Dutchmen are Good Sailors: Generics and Gradability*

Bernhard Nickel · Harvard University

Abstract

Many generics have quite strong interpretations. *Ravens are black* is naturally paraphrased as *all normal ravens are black.*. However, there is a sundry collection of weak generics, ones that can be true even though most members of the kind at issue do not conform to them, including *Dutchmen are good sailors* and *lions have manes*. This paper presents a new semantic theory for generics on which the weakness of the first of these and others like it is due to the interaction between a generic operator and the semantics of gradable predicates, and argues that the account is empirically and theoretically preferable to extant proposals. Methodologically, the

*Thanks to the participants of the first GENIUS conference as well as two anonymous reviewers for this publication for helpful comments.

paper argues that we should not expect a uniform account of all weak generics.

Keywords GENERICS · SEMANTICS · QUANTIFICATION · HOMOGENEITY · GRADABLE PREDICATES · RELATIVE GENERICS

1 Introduction

Paradigmatic generic sentences seem to express quite strong generalizations. They tolerate exceptions, but these exceptions are plausibly considered deviations from the norm, as in (1) and (2).

(1) Ravens are black.

(2) Tigers have stripes.

However, there are also many generic sentences that are true even though only a minority of the members of a kind conform to the generalization, as in (3)-(6).

(3) Dutchmen are good sailors.

(4) Lions have manes.

(5) Sea-turtles are long-lived.

(6) Frenchmen eat horse meat.

These examples apparently show that the core meaning of generics is much weaker than (1) and (2) suggest. Generics really only require some members of the kind to conform to them, and the initial sentences just happen to also be ones in which almost all members so conformed.

Unfortunately, we cannot easily account for the truth (1)-(6) by positing uniformly weak truth-conditions for generics. The problem arises because of generics that are false, even though most members of the kind conform to it as in (7).

(7) a. Prime numbers are odd.

b. Sea-turtles die young.

This leads to a puzzle. If a generic is true just in case at least some members of the kind conform to it, we do well with (1)-(6) but mistakenly predict that (7a) and (7b) are true. If a generic is true iff most members of the kind, or perhaps all normal members of the kind conform to it, we make the right predictions for (1), (2), and (7), but mistakenly predict that (3)-(6) are false. Either way lies trouble.

This paper defends a view on which the basic meaning of a generic sentence is most clearly exhibited in the initial examples (1) and (2). Roughly, a generic is true iff all normal members of a kind conform to it.[1]

[1]An alternative that could equally well be used holds that a generic is true iff most members of a kind conform to it, though I'll retain the first semantic sketch. Nothing

The main work for such a defense is giving an adequate treatment of examples such as (3)-(6). I'll say that these sentences exhibit "weak" truth-conditions or "weak" interpretations, meaning that they are true even if a majority of the members of a kind do not conform to the generalization. Derivatively, I'll call generics with weak truth-conditions "weak generics."

My main contribution in this paper is twofold. Methodologically, I suggest that the class of generics with weak truth-conditions is not homogeneous. We should not expect a uniform semantic account of how these truth-conditions are determined. Such a denial naturally invites the charge of ad-hocery. If potentially every weak generic can be given its own semantic analysis, a semantic theory for generics becomes untestable. Thus, my claim that the class of weak generics isn't semantically uniform needs to be accompanied by some principled way of partitioning that class.

Substantially, I will discuss generics with gradable predicates such as (3). I will provide a straight-forward compositional semantic theory on which the weak truth-conditions of these generics are due to the interaction of a generic quantifier with independently motivated semantics for gradable predicates. The case study is designed to make my strategy seem less *ad*

turns on this difference.

It also doesn't matter where in the LF of generics the generic operator appears, whether as a phonologically null nominal determiner of a bare plural [e.g. Asher and Morreau, 1995], as an adverb of quantification [e.g. Cohen, 1999a,b, Schubert and Pelletier, 1989, Wilkinson, 1991], or somewhere else. I have argued for the "somewhere else" option in Nickel [2010b, Forthcoming].

hoc because it allows me to capture a linguistically significant generalization: generics with what I will call relative gradable predicates (following Kennedy 2007, Kennedy and McNally 2005) always allow for weak interpretations, while semantically quite similar generics do not always allow for corresponding weak interpretations. In this respect, my theory is superior to some extant alternatives, such as Ariel Cohen's [1999a, 1999b, 2001, 2004], which cannot capture this generalization. Moreover, in the cases for which both theories predict weak truth-conditions, the interpretation my account yields is more closely in accord with our intuitions.

Naturally, my theory needs to be supplemented with further accounts to explain the data in (4)-(6). Nonetheless, the advantages of my account will hopefully make the case that this account should be part of a complete semantic package for generics.

1.1 A Survey of Options

It will be useful to spell out the assumptions that make (3) problematic. This will allow us to organize the theoretical responses to the problem in terms of which assumption is rejected. Three assumptions are required. In order to have a clean representation, assume that a generic *As are F* is represented more formally in terms of a generic operator GEN, as in (8).

(8) [GEN x: $A(x)$]($F(x)$)

The first assumption holds that GEN has a strong meaning, roughly *all normal*, so that a representation of the form (8) is true iff all normal As are F.[2] The second assumption about (3) in particular holds that the restrictor of the generic operator just consists of the property denoted by the bare plural subject, so that (3) should be interpreted as (9).

(9) [GEN x: Dutch(x)](Good.Sailor(x))

The third assumption concerns the interaction of the generic operator and the semantics of the gradable predicate. To state it, I need to say something about these semantics. A plausible theory relates the predicate's argument to an abstract representation of sailing skill, which we can represent as a degree on a scale that allows us to compare various skill levels.[3] The context then supplies another value on that scale as a standard against which the sailing skill of the argument is measured, and the

[2]There are many well-known problems for a normality-based analysis of generics, chronicled already in Carlson [1977]. But for the purposes of this discussion, I'll retain the simple account since I aim to defend a normality analysis. Assuming the simplest (and thus worst) such analysis does not illegitimately stack the deck in my favor.

If framing the discussion in terms fo normality is too distracting, one could re-construe the present discussion as an argument for the conclusion that as far as (3) is concerned, we can rest content with any strong analysis of generics, such as Cohen's absolute generics (see below, p. 10ff). Based on other considerations, I prefer a normality-based analysis [see Nickel, 2008, Forthcoming].

[3]The view that gradable adjectives should be interpreted in terms of scales is extremely widespread in the literature on these predicates. See, e.g., Bartsch and Vennemann [1973], Bierwisch [1989], Cresswell [1976], Heim [2000], Kennedy [1999, 2007], Kennedy and McNally [2005], Klein [1991], Seuren [1973], von Stechow [1984].

predicate is true of its argument just in case the degree to which that argument is proficient at sailing exceeds the contextually supplied standard. For example, *John is a good sailor* is true just in case the degree to which John is a good sailor exceeds the contextually supplied standard.

The third assumption is the following claim. A quantified claim $[Qx\text{: } Fx](\text{Good.Sailor}(x))$ is true iff Q many Fs all posses a degree of sailing proficiency that exceeds the contextually given one. These three assumptions together entail that (3) is true iff all normal Dutchmen possess a sailing skill that exceeds a contextually given and presumably quite high degree of skill, an interpretation that is too demanding.

We can thus consider three ways of trying to solve the problem. First, the restrictor is richer than in the representation (9), an option discussed in §2. Second, the semantics of GEN are weaker than assumed—§3. Third, the assumed interaction between gradable predicates and quantifiers is mistaken. The rejection of the third assumption forms the basis of my account, presented in §4.

2 Accommodation into the Restrictor

We might try to account for the weak truth-conditions of weak generics by suggesting that extra material is accommodated into the restrictor of GEN.

If the additional material is sufficiently demanding, the generic will be correspondingly weak [see, e.g., Geurts, 1985, although the possibility is floated often in conversation].

The concern about this proposal is that, to the extent that mechanisms for its implementation are understood, it cannot account for the data. As a first point, it is well known that the generic quantifier is not contextually sensitive in the same way as other quantifiers. Even in a context in which a possible restriction is very salient, that restriction still cannot be used to interpret the generic. Consider a circus in which all bears are tame and the circus bears are extremely salient because they're currently performing. There is still a significant difference between (10a) and (10b).

(10) a. All the bears are tame.

b. Bears are tame.

The former is clearly true, with *all the bears* restricted to the bears in the circus; the latter is clearly false, in spite of the salience of the tame circus bears.

This observation shows that if additional material is present in the restrictor of the generic quantifier, there must be some other mechanism than quantifier domain restriction based on contextual salience. One such mechanism that we understand fairly well is sentence-internal. Many

quantificational theories of generics hold that the generic quantifier is restricted by some property that is determined as a function of the predicate, motivated by the need to account for examples such as (11a) and (11b).

(11) a. Lions have manes.

b. Lions give birth to live young.

Without a further restriction of the generic operator, we would predict that (11a) and (11b) entail that all normal lions have manes and give birth to live young. Instead, the predicate in (11a) restricts the generic quantifier in such a way that only male lions remain in its domain, the predicate in (11b) so that only female lions do.[4]

This strategy as applied to (3) most plausibly yields the interpretation (12).

(12) Dutch sailors are good sailors.

However, that doesn't seem to be the correct analysis of (3), either. It doesn't seem to be required for (3) to be true that all normal Dutch sailors are good sailors. A fair chunk of them may be quite ordinary. Thus, on the strategy of accounting for the truth-conditions of (3) by restricting the generic quantifier, a better proposal interprets it as (13).

[4]Different theories differ on how this is accomplished, but these differences need not concern us.

(13) Good Dutch sailors are good sailors.

Krifka et al. [1995, 82] consider the possibility that the extra material is accommodated into the restrictor via focal stress on *good*. But as they rightly point out, (3) has a weak interpretation even with neutral intonation. We are thus left without any good account of how the extra material ends up in the restrictor. Thus, the contextually determined restriction strategy is at best an unattractive fallback position as a solution to the problem (3) poses.

3 Cohen: Absolute and Relative Generics

To introduce Cohen's system, let me begin with his treatment of the core examples (1) and (2), which he calls *absolute* generics. He interprets the generic quantifier in such a way that its restrictor is determined, in part, by the predicate via its association with a set of alternatives. To interpret *As are F*, we have to compute the set of alternatives ALT(F). In most cases, F is included in ALT(F), and in most cases, the alternatives are mutually exclusive. For example, to interpret (1), *ravens are black*, we associate the property of being black with alternative colors. With that set in hand, Cohen gives the following truth-conditions.[5]

[5]See Cohen [1999b, 37].

(14) *As are F* is true iff the probability that a randomly chosen A that also satisfies at least one of the properties in $\textsc{alt}(F)$ is F is greater than .5.

These truth-conditions deal well with the lions/manes problem, because the alternatives to the property of having manes are other ways of having sexually selected ornamentation, and no female lions satisfy any of these alternatives. Thereby, the scope of the generic quantifier is restricted so as to exclude them, and within the restricted domain, it is true that the relevant probability is sufficiently high. However, Cohen accepts that these semantics, when applied to (3), yield the truth-conditions that the probability that a randomly picked Dutch sailor is a good one is greater than .5, and that is too strong.

In response, Cohen introduces relative generics, generic sentences that are analyzed in terms of an alternative generic operator. Generics are therefore systematically ambiguous, depending on whether they are analyzed as absolute or relative. The core idea behind relative generics is that they reflect ways in which the members of the kind at issue might be distinguished from members of other kinds that are relevant in the conversation. The relative generic says that the members of the kind at issue are more likely than members of these other kinds to satisfy the predicate of the generic. More formally, when *As are F* is analyzed as a

relative generic, we do not just consider the alternatives to F, but also the alternatives to A, ALT(A). In the case of (3), ALT(A) might include other nationalities. Relative generics have the truth-conditions in (15).[6]

(15) *As are F* is true iff the probability that a randomly chosen A that satisfies one of the alternatives in ALT(F) is F is higher than the probability that an arbitrarily chosen object that satisfies one of the members of ALT(A) and one of the members of ALT(F) is F.

As applied to (3), these semantics yield the truth-conditions that an arbitrarily chosen Dutch sailor is more likely to be a good sailor than an arbitrarily chosen sailor from one of the alternative nations. We can see why this interpretation is aptly called *relative*. A relative generic requires for its truth that the relevant members of the kind be more likely to satisfy the predicate than members of some other kind(s), i.e., we are relating different kinds, such as Dutchmen and Swiss.

So that we may evaluate this proposal, let me say how claims about likelihoods are related to facts "closer to the ground." For our purposes, we can simply translate talk of probabilities into talk of ratios. To say that the likelihood that a randomly picked Dutch sailor is good is higher than the

[6]See Cohen [1999b, 55f]. One benefit of Cohen's strategy is that it makes the difference between relative and absolute generics not completely *ad hoc*. As he argues in Cohen [2001], the difference between the two readings can be reduced to a difference in the setting of one parameter, and he argues that we can see the same difference in parameter setting in some non-generic cases involving *many* and *often*.

likelihood that a randomly picked sailor from some other country is amounts to the claim that the ratio of good Dutch sailors to Dutch sailors of any skill is higher than the ratio of good sailors from other countries to sailors from these countries of any skill.[7]

The introduction of the relative generic quantifier is only half of Cohen's theory. If that was all there was to it, Cohen would predict far more true readings of generics than are intuitively attested. Consider, for example, (16).[8]

(16) Bees are sterile.

It might appear as if the truth-conditions for a relative generic reading of (16) are satisfied, so that (16) ought to be true. After all, the vast majority of bees are worker bees, and worker bees are sterile. Thus, the likelihood that a randomly picked bee is sterile is far higher than the likelihood that a randomly picked member of some alternative kind is.

For this reason, Cohen introduces the *homogeneity* requirement. It is part of his account that when we assess the truth-value of a generic by assessing a ratio within a reference class, that reference class has to be

[7]Within the context of Cohen's theory, the initial interpretation of generics in terms of probabilities plays other roles than simply introducing ratios. It also allows him to motivate various constraints on the classes within which the relevant ratios are assessed, what he calls homogeneity constraints, of which far more immediately below. See also Cohen [1999a, 2004].

[8]This example has also been cited by Leslie [2007], although she does not seriously consider the appeal to homogeneity Cohen makes.

admissible. And a class is admissible only if it is homogeneous: if the relevant ratio is a certain value in the class as a whole, then it needs to be the same in all cells of all psychologically salient partition. If this homogeneity constraint is not satisfied, then the sentence is unacceptable—perhaps it suffers from presupposition failure, perhaps it is simply false.

In the case of (16), the homogeneity constraint is violated because a psychologically salient partition partitions the bees according to sex. And while the (female) workers are sterile, the (male) drones are not. Because (16) would only be true if homogeneity were satisfied, and because homogeneity is violated in this case, (16) is unacceptable.

Cohen's theory of weak truth-conditions thus has two interlocking parts: a relative generic quantifier and the homogeneity constraint. Impressionistically, the very weak relative generic quantifier makes weak truth-conditions available for all generics while the homogeneity constraint pares the class of available readings back so that the account as a whole does not overpredict.

3.1 Problem—Polarization

The truth-conditions of a relative generic are weaker than those of absolute generics by requiring that fewer members of the kind have to satisfy a

contextually given standard. Rather than it being most, it is enough that relatively more than among other kinds do. But this also means that the truth of a generic with a relative gradable predicate is insensitive to what the facts are with regard to those members of a kind that fall below the contextually determined standard. Let me put this in terms of the example (3), *Dutchmen are good sailors*. So long as relatively more Dutchmen exceed the contextually determined standard of sailing skill than do members of another kind, the sentence is predicted to be true. It simply doesn't matter how good or bad the Dutchmen are that fall below the standard.

This commitment of Cohen's semantics leads to a tension. To bring it out, consider the following situation in which (3) is clearly false. There are some very, very good Dutch sailors, with the kind of sailing skill required to have a maritime colonial empire in the Seventeenth Century. Indeed, there are more such excellent sailors, relative to the population as a whole, than there are French and German sailors. But all the other Dutch sailors are simply abysmal.

One way to support this judgment is to note how odd (3) would be in the context of an explanation that is explicitly contrastive. Why, one might ask, did the Dutch manage to have a maritime empire while the Germans did not? If in the kind of polarized situation I just asked you to imagine, one simply offered (3) in response, the proposed explanation would be

awkward—just compare it with the much better alternative *there were some outstanding sailors*. Another piece of support derives from the analogously polarized case (17).

(17) Americans are good athletes.

The wide swath of Americans are horrible athletes, leading an extremely sedentary lifestyle and consuming a poor diet. The fact that there are also quite a few athletic Americans, the kind who regularly exercise a lot, take part in races, competitions, and the like, is insufficient to make (17) true.

The proponent of the relative generic strategy now faces a problem. The fact that all of the non-excellent Dutch sailors are abysmal has to matter to the interpretation of (3). But given the resources available to her, she can only make them relevant by making the truth-conditions of the sentence as a whole too weak. For the sake of illustration, suppose that the only other relevant nationality are French sailors. Recall that the relative generic operator delivers weak truth-conditions not be restricting the domain over which it quantifies, but by essentially introducing a weak condition on that whole domain: that a larger proportion of that whole domain exceeds the standard than does a corresponding population in a corresponding domain. Thus, we cannot make the interpretation of (3) more sensitive to the fact that all of the non-good Dutch sailors are abysmal by including them in the

domain where they weren't included before. They were in the domain all along.

The only thing to do is to try and exploit the fact that the not-so-good French sailors are better sailors than the not-so-good Dutch sailors. To do that, we need to lower the standard required to count as a good sailor sufficiently for more of the French sailors to count as good sailors. That increases the ratio of good French sailors to sailors as a whole. But if we lower the standard in this way, chances are that the not-so-good Dutch sailors still don't count as good, since they're truly abysmal. In other words, it should be possible for us to choose a standard that's low enough so that not-good French sailors can meet it, but that's still high enough so that not-good Dutch sailors cannot meet it.

But there is also a cost to this maneuver. If the standard a sailor has to meet in order to count as good is relatively low, then a completely homogeneously mediocre population of Dutch sailors will verify these truth-conditions. The proponent of the relative generic strategy thus faces a tension. If (3) is interpreted with a reasonably high standard, then a thoroughly polarized population verifies the sentence, even though it's intuitively false. If (3) is interpreted with a reasonably low standard in order to deal with the problem posed by a polarized population, it is verified by a population that uniformly just barely exceeds the standard,

again going against the intuitive judgment that the sentence is false.

3.2 Problems—Explicit vs. Implicit Standards

Cohen's combination of a relative generic quantifier and a homogeneity constraint formulated in non-semantic terms (viz., salient partitions) predicts that certain classes of generics pattern alike. This prediction is mistaken. I'll focus on a constellation of data concerning non-gradable predicates and two kinds of gradable predicates, what are sometimes called absolute and relative gradable predicates.[9]

Let me begin with a distinction among gradable predicates. I will assume that all gradable predicates are interpreted along the lines I sketched in §1.1: the predicate maps its argument to a point on a scale, intuitively the point that represents the degree to which the argument has the relevant property. The predicate is true of its argument just in case the point to which the argument is mapped is appropriately related to a standard on that scale. For some predicates, it may be required that the argument is mapped to a point above the standard, as for *heavy*, while for others it may have to be mapped to a point below that standard, as for *light*.

[9]The terminology for the distinction among gradable predicates is unfortunate in this context since it is liable to lead to conflations with relative and absolute readings of the generic. However, I retain both terminological conventions since they are established in the literature on these topics.

Given these basic assumption, we can now define a gradable predicate as *relative* just in case the standard with respect to which the truth of the predication is evaluated can vary from context to context. These include predicates such as *big*, *tall*, *fat*, *expensive*, and *good F*. What counts as big, for example, can vary from context to context (even holding fixed the dimension we're evaluating), which is just to say that the standard value on the size scale can vary. By contrast, *absolute* gradable predicates are predicates for which the standard is obligatorily set to some privileged value, usually an extreme value on the associated scale.[10] Absolute gradable predicates include *empty*, *full*, *straight*, *wet*, *bent*, and color predicates.[11] To complete the typology, genuinely non-gradable predicates are ones that do not tolerate comparative morphology or intensification, such as *locked*, *political*, *costs $100*, and *is 7ft tall*.

Given this typology, we can describe some patterns regarding the availability of weak truth-conditions. First, weak truth-conditions are always available for relative gradable predicates, as the examples in (18) illustrate.

(18) a. Europeans are tall.

b. Americans are fat.

[10] See also Kennedy and McNally [2005], Rusiecki [1985], Unger [1975].

[11] For evidence that the absolute and gradable predicates really are semantically distinct, see Kennedy [2007, §3.2].

c. Academics type quickly.

d. Math problems are hard.

e. Academics type quickly.

Strikingly, generics that differ from those in (18) only in that the standard of comparison that is left implicit there is made explicit do not have interpretations in terms of weak truth-conditions.

(19) a. Europeans are 1.95m tall.

b. Americans weigh 250lbs.

c. Academics type at 80 words per minute.

d. Math problems require a 160 IQ to solve.

e. Academics type at 80 words per minute.

The question for Cohen's account is whether he can consistently account for this contrast. The availability of the weak readings for relative gradable predicates requires that when we evaluate (18a)-(18d), the homogeneity constraint is satisfied. The fact that weak readings aren't available for the non-gradable versions in (19a)-(19d) requires that when interpreting them, homogeneity is not satisfied. I'll argue that Cohen in fact cannot do so. In outline, my argument will be this. Cohen's best explanation for the difference is that it is the explicitness of the standard in (19a)-(19d) that

makes certain partitions salient, and these partitions aren't homogeneous in the way his semantics require. On that strategy, Cohen must predict that absolute gradable predicates should always have interpretations with weak truth-conditions, since absolute gradable predicates do not explicitly mention the relevant standard.

However, weak truth-conditions aren't available for all generics with absolute gradable predicates, as the examples in (20) show.

(20) a. Luxury cars are black.

b. Buses are empty.

c. Subways are full.

d. German wines are dry.

Let me now spell out this argument in greater detail. Consider (18a) and (19a), repeated here.

(18a) Europeans are tall.

(19a) Europeans are (at least) 1.95m.

The first sentence is true, which means that homogeneity is satisfied. The second is not true, and that must be because homogeneity isn't satisfied, because relatively more Europeans are (at least) 1.95m tall than populations on other continents.

The difference must be due to the way in which we refer to the standard of tallness, in the one case allowing the context to set it implicitly, in the other mentioning it explicitly. This is also the diagnosis Cohen gives.

> whenever a generic predicates a property of some concept, and the predicate contains reference to a value on a scale, the concept is represented as a multidimensional space, with this scale as one of its dimensions. The prototypical case of a value on a scale is, of course, a number word [Horn, 1972]. Hence, any sentence whose predicate contains an explicit number word will be ruled out, by failure of homogeneity, so long as there are *any* exceptions to the predicated property.[12]

That is, a partition along degrees on a scale can be made salient by mentioning one such degree, in our case, a particular height. If we partition the class of Europeans according to height, then the incidence of being at least 1.95m tall will be universal for those partitions that consist of Europeans taller than that, and it will be nil for those partitions that consist of Europeans shorter than 1.95m. Hence, homogeneity fails.

Correlatively, so long as we only implicitly refer to a degree on a scale, as we do in (18a), homogeneity can still be satisfied. Cohen is committed to

[12] Cohen [2004, 548].

this by the logic of his position, and he endorses a very similar diagnosis of the contrast between (21a) and (21b).

(21) a. Mammals have a placenta. (Homogeneity satisfied)

b. Mammals are placental mammals. (Homogeneity violated)

Cohen suggests that when we use the predicate *have a placenta* in (21a), we do not make salient a partition that separates the mammals into placental mammals and marsupials. We make precisely that partition salient in (21b). Hence, there is a difference in partitions made salient by the predicate, *even though* the predicates that appear in the two sentences are necessarily coextensive.

Now the problem: this proposal predicts that generics with absolute gradable predicates should always have interpretations in terms of weak truth-conditions, just as relative gradable predicates do. In both cases, we only implicitly refer to a point on a scale, so homogeneity should be satisfied. This prediction isn't borne out by the data, as the examples in (20) show.

3.3 Taking Stock

The discussion of polarized cases in §3.1 argued that the distribution of the property associated with a scale across the whole population is important

to the evaluation of generics with (relative) gradable predicates, and that Cohen's theory could not account for this fact. The discussion of the interaction between how a standard of comparison is determined and whether weak truth-conditions are available for a generic is intended to show that there is a significant distinction among predicates, but that Cohen's system cannot account for that distinction because he needs to make use of an essentially non-linguistic mechanism—salience—that isn't naturally sensitive to different ways in which an argument might be referred to implicitly, whether by allowing the context to set it or by defaulting to an extreme value of a scale. There may well be responses to each of the problems individually. But the critical discussion should also incline us to consider seriously an alternative theory that accounts for both kinds of facts more smoothly.[13]

In §1.1, I suggested that three hypotheses jointly lead to predicting far too strong truth-conditions for (3), *Dutchmen are good sailors*. The sentence is interpreted as about the whole kind; it is interpreted in terms of a generic operator with a strong meaning such as *all normal*; and gradable predicates interact with quantifiers in such a way that a quantified claim

[13]A point that I do not stress in the main paper, but that is worth bearing mind: an account of generics with gradable adjectives in their positive form should also integrate smoothly with a treatment of generics that involve comparisons, such as *girls do better than boys in grade school*. In other work [Nickel, 2010b], I argue that my account does exactly that.

$[Qx\text{: } Fx]$(Good.Sailor(x)) is true iff Q many Fs all posses a degree of sailing proficiency that exceeds the contextually given one. §2 argued that we shouldn't solve the problem by rejecting the first hypothesis, §3 that we shouldn't solve it by rejecting the second. That leaves the third.

4 Binding the Standard

In fact, quantifiers do not (only) interact with gradable predicates in the way the third hypothesis suggests. Consider the non-generic example (22).[14]

(22) Everyone in my family is tall.

On its most natural reading, (22) means that every member of my family is tall *for the kind of person they are*: I am tall for an adult male, my daughter tall for a five-year-old, etc. The sentence does not require that all of us exceed a single, contextually determined standard. The argument for this conclusion is completely parallel to the argument I discussed to show that the truth of (3) is sensitive not only to facts about the sailors that exceed the contextually salient standard required for the evaluation of a sentence of the form *x is a good sailor* (where x is replaced by a name for a particular person), but also the sailors that fall below that standard.

[14]The example is from Kennedy [2007]. The general point that the operative standard used in the evaluation of a gradable adjective can be bound by a higher quantifier is also discussed in Ludlow [1989], Stanley [2002], though the three accounts differ in the particular implementation.

Clearly, (22) is false if I am tall for an adult male, but the rest of my family is short for the kind of person they are. If we wanted to maintain the semantic principle I just mentioned, we would have to analyze (22) in terms of a contextually salient standard that my five-year-old exceeds. But if we did that, we would predict that the sentence is true, no matter how tall I am, since even a very short adult male will exceed the standard of being tall for a five-year-old (about three-and-a-half feet). So we need to analyze the sentence in such a way that I have to meet a standard that an adult male can reasonably fail to meet, while my daughter has to meet a standard that she can reasonably meet. The only way to do this is to analyze the sentence in such a way that we each have to meet different standards.

Formally, we can implement this desideratum by allowing the quantifier that binds the argument of tall to also bind the standard that the argument has to satisfy. To say that the standard with respect to which an element in the domain of the quantifier is evaluated is bound by the quantifier is not to say that there is no role for context to play. Context can still determine, for each person the quantifier ranges over, just how exacting the standard is. If we're in a conversation discussing the possibility that my offspring has a future as a professional basketball player, I might not count as tall. Formally, I'll assume that the context provides a function from members of the domain to standards with respect to which they're evaluated. A rough

representation of (22) is (23).

(23) $[\forall x\text{: Member.Of.My.Family}(x)](deg_{\text{Height}}(x) > f_i(x))$

In words: each member of my family x is such that x's height-degree exceeds the height-degree the context returns for x via the contextually provided function f_i. Demanding contexts will determine a function that returns a relatively higher standard for each element in the domain, less demanding contexts a relatively lower standard.

We can have an empirically and theoretically attractive treatment of (3) by directly transferring this theory to generics. That is, (3) is given the semantic representation (24).

(24) $[\text{GEN } x\text{: Dutch.Sailor}(x)](deg_{\text{Good.Sailor}}(x) > f_i(x))$

In words: GEN-many Dutch sailors x are such that x's sailing skill exceeds a standard of skill the context returns for x.[15]

We can see how the account deals with the problems posed by (3), and specifically why it doesn't fall prey to the problem of polarization from §3.1. Because the interpretation (24) requires different sailors to satisfy different standards, the very good sailors have to satisfy a relatively high standard while the not-so-good sailors have to satisfy a lower one. Thus, in

[15]The restriction to dutch *sailors* comes about via the general sentence-internal mechanism of restricting the generic operator as a function of the predicate (cf. §2).

a polarized case, (3) is still false because the bad sailors need to satisfy a standard appropriate to not-great sailors, and as a matter of fact, they fail to satisfy that standard. But the relevance of these not-so-good sailors is mainted without making it too easy for the best sailors to count as good.

This proposal has two important benefits. First, it is theoretically simple in that it makes do with a minimum of machinery. We know from the non-generic case that quantifiers can bind standards, and that is all that is required to account for the truth-conditions of (3). We also immediately account for the difference between relative predicates on the one hand and absolute gradable and non-gradable predicates on the other. Only the first have a free variable that a higher quantifier can bind. That is why on the account I am putting forth here, there is no reason to expect weak truth-conditions to be as freely available for generics containing the last two sorts of predicates.[16]

Let me end this paper by considering some concerns one may have about this proposal, one more theoretical, one more directly empirical. First, one might worry that my proposal places undue strain on the context and the background information available to speakers who use (3). To make

[16]I want to emphasize that the mechanism for generating weak truth-conditions I discuss here in terms of binding a standard isn't the only one that leads to weak interpretations. There are others that are in play, for example, for *sea turtles are long-lived* and *Frenchmen eat horse-meat*. My point is only that I don't predict, as Cohen does, that relative and absolute gradable predicates should pattern together in freely allowing weak interpretations.

this problem vivid, let me compare (3) with (22). In the latter, it was relatively clear what feature of each of the objects in the scope of the quantifier determined the output of the function: being an adult male, an adult female, a three-year-old child, and so on. We should ask which analogous features determine the standard a given Dutch sailor in the scope of GEN has to satisfy in order for (3) to be true. In the form of an objection, the proposal is implausible because it requires the context to fix on a function that returns a value for a large number of people with unknown features. How could context possibly do that?

There is a feature that is readily available in the context, however, that can serve the purpose and that leads to the proper results. The standard each Dutch sailor has to satisfy in order for (3) to be true is determined by his or her position in the distribution of sailing skill. Suppose for illustration that we're considering a relatively rough way of carving up the population of Dutch sailors in the scope of GEN into three tiers, top, middle, and bottom. For each sailor, the contextually determined function returns a standard based on her position in the distribution. Sailors at the top end have to satisfy a high standard, sailors in the middle tier a middling standard, sailors in the bottom tier a lower standard.

Evidence for this proposal comes from considering what happens when the sailors in the top tier vanish, perhaps because of a shift in the general

population's interest in sailing so that the most talented members of the next generation pursue something other than sailing. In that case, the not-so-good sailors form the top tier, and they now have to satisfy the standard appropriate to the top tier of sailors, i.e., a much higher one than before. The same holds, *mutatis mutandis* for the bottom tier. Everybody moves up a tier and now has to satisfy a higher standard than before. In that case, the sentence becomes clearly false—Dutchmen used to be good sailors, but they aren't anymore.

One might further worry that the truth-conditions I predict for (3) are too weak for all that I have said here. I predict that (3) is true so long as within each subdivision according to skill, the Dutch sailors in that subdivision reach a skill level that exceeds the standard appropriate for that subdivision. It might appear that the truth-conditions do not require anything in particular about the size of the subdivisions relative to each other. That means that my truth-conditions are insensitive to whether the good Dutch sailors make up only a very small proportion of the Dutch sailors as a whole. Thus, the concern continues, my account predicts that (3) is true even if there are only very few good Dutch sailors. But intuitively, the sentence is false in that kind of situation.

The force of the objection depends on the fact that the situation is underspecified, and filling out the description of the case in different ways

leads to different intuitive truth-value judgments. First way of filling out: the few good Dutch sailors are outliers. It is somehow a blip or an accident that there are some exceptional Dutch sailors (much the way that it would be an accident if a Swiss sailor won an important race today). In that case, (3) strikes me as false, but my account predicts that it is false, too. That's because we're only considering what is true of the normal Dutch sailors. So if the best sailors are all outliers or for some other reason disqualified from the scope of the generic operator, the top of the skill distribution relevant for evaluating the truth-value of (3) is not made up of the few outliers, but of the middling sailors. And they do not exceed the relevant level of skill.

Second way of elaborating the case: normally (in the same sense that is to be spelled out by a theory of generics), there are lots of quite good Dutch sailors, it just so happens that right now, there aren't very many. Perhaps there was a way. In that case, the few remaining good Dutch sailors do qualify for being in the scope of GEN, and hence the top of the skill distribution is made up of them, not their less skillful compatriots. The relatively small size of the top of the distribution is a temporary aberration. And in that case, (3) strikes me as true, as predicted.

Finally, let me consider a counterpart to the polarization worries I urged in §3.1. It seems as if there are sentences for which only the top of a distribution matters, and my semantics is unable to account for them.

Consider the two well-known examples in (25).

(25) a. Brazilians are good soccer players.

b. Italians are good skiers.

Let me set aside one possibly distracting issue. We often *judge* these examples as true based solely on our acquaintance with the very top performers. But these judgments would be completely unreasonable if my semantics were correct, since we'd be committing ourselves on the quality of soccer play throughout Brazil. I think this epistemic concern is misplaced, however, since we have general background beliefs about how skill at sports is distributed throughout a population, and given these background beliefs, we can reasonably take the top of the distribution as evidence for how the rest looks.

The real objection is semantic. Allegedly, the truth of (25a) depends only on facts about the very top of the soccer skill distribution, that is, that the Brazilian national team is usually among the very best in the world and that Brazilian players are among the most dominant players in club soccer. Soccer skill lower down in the distribution is taken to be irrelevant to the truth of (25a).

(25a) is a tricky example in this case, because it's true that soccer skill is relatively high throughout the population. Consider instead (26).

(26) Americans are good chess players.

Even when Bobby Fisher was world champion, we wouldn't accept (26), precisely on the grounds that chess skill throughout the distribution, except at the very top, is relatively low. More generally, the way to test my semantics is with examples in which members of a population excel at something that is rather unpopular within the population as a whole. And as (26) indicates, my semantics fit well with our intuitions about these cases.

5 Conclusion

I have argued that one of the persisting problem cases for quantificational accounts of generics, *Dutchmen are good sailors*, can be addressed using only the barest resources to account for generics—a strong interpretation of the generic quantifier—and exploiting how other elements of the sentence interact with quantification.

I said at the outset that the class of weak generics is not semantically homogeneous, at least to the extent that different mechanisms of interpretation are responsible for the weakness of different subclasses among the weak generics. Indeed, given the extreme simplicity of many weak generics, it is overwhelmingly likely that at least some weak generics

are weak not because the weakness is compositionally derived (as I've argued it is in the case of generics with relative gradable predicates), but because the interpretation of these generics interacts with non-semantic components of cognition. We've already seen one such approach, that of Ariel Cohen. Salience and homogeneity are the provenance of non-linguistic cognition. Another [defended in Nickel, 2010a] holds that what is normal is a matter of categorization, and at least sometimes, what is normal is statistically abnormal, even in the long run.

If anything like these semantic-cum-cognitive approaches are correct, it's important to have evidence that gives us some indication of which generics should be analyzed purely within semantic interpretation and which as an interaction effect. If we can, as I've suggested in this paper, show that a certain class of weak generics is captured by a generalization couched purely in semantic terms, such as the presence of a gradable adjective in which the standard can be bound by a higher quantifier, that gives us reason to think that the weak interpretation of these weak generics is not the result of an interaction between linguistic and non-linguistic components of cognition. Weak generics that resist generalizations in purely linguistic terms are a better bet for being interaction effects.

References

N. Asher and M. Morreau. What some generic sentences mean. In G. N. Carlson and F. J. Pelletier, editors, *The Generic Book*, pages 300–339. University of Chicago Press, Chicago, 1995.

R. Bartsch and T. Vennemann. *Semantics Structures: A Study in the Relation between Syntax and Semantics*. Athäeneum Verlag, Frankfurt am Main, 1973.

M. Bierwisch. The semantics of gradation. In M. Bierwisch and E. Lang, editors, *Dimensional Adjectives*, pages 71–262. Springer, Berlin, 1989.

G. N. Carlson. *Reference to Kinds in English*. PhD thesis, University of Massachusetts, Amherst, 1977.

A. Cohen. Generics, frequency adverbs, and probability. *Linguistics and Philosophy*, 22:221–253, 1999a.

A. Cohen. *Think Generic!* CSLI Publications, Stanford, CA, 1999b.

A. Cohen. Relative readings of *many*, *often*, and generics. *Natural Language Semantics*, 9(1):41–67, 2001.

A. Cohen. Generics and mental representation. *Linguistics and Philosophy*, 27(5):529–556, 2004.

M. Cresswell. The semantics of degree. In B. Partee, editor, *Montague Grammar*, pages 262–292. Academic Press, New York, 1976.

B. Geurts. Generics. *Journal of Semantics*, 4:247–255, 1985.

I. Heim. Degree operators and scope. In B. Jackson and T. Matthews, editors, *Proceedings of Semantics and Linguistic Theory (SALT) 10*, pages 40–64. CLC Publication, Ithaca, NY, 2000.

L. R. Horn. *On the Semantic Properties of Logical Operators in English*. PhD thesis, University of California, Los Angeles, 1972.

C. Kennedy. *Projecting the Adjective*. Garland, New York, 1999.

C. Kennedy. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1):1–45, 2007.

C. Kennedy and L. McNally. Scale structure and semantic typology of gradable predicates. *Language*, 81(2):345–381, 2005.

E. Klein. Comparatives. In A. v. Stechow and D. Wunderlich, editors, *Semantik: Ein Internationales Handbuch der Zeitgenössischen Forschung*, pages 673–691. Walter de Gruyter, Berlin, 1991.

M. Krifka, 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, editors, *The Generic Book*, pages 1–124. University of Chicago Press, Chicago, 1995.

S.-J. Leslie. Generics and the structure of the mind. In J. Hawthorne, editor, *Philosophical Perspectives 21, Philosophy of Mind*, pages 375–403. Blackwell, Malden, MA, 2007.

P. Ludlow. Implicit comparison classes. *Linguistics and Philosophy*, 12(4): 519–533, 1989.

B. Nickel. Generics and the ways of normality. *Linguistics and Philosophy*, 31(6):629–648, 2008.

B. Nickel. *Ceteris Paribus Laws*: Genericity, and natural kinds. *Philosophers' Imprint*, 10(6):1–25, 2010a.

B. Nickel. Generic comparisons. *Journal of Semantics*, 27(2):207–42, 2010b.

B. Nickel. Generically free choice. *Linguistics and Philosophy*, Forthcoming. URL `http://www.bernhardnickel.net/papers/Generic-Predication.pdf`.

J. Rusiecki. *On Adjectives and Comparisons in English*. Longman Linguistics Library, New York, NY, 1985.

L. K. Schubert and F. J. Pelletier. Generically speaking, or, using discourse representation theory to interpret generics. In G. Chierchia, B. H. Partee,

and R. Turner, editors, *Properties, Types, and Meaning, Vol. II*, pages 193–268. Kluwer Academic Publishers, Dordrecht, 1989.

P. A. Seuren. The comparative. In F. Kiefer and N. Ruwet, editors, *Generative Grammar in Europe*, pages 528–564. Reidel, Dordrecht, 1973.

J. C. Stanley. Nominal restriction. In G. Preyer and G. Peter, editors, *Logical Form and Language*, pages 365–388. Oxford UP, Oxford, 2002.

P. Unger. *Ignorance*. Clarendon Press, Oxford, 1975.

A. von Stechow. Comparing semantic theories of comparison. *Journal of Semantics*, 3(1-2):1–77, 1984.

K. Wilkinson. *Studies in the Semantics of Generic Noun Phrases*. PhD thesis, University of Massachusetts, Amherst, 1991.