

Do Ducks Lay Eggs? How People Interpret Generic Assertions

Sangeet Khemlani¹, Sarah-Jane Leslie², Sam Glucksberg¹, and Paula Rubio Fernandez¹

{khemlani, sjleslie, samg, prubio}@princeton.edu

¹Department of Psychology

²Department of Philosophy

Princeton University

Princeton, NJ 08540 USA

Abstract

Generics are statements that express generalizations, such as ‘ducks lay eggs’. Intuitively, such statements seem true. Even the universal form of such statements e.g., ‘all ducks lay eggs’ seems true, despite our knowing that the majority of ducks do not. We conducted an experiment to verify these intuitions, and found that people overwhelmingly judged generic assertions true. People also judged universally quantified assertions true, but to a lesser extent. These results are consistent with the hypothesis that generics express cognitively primitive generalizations, and so require fewer cognitive resources than do explicitly quantified assertions (Leslie, 2007). Hence when people encounter universally quantified assertions they treat them as generics in order to minimize cognitive effort.

Keywords: generics; semantics; linguistics; philosophy

Introduction

Generics are statements such as ‘tigers are striped’, ‘cars have radios’ and ‘ravens are black’. Such statements lack explicit quantificational operators (e.g., ‘some’, ‘all’, or ‘most’) but nevertheless express generalizations, rather than claims about specific individuals. This is true not just for English, but for all known languages for which generics have been explored (Krifka et al., 1995; Dayal, 1999). Generic statements are interesting for a number of reasons. Though they occur frequently in everyday speech, it is not immediately apparent how generics are interpreted. They do not express exceptionless universal generalizations; while it is false that all tigers are striped (a few unfortunate albino tigers are quite stripeless), the generic claim ‘tigers are striped’ remains true in the face of these exceptions. The generic claim ‘tigers are striped’ is thus not equivalent to its universal counterpart ‘all tigers are striped’. Perhaps, then, generics might be assimilated to the quantifier ‘most’ — certainly it is true that most tigers are striped. Such a proposal runs into difficulties when we consider generics such as ‘lions have manes’ and ‘mosquitoes carry the West Nile virus’. The mature maned males make up less than half of the lion population, so it is false that most lions have manes, yet the generic strikes us as true. Further, over ninety-nine percent of mosquitoes are perfectly virus free, yet ‘mosquitoes carry the West Nile virus’ is a true generic.

If one were tempted to understand generics as meaning “all normal”, the last two examples should give one pause for thought. There is nothing abnormal about mane-less female lions, so it is surely not true that all normal lions have manes. Nor is it true that all normal mosquitoes carry the

West Nile virus; if anything, the norm is for a mosquito to be virus-free. Many theorists have proposed variations on such an account (e.g. Asher & Morreau, 1995; Pelletier & Asher, 1997; Greenberg, 2003). These accounts are discussed at much greater length in Leslie (2007; forthcoming a).

One might suppose that generics are interpreted as existential statements, such that the generic ‘*Ks are Fs*’ is said to be true whenever some *Ks* are *Fs*. One might hope to thus explain the otherwise puzzling truth of ‘mosquitoes carry the West Nile virus.’ Yet this interpretation fails to account for generics such as ‘birds are female’, which is intuitively false even though, of course, some birds are female.

As these examples suggest, generics cannot be accounted for in terms of quantifiers such as ‘all’, ‘most’, or ‘some’. Further, language acquisition studies suggest that while generics lack explicit quantification, children acquire the conceptual framework of generics and employ them in their speech at least a year before they use explicit quantifiers (Hollander, Gelman, & Star, 2002; Gelman, Geotz, Sarnecka, & Flukes, forthcoming).

These observations suggest that generics are not quantificational. Quantificational statements are about how much or how many in a way that generics are not. Notice that, upon being asked ‘how many tigers are striped’, one might reply ‘most tigers are striped’, or ‘some tigers are striped’, but one cannot reply ‘tigers are striped’. The generic is not an appropriate answer to this question (Carlson, 1997). Leslie (2007; forthcoming b) argues at length that the truth and falsity of generics does not depend on how many of the relevant individuals possess the predicated property. There is no sense, she claims, in which generics are dependent on such quantitative considerations. The results presented here offer empirical support for this claim.

Leslie’s Account of Generics

As the foregoing discussion might suggest, the truth conditions of generics are an extremely complex matter. Generic truth conditions are so divorced from quantitative considerations that ‘mosquitoes carry the West Nile virus’ is true even though less than 1% of mosquitoes carry the virus, while ‘books are paperbacks’ is false, even though over 80% of books are paperbacks. The philosophical and linguistic literature on generics offers accounts that are often so baroque as to take a half a page to simply state. These accounts make use of everything from comparative probability to non-standard

logics. In short, they are exceedingly complex.

In this respect, semantic accounts of generics contrast sharply with semantic accounts of quantifiers such as ‘all’, ‘some’ and ‘most’, which may be given simply and concisely (e.g. ‘all K s are F ’ is true iff $\{x: x \text{ is } K\} \subset \{x: x \text{ is } F\}$). The contrast, while remarkable enough on its own, becomes quite vexing once we consider data from language acquisition. Generics, it appears, are considerably easier for children to acquire than are quantifiers (Gelman, 2003; Roeper, Strauss, & Pearson, 2006). We should wonder, then, how it is that children find the theoretically vexing generic easier to acquire than the far more tractable quantifiers.

Leslie (2007; forthcoming a) argues that this tension is resolved if we understand generics to give voice to our cognitively primitive generalizations. The generalizations that generics express are produced by a cognitively fundamental mechanism, which is likely present even before the advent of language acquisition. Quantifiers, in contrast, express generalizations that are more cognitively sophisticated.

Leslie argues that this mechanism deals with different types of information in different ways. In particular, she argues that the mechanism generalizes information differently depending on whether the information is characteristic of the kind in question, whether it is striking (often horrific or appalling), or neither. Broadly speaking, a generic ‘ K s are F ’ is judged true if a) the property of being F is characteristic of the kind K , else b) being F is a particularly striking property, and some members of the kind K possess it, else c) the majority of the kind K possess the property of being F (for far greater detail, see Leslie (2007; forthcoming a; forthcoming b)).

Prasada and Dillingham (2006) discuss what they call k -properties and t -properties. K -properties, on their view, are statistically frequent properties that individuals are expected to possess in virtue of belonging to a particular kind, while t -properties are properties that are merely statistically frequent amongst the members of that kind. It should be noted that their taxonomy, while related, is quite different from the taxonomy presented here. One key difference between Leslie’s characteristic properties and Prasada and Dillingham’s k -properties is that, on Leslie’s view, there is no expectation that characteristic properties will be statistically frequent amongst the members of the kind. (Quite the contrary: ‘lions have manes’ is an example of a characteristic generic in which the property is possessed relatively infrequently.) Also, the category of striking properties does not appear in Prasada and Dillingham’s taxonomy.

Thus in Leslie’s taxonomy, there are three predicate types that are relevant to understanding when generics are judged true:

- The predicate in question is *characteristic* of the type in question, e.g., ‘ducks lay eggs’, ‘horses give live birth’, ‘lions have manes’.
- In matters of *striking* or *dire* consequences only the existential is satisfied, not necessarily the majority, e.g.,

‘mosquitoes carry the West Nile virus,’ ‘sharks attack bathers’, ‘ticks carry Lyme disease’.

- The *majority* of the type in question satisfies the given predicate, e.g., ‘cars have radios,’ ‘barns are red.’

While a generic statement may be true if it meets one of these conditions, the same is obviously not true for universally quantified statements, which are strictly true only when 100% of the kind possess the relevant property — regardless, of course, of whether that property is characteristic, striking, or what have you. Nonetheless, we conjectured that people might be inclined to mistakenly assent to universally quantified statements, especially when the predicate in question is characteristic of the kind. Some characteristic predications, we supposed, might pass muster when quantified universally, e.g., ‘all ducks lay eggs.’ In contrast, universally quantified versions of majority and striking predications seem to be false. For example, ‘all cars have radios’ and ‘all ticks carry Lyme disease’ don’t ring true.

‘All ticks carry Lyme disease’ should be judged as false, because each and every tick does not carry the disease, only some individual ticks do. We should thus expect that striking/dire predicates should always be judged true when existentially quantified, but rarely judged as true when universally quantified. Majority predications may be the simplest case. Existentially quantified assertions are counted true, while universally quantified ones are counted false whenever the majority is less than 100%.

Characteristic predications present a more complex picture. Such predications may be interpreted as predicating a property of a kind, yet it is not a property had by all individual members of that kind. Therefore the assertion ‘all ducks lay eggs’ is false, and should be judged as such whenever any counterexamples come to mind, e.g. when it is recognized that male ducks and immature female ducks do not lay eggs.

Apparently, although such counterexamples are available to people, they do not seem to be spontaneously accessible. If they were spontaneously accessible, people should not agree to universally quantified characteristic statements such as ‘all ducks lay eggs’. Yet casual observation suggests that people do agree with such statements. Leslie argues that understanding characteristic generics is a cognitively primitive operation, and so is easier than understanding quantified assertions (cf. Hollander et al., 2002, on children’s acquisition of generics prior to quantified assertions). If characteristic generics are easier to process than quantified assertions, then people might default to the generic form of characteristic predications when they encounter universal characteristic predications. Specifically, when encountering an assertion such as ‘all ducks lay eggs’, people treat it as if it were ‘ducks lay eggs’ in order to minimize cognitive effort.

In order to examine whether people do in fact accept universally quantified characteristic predications (at least those that are true in generic form), while rejecting universally quantified majority and striking consequence predications,

we conducted an exploratory experiment to see how people judge the truth value of these three types of predications as a function of statement type, be it existential, generic, or universal.

Experiment

We asked a sample of Princeton University undergraduates to judge the truth value of nine different kinds of assertions: three types of predicates, each appearing in generic form, the existential quantifier ‘some’, or the universal quantifier ‘all.’ That is, participants could see the majority-predicate ‘barns are red’ as ‘barns are red’, ‘some barns are red’, or ‘all barns are red’. The study thus employed a 3 (predicate-type: characteristic, majority, striking) × 3 (statement-type: existential, generic, universal) repeated measures experiment.

Method

Participants. Twenty-eight undergraduate students at Princeton served as participants. All spoke English as their first language and none had participated in experiments concerning generics before.

Procedure and Materials. Participants were first asked to judge the truth-value of nine different types of assertions, using a PC running LispWorks 4.4 Professional Edition. They pressed one key designated as ‘yes’ and another designated as ‘no’ for each statement when it appeared on the computer screen. An initial training phase consisting of responding to the words ‘yes’ and ‘no’ respectively was used to familiarize participants with the keyboard key assignments. Participants were given nine types of assertions: the three predicate types, each appearing in existential, generic or universal form. For each item, the generic version of the statement was intuitively true. These statements were counterbalanced via four distinct 9 × 9 Latin squares. In addition to the nine experimental items in each Latin square, nine false filler statements (such as ‘notebooks are pencils’ and ‘rats have stripes’) were presented to roughly equalize the number of ‘yes’ and ‘no’ responses.

Participants first completed the statement-verification task using the nine experimental items described above. They were then given each statement that they had seen before to estimate the proportion of members of each category that shared the property in question. For example, if a participant had seen the statement ‘ducks lay eggs’ in the verification task, then he or she was asked to estimate the proportion of ducks that lay eggs, e.g., “What percent of ducks lay eggs?”

Participants also completed a variety of other tasks related to their understanding of generic assertions to provide information for future experiments, but they are not relevant to this paper and so will not be discussed further.

Results and Discussion

The proportions of ‘yes’ responses to each of the nine types of assertions are provided in Table 1 and Figure 1. As expected, virtually all agreed that existentially quantified state-

ments were true, regardless of predicate-type. The preponderance of characteristic predications in generic form were judged true (.89), with somewhat fewer agreeing that striking and majority were true in generic form (.68 for each type). This difference between the characteristic predicates and the other two predicate types may not hold in general because we sampled only a few items of each type. Future work should sample a wider variety of predicates of each type before we can conclude that characteristic predicates are more readily viewed as true than majority and striking generics.

Table 1: Mean proportions of ‘yes’ responses as a function of predicate-type and statement-type (mean estimated proportion of category members with relevant property).

	Existential	Generic	Universal
Characteristic	.93 (.65)	.89 (.70)	.46 (.97)
Majority	1.00 (.66)	.68 (.80)	.07 (.97)
Striking	.93 (.26)	.68 (.36)	.07 (.95)
Total	.95 (.52)	.75 (.62)	.20 (.96)

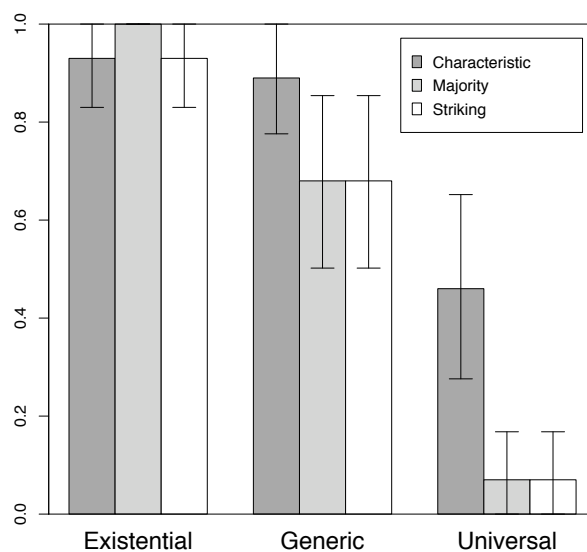


Figure 1: Mean proportions of ‘yes’ responses as a function of predicate-type and statement-type. Bars reflect 95% confidence intervals.

Finally, consider predications quantified by ‘all’. Characteristic predications were far more likely to be judged true in universal form than were majority and striking (46% vs. 7% for the latter two types). In our set of items, the universally quantified majority predications are of course false, since we chose them to be so. Universally quantified striking predications will be judged as true or false depending upon either of two conditions. First, if an item is construed as referring to the type of category member rather than to individual members, then it should be judged as true (e.g., ‘all

cigarettes cause cancer' if construed as 'all (kinds or brands) of cigarettes cause cancer'). Second, if an item is construed such that each individual member of the category contributes to the striking consequence, then it should also be judged true in 'all' form (e.g., 'all cigarettes cause cancer' is true given this construal).

Obviously, very few of our striking predications were so construed; only 7% of responses to striking predications in universally quantified form were judged true. Again, given the small sample of items used in this exploratory study, we cannot generalize to the class of striking consequence predications, but the evidence is certainly suggestive. It is very likely that counterexamples to such assertions are readily accessible that is, they come readily to mind, and so are rarely if ever judged to be true.

The proportion of agreement data were subjected to a 3×3 within-subjects ANOVA. The test revealed a main effect of predicate-type, $F(2, 216) = 8.47, p < .05, \eta_p^2 = .071$, reflecting the higher agreement rate to characteristic generic statements than to either of the other two types (.76 vs. .58 and .56 for majority and striking, respectively). There was also a main effect of statement-type, $F(2, 216) = 104.24, p < .05, \eta_p^2 = .491$, reflecting the highest agreement for 'some' statements (.95) than for generic statements (.75), with the lowest agreement for 'all' statements (.20). Finally, there was an interaction between predicate-type and statement-type, $F(4, 216) = 3.66, p < .05, \eta_p^2 = .063$. This can be interpreted in terms of the differences in agreement rate for statement-type as a function of predicate-type. Agreement rates for 'some' quantified assertions were virtually equivalent across predicate types. In contrast, only characteristic predicates yielded a robust agreement rate in the universal condition (.46 compared to .07 for the other two predicate types). Finally, generic statements did differ as a function of predicate-type, with characteristic assertions eliciting a somewhat higher agreement rate (.89) than either majority or striking (both .68).

Because agreement rates in the existential condition were uniformly high, a 2×3 ANOVA omitting the existentially quantified condition was conducted. The test revealed significant effects for statement-type, $F(1, 135) = 80.96, p < .05, \eta_p^2 = .733$, and predicate-type, $F(2, 135) = 11.05, p < .05, \eta_p^2 = .429$, though the interaction was not significant, $F(2, 135) = 0.95, p = .38, \eta_p^2 = .061$.

The participants' estimates of the prevalence of the attributed properties for each item type are provided in Table 1. These estimates are quite consistent for characteristic and striking predications in 'some' or generic form. With universally quantified predications, prevalence estimates are generally high, but the data for majority and striking predications are based on only two responses each and so can be ignored. The estimate for universally quantified characteristic predications is meaningful. In generic form, the mean prevalence estimate for these types of predications was .70, with a range from .52 to 1.0. This was substantially lower than the esti-

mate for their universally quantified counterparts (.97 with a range from .95 to 1.0). Clearly, people's acceptance of a characteristic predication in generic form is not based on their beliefs about the number of category members with the relevant property, and this is also the case for striking predications in generic form, for which the mean prevalence estimate was merely .36.

So, what are people doing when they agree with statements of the sort presented? For existentially quantified predications, the answer is simple: if even one member of a category has the attributed property, then the assertion is true, and this is reflected in both the agreement rates and the prevalence estimates that we obtained.

For universally quantified predications, majority and striking predicate types are virtually always rejected. In contrast, universally quantified assertions involving characteristic predicates were accepted 46% of the time, with a .97 prevalence estimate. With respect to the relation between agreement and prevalence rate estimates, the interesting cases are the predicate statements in generic form. For these items, people's acceptance rates seem independent of their prevalence estimates: .89 acceptance for characteristic predications with a .70 prevalence estimate, .68 acceptance for majority predications with a .80 prevalence estimate, and a .68 acceptance rate for striking predications with a mere .36 prevalence estimate (range .01 to 1.0). Clearly, people do not treat generic assertions as assertions about how many of the relevant individuals have the predicated feature.

Analogously, consider how people might be interpreting universally quantified characteristic predications such as 'all ducks lay eggs'. Even though the prevalence estimates for these assertions are high, a moment's reflection reveals that these estimates cannot be accurate. They may well have been inflated by participants' prior agreement with each item in universal form, and so do not accurately reflect prevalence beliefs. Instead, the estimates in this condition may reflect participants' trying to respond consistently: if an item had been agreed to in universal form, then the only consistent prevalence estimate would be very close to 1.0.

A more interesting question is, why do people agree to universally quantified predications in the first place? One possibility is that counterexamples are not accessible at the moment, i.e., they do not come to mind, just as they do not come to mind when interpreting generic characteristic assertions. In effect, when people agree to universal characteristic assertions, they are treating them as if they were generic assertions, perhaps to avoid the cognitive effort that would be required to process quantifiers (cf. Leslie, 2007, forthcoming b).

One might wonder, however, that when people assent to universally quantified characteristic predications, they are not quantifying over individuals, but rather over sub-kinds of the kind in question. On this hypothesis, people would assent to, for example, 'all ducks lay eggs', because they understand the universal quantifier to range over sub-types of ducks rather than individual ducks. 'All ducks lay eggs', then, may be

judged true because mallard ducks lay eggs and Peking ducks lay eggs, and so forth. We do not think that this hypothesis is likely, however. For one, most of our items simply do not have available sub-types (see Appendix). It is unlikely that our participants assented to 'all peacocks have blue tails' or 'all cardinals are red' because they took these statements to be quantifying over all the sub-types of these kinds, since it is unlikely that our participants have any idea what, if any, sub-kinds of peacocks or cardinals exist. Further, this hypothesis does not explain why we only found a significant degree of assent to universal statements when the predicate was characteristic. Why would they not employ this strategy for the other predicate types? If 'all lions have manes' is judged true because it is interpreted as quantifying over sub-types of lions, why would 'all lions claw people' not be open to a similar interpretation?

General Discussion

Generics express generalizations that are not fundamentally about how much or how many. Our subjects' judgments of their truth and falsity were not determined by their estimates of prevalence. Unlike quantifiers, generics do not depend on such quantitative considerations. Rather, the truth conditions of generics are sensitive to factors such as whether the predicate in question expresses a property that is striking, or characteristic of the kind.

Hollander, Gelman, and Star (2002) found that three-year olds treated both existentially and universally quantified statements as though they were generics. Four-year olds performed as did adults, discriminating between these two statement types. In light of these data and others, Leslie (2007) hypothesized that that generics give voice to cognitively primitive generalizations, while quantified statements give voice to more cognitively complex ones. That is, she hypothesized that quantified statements are more cognitively demanding to process than generics. Thus, Hollander, Gelman, and Star's three-year olds were falling back on the less taxing generic interpretation instead of attempting to process the more demanding explicit quantification.

Leslie further speculates that, under some circumstances, adults may be susceptible to similar errors. We found this to be true in the case of universally quantified characteristic predications. When faced with these statements, adults have a tendency to judge the statement to be true if the corresponding generic statement is true. This, we conjecture, is because, like the three-year olds, adults are incorrectly relying on the less taxing generic interpretation. That is, instead of evaluating the universally quantified statement, they are evaluating the corresponding generic. This is what leads them to erroneously assent to these false universal claims. We plan to test this hypothesis explicitly in future work.

Acknowledgments

We are grateful to Phil Johnson-Laird, Adele Goldberg, Geoffrey Goodwin, Louis Lee, Carlos Santamaria, and Axit Fumero for their helpful comments.

References

- Asher, N., & Morreau, M. (1995). What some generic sentences mean. In *The generic book* (p. 330-339). Chicago: Chicago University Press.
- Carlson, G. (1997). *Reference to kinds in english*. Unpublished doctoral dissertation, University of Massachusetts, Amherst.
- Dayal, V. (1999). Bare np's, reference to kinds, and incorporation. In T. Matthews & D. Strolovich (Eds.), *Proceedings of SALT IX* (p. 34-51). Ithaca, NY: Cornell University.
- Gelman, S. A. (2003). *The essential child*. New York: Oxford University Press.
- Gelman, S. A., Geotz, P., Sarnecka, B., & Flukes, J. (forthcoming). Generic language in parent-child conversations. *Language Learning and Development*.
- Greenberg, Y. (2003). *Manifestations of genericity*. New York: Routledge.
- Hollander, M. A., Gelman, S. A., & Star, J. (2002). Children's interpretation of generic noun phrases. *Developmental Psychology*, 36(6).
- Krifka, M., Pelletier, F., Carlson, G., Meulen, A. ter, Chierchia, G., & Link, G. (1995). Genericity: An introduction. In *The generic book* (p. 1-125). Chicago: Chicago University Press.
- Leslie, S. J. (2007). *Generics, cognition, and comprehension*. Unpublished doctoral dissertation, Princeton University.
- Leslie, S. J. (forthcoming a). Generics: Cognition and acquisition. *Philosophical Review*.
- Leslie, S. J. (forthcoming b). Generics and the structure of the mind. *Philosophical Perspectives*.
- Pelletier, F., & Asher, N. (1997). Generics and defaults. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language*. Cambridge, MA: MIT Press.
- Prasada, S., & Dillingham, E. (2006). Principled and statistical connections in common sense conception. *Cognition*, 99(1), 73-11.
- Roeper, T., Strauss, U., & Pearson, B. Z. (2006). The acquisition path of the determiner quantifier every: Two kinds of spreading. In T. Heizmann (Ed.), *Current issues in first language acquisition, university of massachusetts occasional papers* (Vol. 34, p. 97-129). GLSA.

Appendix

Table 2: Sentence stimuli in generic form, separated by predicate-type.

Characteristic	Majority	Striking
Ducks lay eggs	Cars have radios	Mosquitoes carry malaria
Pigs suckle their young	Barns are red	Ticks carry Lyme disease
Lions have manes	Shoes have laces	Sharks attack swimmers
Peacocks have beautiful tails	Books have chapters	Tigers claw people
Cardinals are red	Jackets have zippers	Rottweilers maul children
Horses give live birth	Shirts have collars	Pit bulls attack people
Moose have antlers	Clocks are round	Lions eat people
Bees gather honey	Radios have dials	Birds carry avian flue
Elephants have tusks	Trumpets are loud	Cigarettes cause cancer