Contingency and Ideality in Generic Statements

Matt Teichman (University of Chicago)
Melbourne Logic Seminar

August 26, 2016

1 Generics: A Crash Course

Generic statements can be thought of as what appear to be loose generalizations, such as:

(1) Birds fly.

Generalizations like (1) are loose rather than strict, because they are compatible with claims that there exist exceptions to them:

(2) Penguins are birds.

(3) Penguins don't fly.

One core puzzle regarding generic statements is what statements like (1)-(3) could be saying, such that they can all jointly be true.

I won't say much about my approach to that puzzle today, except to note that my semantics for generic statments is a *normality* theory, on which statements like (1) and (2) are statements about what is *normal* for e.g. penguins or birds.

But I will be happy to say more about how my normality theory works during the Q&A session, if there's interest.

2 Bare Plural Generics vs. Indefinite Generics

Generics come in a number of varieties, including:

(4) Grizzly bears hibernate. *bare plural*

(5) A grizzly bear hibernates. *indefinite*

(6) The grizzly bear hibernates. *definite*

Today, we will focus on the difference between bare plural and indefinite generics, which is subtle but (in my estimation) quite philosophically interesting.

The principal semantic contrast between bare plural and indefinite generics is that indefinite generics are choosier about which predicate they will accept in predicate position:

(7) Rap songs have spoken words in them.

(8) A rap song has spoken words in it.

(9) Rap songs are popular.

(10) #A rap song is popular.

For some reason, the bare plural and indefinite are both fine with the predicate *has spoken words*, but only the bare plural generic will accept the predicate *popular*.

Similarly for the following examples: both will admit the predicate *equipped with measuring tape*, but only the bare plural generic will admit the predicate *closed on Sunday*.

(11) Tailors are equipped with measuring tape.

(12) A tailor is equipped with measuring tape.

(13) Tailors are closed on Sunday.

(14) ??A tailor is closed on Sunday.

That is an odd contrast, indeed. What could be causing it?

3 One Standard Answer: Definitional Analyses

The Definitional Anlaysis: Indefinite generics express analytic truths.
'An *F* is *G*' is true just in case being *G* is part of the definition of being *F*.
(Lawler, 1973; Burton-Roberts, 1976; Krifka, 2013)

Why That Helps: Having spoken words is part of the definition of being a rap song, but being popular isn't.

However, the problem with this approach is that it would seem to predict that indefinite generics cannot tolerate exceptions. Whereas indefinite generics are just as exception-tolerant as bare plural generics:

(15) A dog has four legs. Of course, Fido over there is a dog who was injured in a horrible accident, and thus only has three legs.

Examples of this kind of thing abound, and thus the prediction that indefinite generics are exceptionless is undesirable.

4 The Basic Data

I would argue that a semantic analysis of indefinite generics should at least try to capture the following three data points:

Selectivity: Indefinite generics are more selective about which predicates they will accept in predicate position.

Exception Tolerance: Both indefinite and bare plural generics can sometimes retain their truth in the face of counterexamples.

Relative Logical Strength: Any indefinite of the form 'An *F* is *G*' entails the corresponding bare plural of the form '*F*s are *G*.'

To my knowledge, that last point has never been discussed in the literature.

5 My Approach

Bare plural generics have been more thoroughly investigated than indefinite generics.

However, rather than proposing a new analysis of indefinite generics, I am going to recommend carrying previous analyses of bare plural generics over to the indefinite case, then further refining the previous analyses of bare plural generics so as to explain these contrasts.

5.1 The Previous Analysis of BP Generics

In Teichman (2015) and in Teichman (under review), I ascribe the following logical form to bare plural generic statements like 'Bears are furry':

(16)

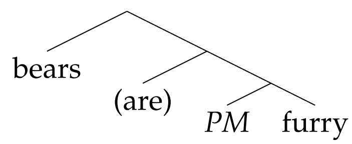

The noun phrase in subject position (*bears*) is taken to denote a kind (bear-kind), the predicate is given an ordinary denotation, and *PM* is a logical operator mapping any property of individual objects to a property of kinds.

The truth conditions for 'Bears are furry' end up being the following, where Base is a modal base function taking any kind at a world of evaluation to the set of outcomes that are ideal for it:

(17) Bears are furry at w iff:
$\forall w' \in \textsc{Base}(\textit{bear-kind})(w)$ ...
... $\exists x\ (\textit{member}(x)(\textit{bear-kind})(w') \wedge \textit{furry}(x)(w'))$

These truth conditions have it that bears are furry just in case at all world at which things go as they are supposed to for bear-kind, certain members of bear-kind are furry.

For convenience, I will refer to the above using the shorthand 'being furry is part of the blueprint for being a bear.'

5.2 Refining the Previous Analysis

The next step is to refine the previous analysis so as to predict the contrast between indefinite and bare plural generics.

Here is the general idea behind the refinement:

An F is G: true just in case being *G* is part of the blueprint for being an F

Fs are G: true just in case some property *P* closely related to being *G* is part of the blueprint for being an *F*.

To spell out what 'closely related' might mean, we turn to Kratzer (1989)'s version of situation semantics.

5.2.1 Situation Semantics, Kratzer Style

The framework in Kratzer (1989) evaluates propositions not at possible worlds, but at parts of possible worlds called *situations*, which we will symbolize using variables s, s', s'', etc..

In this version of the framework, situations are theoretically primitive and ordered under a partial ordering (henceforth, $\sqsubseteq$).

The possible worlds, which we will symbolize using the variables w, w', w'', etc., are the maximal situations under this ordering.

Using these basic ideas, Kratzer defines the following *lumping* relation on proposition:

(18) **(Strong) Lumping**
p (strongly) lumps q at w just in case:

- p is true at w
- every sub-situation of w in which p is true is one in which q is true

I recommend thinking of lumping intuitively via the following two metaphors:

Metaphor #1: p lumps q at w when the 'way in which' p is true, in w, is by q being true. (*Example*: My having painted a still life lumps my having painted a banana at the actual world, because the 'way in which' I painted a still life was by painting a banana.)

Metaphor #2: p lumps q at w when, if you've looked at enough of w to see that p is true, you've thereby looked at enough of w to see that q is true. (*Example*: My having painted a still life lumps my having painted a banana at the actual world, because there's no way to look at enough of the world to see that I've painted a still life without also looking at enough of the world to see that I've painted a banana.)

One final thing we need to do, for reasons that will become apparent later, is eliminate the first condition on lumping, thus making it a reflexive relation:

(19) **Lumping** ($p \rightrightarrows_w q$)
p lumps q at w just in case:
every sub-situation of w in which p is true is one in which q is true

5.2.2 Deriving One Property From Another

Let's work our way up to generic statements about rap songs by starting with a particular statement about a rap song:

(20) *Thrift Shop* is popular.

Is there another feature P such that *Thrift Shop*'s being P lumps its being popular in the actual world?

I would argue that there is: the feature of being socially relevant to its historical moment.

Indeed, unlike being popular, that feature is a pretty good candidate for being part of the blueprint for being a rap song. Arguably, rap always has the aspiration for social relevance at its core.

What about a particular tailor?

(21) *Belinda the tailor is closed on Sundays.*

Being closed on Sundays is not plausibly thought of as part of the blueprint for being a tailor. But there is a closely related property such that Belinda's having that property, in the world we are imagining, lumps her being closed on Sundays: namely, observing the de facto conventions applicable to small businesses in her community.

Let's second-order generalize over these cases. The strategy will be to move from a statement about G being part of the blueprint for being an F to a statement about some closely related P being part of the blueprint for being an F. Informally:

(22) $\exists P\ (P(\mathit{thriftShop}) \rightrightarrows_w \mathit{popular}(\mathit{thriftShop}))$

5.2.3 Pulling a Property from the Actual World

For any bare plural generic of the form '*F*s are *G*,' that statement is going to 'pull' a property from the actual world using this second-order quantifier and whatever *F*s are actually *G*, then say that *that* property (rather than being *G*) is part of the blueprint for being an *F*.

The truth conditions of (9) will be the following:

(23) $\exists P\ (\forall x\ (member(x)(rapKind)(w) \rightarrow ...$
$...\ (\forall s \sqsubseteq w\ (P(x)(s) \rightarrow popular(x)(s)))) \wedge ...$
$...\ \forall w' \in \textsc{Base}(w)(rapKind)\ (\exists y(member(y)(rapKind)(w') \wedge P(y)(w')))$

This formula is a bit unwieldy, but let's walk through it step by step.

It is a long conjunction in the scope of a second-order quantifier, which means that it has the overall form $\exists P(\phi \wedge \psi)$. The first conjunct states that for every actual rap song r, there is some property P such that every actual situation in which r is P is a situation in which r is popular.

The second conjunct takes the property P, on which the first conjunct zeroed in, and states that at every possible world at which things go as they are supposed to for rap song-kind, some rap songs are popular.

To derive those truth conditions fully compositionally, we make the simplifying assumption that *rap songs* is an individual constant denoting a kind, then give *PM* the following denotation, indicating the types of variables via subscripted Montague notation:

(24) $[\![\mathbf{PM}]\!]^{M,w,g} = \lambda f_{\langle e,\langle s,t\rangle\rangle}\ .\ \lambda K_k\ .\ ...$
$...\ \exists h_{\langle e,\langle s,t\rangle\rangle}\ (\forall y_e(member(y)(K)(w) \rightarrow ...$
$...\ (\forall s_s \sqsubseteq w\ (h(y)(s) \rightarrow f(y)(s)) \wedge ...$
$...\ \forall w'_s \in \textsc{Base}(w)(K)\ (\exists x_e(f(x)(w') \wedge h(x)(w'))))))$

A fully compositional analysis of indefinite generics is not quite within our reach, due to complications involving the many different meanings of *a*. However, as a stepping stone on the way to such an analysis, we can ascribe to indefinite generics the same truth conditions the old analysis of bare plurals ascribed to bare plurals.

This delivers, right off the bat, the happy result that any indefinite generics will entail the corresponding bare plural generic:

(25) **Bare Plural Generic:** 'Rap songs have spoken words.'
$\exists P\ (\forall x\ (member(x)(rapKind)(w) \rightarrow$ ...
... $(\forall s \sqsubseteq w\ (P(x)(s) \rightarrow hasSpokenWords(x)(s)))) \wedge$ ...
... $\forall w' \in \textsc{Base}(w)(rapKind)\ (\exists y(member(y)(rapKind)(w') \wedge P(y)(w')))$

(26) **Indefinite Generic:** 'A rap song has spoken words.'
$\forall w' \in \textsc{Base}(w)(rapKind)$ ...
... $(\exists y(member(y)(rapKind)(w') \wedge hasSpokenWords(y)(w')))$

Why? If we look at the truth conditions side by side, we see that since the lumping relation we are using is reflexive, having G be part of the blueprint of F-kind is just a special case of having some property P that lumps being G at the actual world be part of the blueprint for F-kind.

6 A Problem

The analysis under consideration here comes very close to being able to satisfy the three desiderata indicated earlier. However, in the form just outlined, it has a serious bug.

The bug comes to light as soon as we are in a scenario where being G is part of the blueprint for being an F, but due to unusual circumstances, no F actually are G.

In that scenario, all generic statements about Fs become trivially true.

For instance, suppose a fanatic has tracked down every last bear in the world and shaved it. Perhaps, in that scenario, bears are still furry, because being furry is part of the blueprint for being a bear. The trouble is that this sentence comes out true:

(27) Bears are fluent in English.

Why? Because trivially, for every bear, every actual situation in which it is furry is a situation in which it is fluent in English, given that in no actual situation is any bear furry. But being furry is part of the blueprint for being a bear.

6.1 Option 1: Relevant Logic

My preferred option for dealing with this problem is to replace the material conditional operator in the semantics given above with a relevant conditional operator.

The general motivation for doing so is the following. Suppose we can find a conditional operator $\succ$ with the following two features:

- When q entails p, $q \succ p$ is logically valid.
- When q does not entail p and q is false, $q \succ p$ is false.
- In other respects, $\succ$ behaves like the material conditional.

If that's the case, then replacing the material conditional operator in the above definition of *PM* should have the following beautiful result:

Beautiful Result: When no *F*s actually are *G*, the truth conditions for '*F*s are *G*' flip to being those of 'An *F* is *G*.'

Why? Because when no *F*s actually are *G*, then the only property which lumps *G* in the actual world will be *G* itself.

However, I have not quite gotten this solution to work as of yet. I'm still fairly new to relevant/ce logic!

Let me walk you through my difficulty. Suppose we assume that $\succ$ is at least as strong as the conditional from Priest & Sylvan (1992)'s **B**. We could define a version that conditional in our type logic, with *R* as a special predicate symbol for the ternary accessibility relation of the Routley-Meyer semantics:

(28) $\succ^{\mathbf{B}} =_{def} \lambda s_s \,.\, \lambda \phi_{\langle s,t \rangle} \,.\, \lambda \psi_{\langle s,t \rangle} \,.\, \forall s' \forall s'' \; (R(s)(s')(s'') \rightarrow (\phi(s') \rightarrow \psi(s'')))$

Then, if we adopt the following notational abbreviation:

(29) $\succ^{\mathbf{B}}_s =_{def} \succ^{\mathbf{B}}(s)$

...we can essentially just stick this relevant conditional in exactly where the material conditional used to be:

(30) $[\![\mathbf{PM}]\!]^{M,w,g} = \lambda f_{\langle e,\langle s,t \rangle\rangle} \,.\, \lambda K_k \,.\, ...$
$... \; \exists h_{\langle e,\langle s,t \rangle\rangle} \; (\forall y_e(member(y)(K)(w) \rightarrow ...$
$... \; (\forall s_s \sqsubseteq w \; (h(y) \succ^{\mathbf{B}}_s f(y)) \wedge ...$
$... \; \forall w'_s \in \textsc{Base}(w)(K) \; (\exists x_e(f(x)(w') \wedge h(x)(w'))))))$

The trouble is that in order for the generalization over sitiuations in that part of the definition to continue to capture lumping, it needs to be over sub-situations of the world of evaluation.

One way to secure that result would be to institute the following reflexivity constraint on R:

(31) **Reflexivity:**
For every situation s, $R(s)$ is reflexive.

However, I have the impression that this is not a constraint we want R to have in relevance logic—I seem to recall reading a remark by Greg Restall to the effect that subjecting R to this constraint means letting the following woefully irrelevant entailment in:

(32) $A \vdash B \rightarrow B$

So it seems like I'll have to find another way to integrate lumping with the relevant conditional in type logic.

Suggestions would be welcome!

6.2 Option 2: Ambiguity in *PM*

A fallback option, in case that doesn't work, is to make bare plurals systematically ambiguous between the old truth conditions, wherein '*F*s are *G*' is true just in case being *G* is part of the blueprint for being an *F*, and the new definition, wherein '*F*s are *G* is true just in case for some property *P* that lumps being *G* at the actual world, being *P* is part of the blueprint for being an *F*.

This option will yield the result that the indefinite entails the corresponding bare plural, but in a comparatively uninteresting way.

It will do so by making *PM* lexically ambiguous—though some explanation for why indefinite generics don't have the 'lumping' interpretation will still be needed.

7 Concluding Remarks

The foregoing might seem like a set of somewhat abstruse considerations. However, I think that a deep philosophical issue underlies the contrast between indefinite and bare plural generics.

We saw how earlier authors tried to make the indefinite/bare plural distinction into a distinction between lawlike and accidental statements.

The problem was that the supposedly lawlike statements—the indefinite generics—were not as strict and lawlike as that analysis made them out to be, because they can still be true even when some instances fail to exemplify the relevant property.

And the supposedly accidental statements—the bare plural generics—were not as accidental as that analysis made them out to be, as we can see from the fact that they can still be true even when none of the instances exemplify the relevant property.

That raises a puzzle: both indefinite and bare plural generics describe ideal features of kinds, but one of them seems to 'involve' historical contingencies somehow even as it describes ideals. As a result, we get e.g. the contrast between having spoken words and being popular.

I think this is a symptom of the fact that our ordinary vocabulary for talking about ideality and contingency is too expressively impoverished to get at interesting distinctions like this.

Bare plural generic statements depend for their truth on what is ideally the case, but indirectly, by way of information about historical accidents in the actual world. Which property they implicitly ascribe as part of an ideal is affected by the way things actually turned out.

Many arguments that lean on the philosophical significance of generic statements—especially investigations into whether they are 'normative' or 'descriptive' (Rescher, 1994; Thompson, 2009)—pass a bit too quickly over these important subtleties.

The result, I would argue, is yet another case study in how carefully examining the workings of logic and language can push us to develop terminology for talking and reasoning about distinctions we never previously knew were there. If we want to have a good understanding of what's epistemically (and sometimes morally!) at stake when we make generic statements, we can't brush these subtleties to one side.

References

Burton-Roberts, Noel. 1976. On the Generic Indefinite Article. *Language*, **52**(2), 427–448.

Kratzer, Angelika. 1989. An Investigation of the Lumps of Thought. *Linguistics and Philosphy*, **12**(5), 607–653.

Krifka, Manfred. 2013. Definitional Generics. *In:* Mari, Alda, Beyssade, Claire, & Del Prete, Fabio (eds), *Genericity*. Oxford University Press.

Lawler, John Michael. 1973. *Studies in English Generics*. Ph.D. thesis, The University of Michigan.

Priest, Graham, & Sylvan, Richard. 1992. Simplified Semantics for Basic Relevant Logics. *Journal of Philosophical Logic*, **21**(2), 217–232.

Rescher, Nicholas. 1994. *Philosophical Standardism*. University of Pittsburgh Press.

Teichman, Matt. 2015. *Characterizing Kinds*. Ph.D. thesis, University of Chicago.

Teichman, Matt. under review. The Sophisticated Kind Theory.

Thompson, Michael. 2009. *Life and Action: Elementary Structures of Practice and Practical Thought*. Harvard University Press.