and Logical Form*. Ph.D. thesis, MIT (1976).

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

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