

Generics and the Ways of Normality*

Bernhard Nickel · Harvard University

May 28, 2008

Abstract

I contrast two approaches to the interpretation of generics such as ‘ravens are black:’ majority-based views, on which they are about what is the case most of the time, and inquiry-based views, on which they are about a feature we focus on in inquiry—an inductive target. I argue that while majority-based views are preferable based on the most basic data about generics, only inquiry-based views can account for a systematic class of sentences: generics with logically complex predicates, such as ‘cats are black, white, and ginger.’ Thus, inquiry-based views should carry the day. I then go on to sketch a theory of inductive targets.

1 Introduction

Generics fall into many different semantic categories. Some seem to make direct reference to kinds—*genera*, hence generics—in order to say something about them, for example, that the kind is extinct. Other generics seem to speak more directly about the members of a kind and make a general claim about them. The examples in (1) illustrate the latter pattern.

- (1) a. Ravens are black.
- b. Turtles grow old..
- c. Chickens lay eggs.

In this paper, I will be concerned with sentences like these, which I will call *characterizing sentences*.¹ Along with many other theorists, I will assume that characterizing

*[[Acknowledgments suppressed for blind review]]

¹In this terminology, I follow Krifka et al. (1995).

sentences make some kind of general claim about the world, so that we should give truth-conditional, broadly quantificational semantics.² Assuming this much still leaves open just what kind of claim one makes when one asserts a characterizing sentence.

We may define a spectrum of views. On one end are majority-based views on which a characterizing sentence is about what happens most of the time in the actual world. Such views deal with the fact that characterizing sentences tolerate exceptions by taking a long view: though at some time, the aberrations might even dominate, in the long run, the characteristic feature is the one that is present most of the time. However, we may need to “smooth out” the actual patterns, even if we consider the patterns exhibited in the long run, to deal with cases in which the aberrations happen to persistently predominate. Such views still take their departure from the actual patterns, but supplement them with patterns that could have occurred. To determine which such merely possible patterns are relevant to the interpretation of a given characterizing sentence, we have to look towards a “feature of interest” of the kind in question—what I will call an *inductive target*. On the other end of the spectrum are views on which characterizing sentences are only about such inductive targets..

We can illustrate in rough terms what this spectrum comes to by looking at (1a). According to majority-based views, (1a) is true just in case most ravens are black, perhaps not in the actual world at the time of evaluation, but at least when considered over a sufficiently long stretch of time. According to inductive target views, (1a) is true just in case an inductive target, perhaps the genetic endowment of ravens, bears an important relation to their being black. Views that belong between the two extremes combine these two ideas in some way, emphasizing one or the other.

In some cases, all of these views come to the same thing. An inductive target is connected in some important way to a property, as the genetic endowment of ravens is connected to their being black, if and only if most members of the kind in fact have the property at issue, i.e., if and only if most ravens are black.

I will argue, however, that there are extremely widespread counterexamples to majority-based views which are naturally accounted for on an inductive-target view. That there are *some* such counterexamples is well-known, such as (1b), ‘turtles grow old.’ This sentence is true, even though most turtles die almost immediately after they hatch because of predation. And this is not a random aberration in any sense, but a

²This supposition rules out approaches like Carlson’s early kind-centered approach as it is presented in Carlson (1977) and Carlson (1982). It also rules out non-truth-conditional approaches like those proposed in Brandom (2003, chp. 2), Glymour (2002), and Veltman (1996).

completely stable aspect of turtle life.³ However, proponents of a majority-based view might well think of examples like this as an outlier, deserving some more or less *ad hoc* response.

I will focus instead on characterizing sentences with logically complex predicates, such as (2).

(2) Cats are black, white, and ginger.

Adequate semantics for these sentences need to countenance ways of being normal, so that we interpret (2) roughly along the lines of (3).

(3) One way of being a normally colored cat is being black, another way is being white, and another is ginger.

None of these ways of being normal, by itself, is the way most cats are: most cats are not black, nor are most white, nor are most ginger. Instead, characterizing sentences need to be interpreted in terms of inductive targets—they determine ways of being normal. Further, these inductive targets are at best accidentally related to what is the case most of the time. What makes something an inductive target, therefore, is crucially not that it is instantiated most of the time. So, at any rate, I shall argue.

The plan of the paper: in section 2, I begin with some preliminary remarks. I argue in section 3 that any quantificational account of characterizing sentences needs to incorporate two aspects, and that the least committal way of implementing such a minimal quantificational account leads to a majority-based view. Section 4 introduces the majority-based proposal Ariel Cohen has made in recent work, which tries to make do with these minimal features. His work thus provides a sophisticated version of the view I am attacking, and we can see what options are open to its proponents. Section 5 presents the problem of logically complex predicates, while section 6 presents my solution, along with the argument against majority-based views. I sketch a theory of inductive targets and give some more direct evidence for their role in interpreting characterizing sentences in section 7.

³That is why the example is used against so-called normal worlds analyses of characterizing sentences. On such a view, a characterizing sentence is true if the generalization holds always, or at least for the most part, in all relevant normal worlds—see e.g., Delgrande (1987), Krifka (1995), and Morreau (1992). For biological generalizations, for example, we look to biologically normal worlds. But any biologically normal world will precisely be one in which most turtles die because of predation. Hence, such views cannot account for the truth of (1b).

2 Preliminaries

The class of sentences exhibiting some form of genericity is extremely heterogeneous, both syntactically and semantically. To keep the discussion manageable, I will restrict myself to sentences with bare plural subjects that intuitively express some kind of generalization about the members of a kind. How to sharpen this informal gloss is extremely unclear; indeed, it is even a matter of debate how many different readings of bare plurals there are. To deal with this difficulty, I propose to focus on a sufficient condition for being a characterizing sentence. This will give us a class of paradigms that include all of the examples in (1). We can then expand that class as our inquiry warrants.

[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 second clause excludes a number of types of sentences that 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.’⁴

I want to further restrict my discussion to a subset of the characterizing sentences. My reason for this restriction is that many of the phenomena that puzzle investigators of generics are surely interaction effects, results of the interplay of distinctively generic features of a sentence with aspects introduced by non-generic, albeit complex, constructions, such as non-intersective predicates. I will therefore only discuss the *relatively simple* characterizing sentences.

[RELATIVELY SIMPLE CHARACTERIZING SENTENCES]

A characterizing sentence As are F is relatively simple iff

- the main verb is in the present tense,
- the subject is not a mass term,
- the predicate is distributive, and
- the predicate is intersective.

⁴So 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.

All of the examples in (1) are relatively simple characterizing sentences. They are only *relatively* simple, since I do want to allow for logically complex predicates, as in (2). Hopefully, by broadening our scope beyond the completely simple sentences in (1), we can test our theories more thoroughly while still avoiding as many interaction effects as possible.

I need to introduce one more notational convention before we launch into the discussion proper. We sometimes cannot tell from the way a sentence looks what is being characterized. Consider the famous example (4), which can be paraphrased either as (4a) or (4b).⁵

- (4) Typhoons arise in this part of the Pacific.
 - a. It's characteristic for typhoons to arise in this part of the Pacific.
 - b. It's characteristic for this part of the Pacific that typhoons arise there.

The latter, but not the former paraphrase is compatible with there being several hotspots that all are regularly plagued by typhoons. When I talk about a characterizing sentence schematically as *As are F*, I mean to discuss the reading that corresponds to (4a), i.e., one that can be paraphrased as 'it is characteristic for *As* to be *F*.'⁶

3 Minimal Features of a Quantificational Account

The most general schema for a quantificational account is (5).

- (5) *As are F* is true iff in a suitable domain, suitably many *As* are *F*.

Different specific proposals are generated by spelling out the suitability involved. In this section, I argue that any way of spelling out (5) must include two components, what I'll call a modal component and a predicate-induced restriction on the quantifier. I'll also argue that any view that tries to make do with these minimal features must take the form of a majority-based view.

The modal component is the easiest to argue for. Suppose for reductio that the truth of a characterizing sentence depends only on facts about the actual world at the time of evaluation. In that case, (5) reduces to the claim that *As are F* is true iff suitably

⁵The example is originally due to Milsark (1974).

⁶Even though the schema contains the copula, I've included it only for readability. Any predicate can replace 'are *F*,' such as 'lay eggs.'

many *As* are *F*. On this assumption, every way of spelling out ‘suitably many’ will bring in its wake either that *As are F* entails that *some As are F*, or that it is entailed by *all As are F*. But neither entailment goes through in general. For example, (1a), ‘ravens are black,’ could be true even at a time when all the ravens are painted white, so that it cannot entail that some ravens are black. In the same situation, ‘ravens are white’ would be false, even though ‘all ravens are white’ would be true, showing that the universal generalizations doesn’t entail the corresponding characterizing sentence. To deal with this problem, we need to broaden the domain from which we are drawing the *As*. We can do this by waiving the restriction to the actual world, to the time of evaluation, or both.⁷

Next is the restriction induced by the predicate. Consider the pair of sentences (6a) and (6b).

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

According to schema (5), (6a) is interpreted as (7a), (6b) as (7b).

- (7) a. In a suitable domain, suitably many chickens lay eggs.
- b. In a suitable domain, suitably many chickens are hens.

If we hold fixed the domain along with the interpretation of ‘suitably many chickens,’ (7a) entails (7b). That means that on these assumptions, we also predict that (6a) entails (6b), which it clearly does not.⁸

The simplest response is to say that the schematic quantifier ‘suitably many’ is restricted by the predicate, so that we aren’t talking about most chickens *simpliciter*, but rather, most chickens that satisfy some other condition, where that condition is determined by the predicate in some way.

The informal justification for this maneuver is that a sentence like (6a) is about only those chickens that are even in the business of producing offspring. By contrast, (6b)

⁷I include both of these ways of expanding the domain under the heading of *modality* because of the analogy between possibility and time.

⁸Strictly speaking, we need to assume as well that whatever quantifier replaces ‘suitably many’ validates the inference.

- (i) $[Qx : Ax](Fx)$
- (ii) $[\forall x : Ax](Fx \rightarrow Gx)$
- ∴(iii) $[Qx : Ax](Gx)$

Just about any quantifier ever proposed to spell out ‘suitably many’ does.

is about chickens that are in the business of having a gender, a much larger class. On this strategy, the inference from (7a) to (7b) is blocked because in (7b), the domain of the quantifier is far less restricted.

Given these components, the schematic semantics are elaborated as (8).

- (8) *As are F* is true iff in a suitable domain (wider than the actual world at the time of evaluation), suitably many *As* that are also *C* are *F* (where *C* is determined as a function of *F* and possibly the context).⁹

The data I've considered so far are the most secure and widely agreed upon in discussions of characterizing sentences. For that reason, we might try to give a semantic theory that only makes use of the components these data motivate. If we do, we can draw out some further conclusions about such a theory. The most conservative way of spelling out the expansion of the domain is by drawing it just from the actual world, but considering longer stretches of time. If we adopt this further constraint, we can also infer that the quantifier 'suitably many' must be spelled out as 'most.' The quantifier cannot be 'all,' since we will often find cases that would be counterexamples to *all As are F*, even when *As are F* is true. But 'some' is too weak. 'Most' has the best chance at being adequate.

In that case, the schematic semantics turn into (9).

- (9) *As are F* is true iff across a sufficiently long stretch of time, most *As* that are also *C* are *F*.

These semantics are already quite powerful. They deal well with situations in which all ravens happen to be painted white, since these are presumably aberrations from what is true most of the time, to wit, that most ravens are black. They also deal well with sentences like 'mammals give birth to live young,' which is true even though most mammals do not give birth at all—most mammals being either male, immature, or infertile. Thanks to the predicate-induced restriction of the quantifier, none of these mammals are in the range of the quantifier.

Thus, the schematic semantics (9) give good results for many cases, and we can even see that a feature introduced to deal with one problem (the chicken-hen inference)

⁹*C* might be determined purely formally, as is proposed in Cohen (1999b, p. 102) who introduces the notion of a minimal absolute determinable to do the job. *C* might instead be determined by contingent features of our psychology, which is a proposal Cohen has made in his more recent (2004). I don't have a view on the matter, and at any rate, for my purposes here, I can remain neutral.

solves another problem (the preponderance of mammals that don't give birth). In the next section, I'll introduce Ariel Cohen's account as a representative of this approach that retains the basic idea but makes it more sophisticated to deal with recalcitrant cases. There, we'll see some reasons for moving beyond a simple majority-based view to include not the actual pattern we can observe, but "smoothed out" ones, instead.

4 Cohen's Proposal

The view I discuss here emerges from Cohen (1999a) and Cohen (1999b). He presents it as one on which generics in general, and characterizing sentences in particular, are probability claims. As a first approximation, *As are F* is true if and only if the probability that a given *A* is *F* is greater than even chance. My first task in this section is to show that this really is a majority-based view. I'll then explain why we should move towards a smoothed-out version of the view.

The key observation I will cash out in the following paragraphs is that Cohen is a frequentist about probability claims, at least insofar as these claims figure in the interpretation of generics. And the reference class within which the relevant frequency is determined in terms of actual and suitably related histories, what Cohen calls the "admissible" histories.

Cohen represents characterizing sentences schematically as $\text{GEN}(\psi, \phi)$. For example, 'chickens lay eggs' is represented as $\text{GEN}(\text{chickens}, \text{lay-eggs})$. The basic probabilistic notion used in his account is that of a conditional probability, in this case, the probability that something lays eggs, given that it is a chicken, $P(\text{lays-eggs}|\text{chicken})$. A second approximation of Cohen's semantics is then (10).

$$(10) \text{ GEN}(\psi, \phi) \text{ is true iff } P(\phi|\psi) > .5.$$

For the reasons I discussed in section 3, Cohen incorporates a predicate-induced restriction. This restriction works on analogy with focus. The predicate is associated with a set of alternatives, including the predicate itself, and we only consider those objects that satisfy one of the alternatives. In the example of 'chickens lay eggs,' the alternatives consist of ways of producing offspring, such as giving live birth, cell-division, laying eggs, and perhaps some others. The probability involved in determining the truth-conditions of a sentence like 'chickens lay eggs' is then the probability of something's laying eggs, given that it is a chicken and that it satisfies at least one of the

alternatives in the set associated with the predicate. Cohen’s DEFINITION 1 states this generally.¹⁰

[DEFINITION 1] (Generics, first version)

Let $\text{GEN}(\psi, \phi)$ be a sentence, where ψ and ϕ are properties. Let $A = \mathbf{ALT}(\phi)$, the set of alternatives to ϕ . Then $\text{GEN}(\psi, \phi)$ is true iff

$$P(\phi|\psi \wedge \bigvee A) > 0.5$$

At this point, frequentism becomes relevant.

An informal way of stating what frequentism amounts to begins with how we often evaluate probabilities. When asked what the likelihood is that the next flip of a coin lands heads—call that event H —we simply consider all the possible outcomes—heads (H) and tails (T), in this case—and count the favorable outcomes. In the example, there is one favorable among two possible outcomes, making the probability .5.

We can make this model more sophisticated by considering what happens when we are dealing with a setup in which the outcomes don’t all have the same probability, as is the case when we are dealing with an unfair coin or a loaded die. Simply counting favorable outcomes among possible ones isn’t going to do the job, since the possible outcomes of tossing an unfair coin, one that (say) tends towards heads, are exactly the same as the possible outcomes of tossing a fair coin, to wit, H and T . Rather, we need to count favorable outcomes among a set of possible ones that reflects any uneven likelihood of these outcomes. So if the coin is unfair in a way that makes heads twice as likely as tails, the class of possible outcomes should be $\{H, H, T\}$. Counting favorable outcomes among possible ones yields a $\frac{2}{3}$ probability that the next toss yields heads, just as intuition recommends.

Frequentism is the claim that *what makes it the case* that the probability of some event is x is that, in a suitable reference class, the ratio of favorable among possible outcomes is x . Thus, frequentism reverses the order of explanation I tacitly used in the previous paragraph. There, I said that a suitable reference class for the biased coin is $\{H, H, T\}$ *because* it reflects the probabilities involved. By contrast, frequentism says that the probabilities are what they are *because* they reflect the ratios in a suitable reference class.¹¹

¹⁰Quoted from Cohen (1999b, p. 37).

¹¹For technical reasons, frequentists often appeal to infinitely large reference classes (for some of these reasons, see Hájek (1996)). Since our ordinary notion of ratio does not properly apply to infinitely large

The key problem for frequentism is therefore to specify what makes something a suitable reference class for assessing the probability of some event without taking for granted what that probability is. This is the so-called “problem of the reference class.”

We can now apply these considerations to Cohen’s treatment of characterizing sentences. According to (10), $\text{GEN}(\psi, \phi)$ is true iff $P(\phi|\psi \wedge \bigvee A) > 0.5$. By frequentism, this is in turn equivalent to the claim that the frequency of ϕ s in a suitable reference class of ψ s that also satisfy one of the alternatives associated with the predicate is greater than .5. And that just means that most such ψ s are ϕ s.

Like any frequentist, Cohen at this point faces the problem of the reference class. In his case, it takes the form of specifying the domain from which the ψ s are taken among which we evaluate the frequency of ϕ s, since not just any domain will do. For example, if we evaluated ‘ravens are black’ with respect to a reference class drawn only from the ravens that exist in the world of evaluation at the time of evaluation, we would get the wrong result when we consider situations in which all the ravens happened to be painted white. This is where the modal component enters into Cohen’s semantics.

Cohen specifies the reference class in terms of histories. A history is, most generally, a temporal segment of a possible world. It need not be the whole history of a world.¹² The reference class is then generated by combining all histories that are *admissible*.¹³ On Cohen’s view, a history has to satisfy many different constraints to be admissible, and thus to provide the domain from which the relevant ψ s are drawn. I will focus on two. The first is that admissible histories are extrapolations from the history of the actual world as that history has unfolded so far.¹⁴ The second is that the extrapolation occurs not simply by considering ever longer stretches of the future as it will actually take place. Rather, the extrapolation occurs by assuming that history will go on as it has so far, largely unchanged, even if the history of the actual world includes some very significant changes.¹⁵ (11) is an example to make the need for this second aspect of Cohen’s semantics vivid.

(11) Supreme Court Justices have even social security numbers.

collections, frequentists instead employ a generalization of that notion, that of a limiting relative frequency. I mark this fact here only because Cohen casts his discussion in terms of infinitely large reference classes and limiting frequencies. But this complication is inessential to the purposes of this paper, so I will continue to ignore it in the main text, remaining with our ordinary conception of a ratio.

¹²See Cohen’s introduction of the notion at Cohen (1999a, p. 232).

¹³This is the import of Cohen’s DEFINITION 2 from Cohen (1999a, p. 232). In the main text, I am ignoring the complications introduced by using infinitely large reference classes and the ensuing need for relative frequencies.

¹⁴See Cohen (1999a, p. 235).

¹⁵See Cohen (1999a, pp. 235-40).

Let it be the case that all Justices so far have had even social security numbers. Let it also be the case that the Supreme Court is abolished before any Justices with odd numbers are appointed. In that case, *all* Supreme Court Justices, even considered over all of history of the actual world, have even social security numbers.¹⁶ Nonetheless, (11) would still be false. That's why we need to countenance histories that are suitably smoothed out. That is, we need to consider histories that reflect the intrinsic nature of the selection process for Supreme Court Justices insofar as that process doesn't turn on the social security number of the nominees.

Thus, Cohen's view is a majority-based one, but one in which the domain of quantification is determined, at least in part, by reference to an inductive target.

5 Logically Complex Predicates

I now turn to a problem for majority-based views, regardless of whether they are couched in terms of smoothed out histories or not. I'll return to that issue in section 7. Consider (2).

(2) Cats are black, white, and ginger.

I'll argue now that the most straight-forward quantificational semantics lead to a puzzle: they have to treat 'and' as it appears in (2) simultaneously as if it meant 'and' and as if it meant 'or.' I'll then argue that Cohen's more sophisticated semantics cannot solve this problem, either.

Return to the basic schema (5).

(5) *As are F* is true iff in a suitable domain, suitably many *As* are *F*.

The most straightforward way of applying (5) to (2) treats the whole phrase 'are black, white, and ginger' as a single, albeit logically complex, predicate. This is motivated by one of the main inspirations for quantificational account, that characterizing sentences differ from sentences with overt quantifiers only minimally. Both contain a quantificational element as the determiner of the subject noun-phrase. It just so happens that the generic quantifier GEN is unpronounced. On this view, the syntax of (12a) is (12b), just as the syntax of (13a) is (13b).

¹⁶Notice that, because all Justices have even social security numbers, Cohen's appeals to various kinds of homogeneity in the reference class cannot explain why the sentence is false. For these constraints on admissible histories in terms of homogeneity, see Cohen (1999a, pp. 234-50).

- (12) a. Ravens are black.
 b. [DP GEN ravens] [VP are black]
- (13) a. All ravens are black.
 b. [DP all ravens] [VP are black]

In the case of an ordinary quantifier like ‘all,’ a complex verb phrase is then interpreted as placing a single, complex condition on the elements the quantifier ranges over: we interpret ‘all giraffes are tall and yellow’ as saying that each giraffe satisfies the condition of being tall-and-yellow. By parallel reasoning, we interpret ‘are black, white, and ginger’ in (2) as expressing a single condition. Doing so yields (14).

- (14) ‘Cats are black, white, and ginger’ is true iff in a suitable domain, suitably many cats are black, white, and ginger.

But these truth-conditions just cannot be right, since they entail that there are some cats that have all of these colors. (2), however, entails no such thing. Its meaning is captured more closely by (15).

- (15) ‘Cats are black, white, and ginger’ is true iff in a suitable domain, suitably many cats are black, white, *or* ginger.

Indeed, unless we treat ‘and’ as ‘or,’ we cannot avoid predicting that some cats are black, and white and ginger, while maintaining the simple syntactic analysis.

However, if we consider the inferential relations (2) enters into, the ‘and’-as-‘or’ hypothesis makes exactly the wrong predictions. (2) in fact entails (16a) and fails to entail (16b).

- (16) a. Cats are black.
 b. Cats are black, white, ginger, and paisley.

However, on the hypothesis of ‘and’-as-‘or,’ we predict the opposite entailments. In general, we cannot infer from $[Qx : Fx](Gx \vee Hx)$ to $[Qx : Fx](Gx)$, so that the inference from (2) to (16a) no longer goes through. But we can validly infer to $[Qx : Fx](Gx \vee Hx \vee Ix)$, thereby validating the inference from (2) to (16b). So in order to account for this inference pattern, the simple quantificational semantics has to treat

‘and’ as ‘and.’ This gives us the puzzle: the semantics has to treat ‘and’ simultaneously as if it means ‘and’ and as if it meant ‘or.’

One might object at this point to my claims about the inference pattern. Obviously, (2) does not entail (16b), but one might also deny that it entails (16a). After all, the most natural way to understand (1a) ‘ravens are black’ is as saying that black is the *only* characteristic color for ravens. Likewise, out of the blue, the most natural way to hear (16a) is as saying that black is the *only* characteristic cat color. That is simply incompatible with (2), and hence cannot be one of its entailments.

I don’t dispute the contention about how one naturally hears (1a) or (16a). But these are not the only readings available. We can hear (16a) as an appropriate thing to say, as in the exchange (17).

(17) A: Who can tell me what colors cats are?

B: Cats are black.

A: Excellent, any others?

The exchange is completely natural, which means that (16a) in this context is true. I’ll return in section 7 to the question why we don’t usually hear it that way. But for now the important point is that the inference pattern holds, and thus, that the simple quantificational semantics (5) lead to a puzzle.

Let me now turn to Cohen’s account to see whether a sophisticated majority-based implementation can avoid the problem. In thinking through whether Cohen’s semantics can deal with the puzzle, I’ll continue to assume that the whole VP ‘are black, white, and ginger’ expresses a single complex condition. In the next section, I’ll consider what happens when we give up this assumption. There, I’ll argue that if we do, majority-based semantics fall in its wake. That’s why I want to retain the conservative syntactic analysis for now. Doing so also requires us to adopt the ‘and’-as-‘or’ hypothesis, since we cannot otherwise avoid the prediction that at least some cats have all of these colors.

That means that we have to ask the following question: can Cohen’s more sophisticated semantics allow us to retain that hypothesis but do better at predicting the entailments of (2)? Given our assumptions about the syntactic analysis of (1a), Cohen’s semantics renders it as (18), and (16a) and (16b) are interpreted as (18a) and (18b), respectively (note that the ‘and’-as-‘or’-hypothesis is in effect).

(18) In all admissible histories, most *A*-cats are black, white, or ginger.

- a. In all admissible histories, most *A*-cats are black.
- b. In all admissible histories, most *A*-cats are black, white, ginger, or paisley.

Here, the set of alternatives associated with the predicate is represented as ‘*A*,’ so that you should read the expression ‘*A*-cats’ as meaning ‘cats that satisfy one of the alternatives in *A*.’

The only way of blocking the entailment from (18) to (18b) is to suppose that the two sentences are interpreted either with different admissible histories or different alternatives of the predicate. However, no amount of tinkering with the admissible histories or the alternatives will yield the prediction that the inference to (18a) from (18) is valid. And what is worse, even if Cohen’s semantics could somehow be made to predict the inference, the truth-conditions (18a) it predicts for (16a) are suspiciously strong. It doesn’t seem as if someone who says (16a) is committed to saying that most cats are black. Intuitively, they make a weaker claim. So Cohen’s more sophisticated semantics cannot solve the puzzle posed by complex predicates, either.

6 Ways of Being Normal

A first step towards a solution is to jettison the assumption that the whole verb-phrase ‘are black, white, and ginger’ should be interpreted as a single, complex predicate. It may be more useful to treat it as three separate predicates. Informally, that amounts to the contrast between (19a) and (19b) as paraphrases of (2).

- (19) a. It is characteristic for cats to be black, white, and ginger.
- b. It is characteristic for cats to be black, and it’s characteristic for cats to be white, and it’s characteristic for cats to be ginger.

One might try to implement this idea in the framework of a quantificational approach by interpreting (2) as (20).

- (20) In a suitable domain, suitably many cats are black, and in a suitable domain, suitably many cats are white, and in a suitable domain, suitably many cats are ginger.

For many quantifiers that could be used to spell out ‘suitably many,’ however, (20) is still the wrong interpretation, because many such quantifiers still have the consequence

that there are some cats that have more than one of the colors. This is certainly true of ‘most.’ If more than half of the cats are black, and more than half are white, then at least one cat must be both. If we want to use these kinds of quantifiers to spell out ‘suitably many,’ we need to complicate the semantics further.

Let’s introduce the notion of being a normal member of a kind.¹⁷ For some kinds, there’s only one way to be a normal member. Ravens are like that: there is one way to be a normally colored raven, to wit, being black. But for other kinds, including cats, there are many ways to be a normally colored member. I’ll say more about what a way of being normal amounts to in the next section. For now, I want to focus on how introducing such ways helps with the interpretation of complex predicates. In a quantificational setting, we can implement this idea by interpreting (2) as (21).

- (21) There is a way of being a normally colored cat w , such that in a suitable domain, suitably many cats that are w are black, and there is a way of being a normally colored cat w' , such that in a suitable domain, suitably many cats that are w' are white, and there is a way of being a normally colored cat w'' , such that in a suitable domain, suitably many cats that are w'' are ginger.

This interpretation suffers from none of the problems I’ve raised so far. It does not entail that there are any cats that are black, white, and ginger. The entailment patterns are predicted immediately by the ordinary semantics for ‘and.’ And as an added bonus, we no longer predict suspiciously strong truth-conditions for ‘cats are black.’

I can also account for a fact I mentioned earlier, that we ordinarily understand (1a) ‘ravens are black’ as saying that being black is the only way to be a normally colored raven. On my account this is a standard Gricean implicature. (1a) is strictly weaker than any sentence of the form (22).

- (22) Ravens are black and F (where F is some other color).

Suppose now that a speaker utters (1a). In that case an audience might reason as follows. Any sentence of the form (22) is strictly stronger than (1a), yet the speaker did not assert any of these sentences. Moreover, the truth of any such sentence would be

¹⁷The idea of interpreting characterizing sentences somehow in terms of normality is not new. I argue for it in my *Nickel*. It is also suggested, in one way or another, in Asher and Morreau (1995), Dahl (1975), Farkas and Sugioka (1983), Heim (1982), Krifka et al. (1995), and Lawler (1973), along with the theorists I mentioned in note 3. However, the particular proposals differ both from each other and from my own, on the role of references to normality. The idea of appealing to *ways* of being normal is my own.

relevant. Since the speaker is cooperative, she thus must not believe that any such sentence is true. Presumably, the speaker is well-informed on the matter, and we should thus assume that no such sentence is in fact true. And that is just to say that black is the only normal color for ravens.

This account squares well with what happens in the exchange I used to illustrate that ‘cats are black’ can be true, repeated here.

- (17) A: Who can tell me what colors cats are?
B: Cats are black.
A: Excellent, any others?

In this case, it is simply not true that in order to be cooperative, *B* had to say something stronger than she in fact said, even if she thought that something stronger was true. The implicature therefore doesn’t arise, and *B*’s utterance is fine.

The proposal of interpreting (1a) as (21) thus has some nice empirical consequences. But even putting aside for now any concerns about what a way of being a normally colored cat is, one might worry that this interpretation is completely *ad hoc*, since it cannot be predicted by an independently motivated compositional semantics. The problem concerns the relevant scope of the various quantifiers and the main conjunction. As I mentioned earlier, many theorists pursue a quantificational approach by adopting a conservative syntactic analysis, on which the syntax of (2) is essentially (23).

- (23) [_{DP} GEN cats] [_{VP} are black, white, and ginger].

If that’s the right syntax, then there is no way to have the generic quantifier ‘GEN’ scope under ‘and.’ An approach like mine therefore has to reject this syntactic analysis. The easiest way of predicting the interpretation (21) for (1a) would be to have the generic quantifier originate with each of the predicates making up the VP, so that each of ‘black,’ ‘white,’ and ‘ginger’ serves to introduce its own generic quantifier into the interpretation. Whether this idea can ultimately be sustained is an open question, but I do want to point out that the general thought of predicates introducing quantifiers is very common in the interpretation of non-generic plurals, especially to account for distributive readings.¹⁸

Thus, we might interpret (24a) as (24b), with the italicized quantifier originating with the predicate.

¹⁸See, for example, Landmann (2000), McKay (2006), Pietroski (2005), and Schein (1993).

- (24) a. Five boys ate two pizzas.
b. Five boys are such that *for each boy among them*, he ate two pizzas.

If it's reasonable to posit quantifiers introduced by the predicate to account for distributive readings of non-generic plurals, then there can be no *general* problem with positing such quantifiers in the interpretation of characterizing sentences. Supplying the details of such an account has to remain a future project.

Against Majority-Based Views

Consider again (2) and my suggested interpretation (21)

- (2) Cats are black, white, and ginger.
- (21) There is a way of being a normally colored cat w , such that in a suitable domain, suitably many cats that are w are black, and there is a way of being a normally colored cat w' , such that in a suitable domain, suitably many cats that are w' are white, and there is a way of being a normally colored cat w'' , such that in a suitable domain, suitably many cats that are w'' are ginger.

It may well be true that most of the time, most cats are normally colored in one of these ways. But none of these ways, individually, is instantiated most of the time. It is simply not the case that most of the time, cats are normally colored in the black way, as opposed to the white or ginger way. That in turn shows immediately that what makes a way of being a colored cat a way of being a *normally* colored cat is not that it is the way cats are colored most of the time.

Therefore, majority-based approaches to characterizing sentences face far more, and far more systematic, counterexamples than just 'turtles grow old.' Any time that we are dealing with a kind for which there is more than one way of being normal, we have a counterexample to these approaches.

Let me try to be explicit about what I am rejecting. I am not rejecting the claim that we can make sense of characterizing sentences in terms of what happens most of the time across some suitable histories. By itself, the notion of a suitable or admissible history is too flexible to be rejected. What I am suggesting is that we cannot make sense of characterizing sentences in terms of what happens in the actual world most of the time. I suggest that we instead talk about ways of being normal, but my semantics could be couched in the abstract terms of Cohen's account. A Cohen-style interpretation of (2) parallel to my (21) might be (25).

- (25) There is a set of admissible histories H_1 such that most cats in H_1 are black, and there is a set of admissible histories H_2 such that most cats in H_2 are white, and there is a set of admissible histories H_3 such that most cats in H_3 are ginger.

The crucial point is that if one wanted to go this route, one would have to place different constraints on the admissible histories than Cohen in fact places on them. The admissible histories H_1 crucially do not continue an actual pattern of events, since in the H_1 -histories, the black cats predominate. Instead, one would have to constrain the relevant histories in some other way, such as those histories in which a certain way of being normal is predominantly instantiated. In this sense, ways of being normal are fundamental for interpreting characterizing sentences. Whether we actually make use of that notion in giving the logical form of characterizing sentences, or instead give the logical form in terms of some other notion that is then explained in terms of ways of being normal, does not matter to me here.

7 Inductive Targets

I now address question what inductive targets are, and whether we have more direct reasons for believing in them than just that majority-based views fail.

On the approach I am recommending, ‘ravens are black’ is about the genetic endowment of ravens and how it is expressed in some environments, abstracting from certain influences that also play a role in determining color, such as painting. ‘Turtles grow old’ is about the metabolism of turtles, abstracting from other influences that also play a role in determining actual life-span, such as predation.

More generally, characterizing sentences are about a feature of members of a kind that we want to focus on in our inquiry. Usually, such a feature can have a multitude of expressions, depending on whatever other influences a given member of the kind is subject to. A characterizing sentence is then about the result of the interaction between such a feature of interest and some influences, abstracting from others. Which influences we countenance, and which we abstract from, depends on our aims of inquiry.

Once we’ve settled on an inductive target and a set of influences to countenance, we’ve also determined ways of being normal. Depending on the target and these influences, there may be one such way, or there may be more. In the case of cats, for example, the inductive target is a range of genetic endowments which lead to differently colored cats, thereby yielding multiple ways of being normal.

Ultimately, we will need a general theory of inquiry that allows us to see how a goal of inquiry can determine both an inductive target, and a division among the influences between those to countenance and those to ignore. In other words, we need a theory of how we come to categorize the world in the course of our investigations, a theory of natural kinds. But giving the details of such a theory goes far beyond the bounds of this paper.¹⁹ Instead of pursuing that question, let me focus on two reasons for adopting the connection between characterizing sentences and inquiry I have put forth.

Doing so allows us to account for a kind of variability that so far has not been discussed. Consider (26).

(26) Dobermans have floppy ears.

The important fact about dobermans is that they are born with floppy ears that breeders then cut to give them the pointy shape we are familiar with. In the context of evolutionary biology, (26) is true. The text (27) certainly sounds acceptable.

(27) Some breeds of dogs have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, which is why Dobermans have floppy ears.

However, in the context of a discussion of dog breeding, (26) seems clearly false, as the text (28) illustrates.

(28) While Labradors and golden retrievers have floppy ears, dobermans don't. Dobermans have pointy ears.

On the view I am proposing, whether or not (26) expresses a truth depends on the kinds of influences we countenance and which we abstract from, specifically how we treat the interventions of dog breeders. That, in turn, depends on the kind of inquiry we are pursuing. Each of the texts makes a certain inquiry salient, which is why (26) expresses different propositions when embedded in each.

We can also think about my account as the natural end-point of developing majority-based views. To bring this out, let me contrast Cohen's gloss of his view with a corresponding one on mine. Earlier, I suggested that a quantificational account that makes do with a minimum of semantic machinery includes a modal/temporal component and

¹⁹A theory of natural kinds that holds a lot of promise for spelling out the relevant dependence is that of Boyd (1991, 1999) and Kornblith (1993).

a predicate-induced restriction. I also suggested that such a view is implemented most conservatively as a majority-based view without any smoothing out.

On this approach to characterizing sentences, they are not about anything like inductive targets. For example, when we say that ravens are black, what that claims is *about* is the color of ravens over long stretches of time. It's true that we might look at the genetic code of a raven in order to determine what's true most of the time, but that reference to the more or less intrinsic features of ravens only plays an epistemic role. That ravens have a certain genetic endowment which is expressed a certain way under conditions we might call normal, is no part of what the sentence says.

At times, Cohen characterizes his view in these terms. In discussing the sentence 'Mary handles the mail from Antarctica,' which might be true even if so far there hasn't been any mail for Mary to handle, he says this.

[This sentence] does not claim that Mary actually handles mail from Antarctica, but that she is likely to do so. While Mary may never have handled mail from Antarctica, mail from Antarctica may arrive in the future. All that the truth of [that sentence] requires is that in all sufficiently long histories in which mail does arrive, Mary will handle most of it. We may base our prediction that Mary would, indeed, handle Antarctica mail if and when it arrives, on Mary's job description (though we may base it on other things, such as the observed fact that whenever a piece of mail arrived from an exotic place, Mary immediately became curious and asked to handle it); but this is not what the *meaning* of [the sentence], under its descriptive reading, refers to.²⁰

That is to say, Cohen characterizes his view as one on which the sentence is only about what happens in the long run, and Mary's more intrinsic dispositions or her institutional role can only serve as reasons for believing that the long run will take one or another shape.

Cohen goes on to make similar claims about another famous sentence, 'members of this club help each other in emergencies,' which may be true even if no emergence has so far befallen any member of the club. About that sentence, he says:

Again, while the constitution of the club may (but does not necessarily) help us make a prediction about how members behave if and when emer-

²⁰Cohen (1999a, p. 233).

gencies occur, the meaning of [that sentence], under its descriptive reading, does not refer to the constitution.²¹

However, as we have seen, Cohen's view is not about what actually happens over long stretches, but what could happen. I illustrated the need for smoothing out the actual course of history by reference to an inductive target with the example (11).

(11) Supreme Court Justices have even social security numbers.

This sentence is false even if all the Justices that ever exist have even numbers. That the accident is widespread makes it no less of an accident.

To the extent that the admissible histories are determined both by actual patterns, and by inductive targets, characterizing sentences are about both of these things. Thus, Cohen's view is really a hybrid that moves away from the simplest majority-based views. We should go all the way towards a view exclusively about inductive targets.

²¹Cohen (1999a, pp. 233-4).

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.
- Brandom, R. B. *Articulating Reasons* (Cambridge, MA: Harvard UP, 2003).
- 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, (1999a), 221–253.
- . *Think Generic!* (Stanford, CA: CSLI Publications, 1999b).
- . "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).
- Delgrande, J. P. "A First-Order Conditional Logic for Prototypical Properties". *Artificial Intelligence Journal*, 33(1), (1987), 105–130.
- Farkas, D. F. and Y. Sugioka. "Restrictive *If/When* Clauses". *Linguistics and Philosophy*, (6), (1983), 225–258.
- Glymour, C. "A Semantics and Methodology for Ceteris Paribus Hypotheses". *Erkenntnis*, 57(3), (2002), 395–405.
- Hájek, A. "'Mises Redux'—Redux: Fifteen Arguments Against Finite Frequentism". *Erkenntnis*, 45(2-3), (1996), 209–227.
- 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. "Focus and the Interpretation of Generic Sentences". In G. N. Carlson and F. J. Pelletier, eds., *The Generic Book*, 238–264 (Chicago: University of Chicago Press, 1995).

- 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).
- Landmann, F. *Events and Plurality* (Dordrecht: Kluwer Academic, 2000).
- Lawler, J. *Studies in English Generics* (Ann Arbor: University of Michigan Press, 1973).
- McKay, T. J. *Plural Predication* (Oxford: Oxford UP, 2006).
- Milsark, G. *Existential Sentences in English*. Ph.D. thesis, MIT (1974).
- Morreau, M. *Conditionals in Philosophy and Artificial Intelligence*. Ph.D. thesis, University of Amsterdam (1992).
- Nickel, B. "Processes in the Interpretation of CP-Laws and Generics".
- Pietroski, P. M. *Events and Semantic Architecture* (Oxford: Oxford UP, 2005).
- Schein, B. *Plurals and Events* (Cambridge, MA: MIT Press, 1993).
- Veltman, F. "Defaults in Update Semantics". *Journal of Philosophical Logic*, 25(3), (1996), 221–261.