

Sarah-Jane Leslie
Princeton University

Generics Articulate Default Generalizations

Introduction: Our Judgments of Generics

Generics are sentences that express generalizations about kinds, such as “tigers are striped,” “ducks lay eggs,” and “ticks carry Lyme disease.”¹ Quantified sentences such as “all tigers are striped,” “most barns are red”, and “some ticks carry Lyme disease” can also express generalizations about kinds, but unlike generics, they convey information about how many members of the kind have the property in question. Quantified statements have proved highly amenable to semantic analysis – for example, “all tigers are striped” is true just in case the set of tigers is included in the set of striped things. More generally, quantified statements depend solely on straightforward content-neutral information – of the sort that can be represented in set-theoretic terms – and as a result can be readily analyzed using the standard tools of formal semantics (e.g. Barwise & Cooper, 1981).

In contrast to quantified statements, generics have proved difficult for semanticists to analyze. As a simple illustration of why this is so, consider

the generics “ducks lay eggs” and “ducks are female”. There are more female ducks than there are egg-laying ducks (since some female ducks are infertile, immature, and so on), yet the former statement seems intuitively true, while the latter does not. Similarly, consider “ticks carry Lyme disease”; this would seem to be clearly correct, yet only one percent of ticks carry the disease. Conversely, however, “books are paperbacks” does not seem to be correct, yet over eighty percent of books are paperbacks.

In Leslie (2007, 2008), I suggested that our judgments of generics may be sensitive to non-quantitative, content-based factors. In particular, I proposed that our judgments are sensitive to whether the property in question is *characteristic of the kind* – e.g., does the property specify how the kind reproduces? How it nurtures its young? It’s identifying salient physical features? – and also to whether the property in question is *striking* – e.g., dangerous, threatening, appalling; the sort of property about which one would wish to be forewarned. If a property is either characteristic or striking, then the generic may well be accepted at low prevalence levels – hence the disposition to accept “ducks lay eggs” and “ticks carry Lyme disease” respectively. I also proposed that our judgments may be sensitive to the nature of the *exceptions* to the generic claim. That is, amongst the members of the kind that fail to have the predicated property, it matters *how* they fail to have the property – in particular whether they simply lack the property, or whether they have an equally salient, concrete, positive property *instead*. The non-infected ticks simply *do not carry Lyme disease* – they do not have an alternative property instead. However, the books that are not paperbacks are instead hardcover. Consider also “people are right-handed,” “elementary school teachers are female” – in both cases the predicated property is highly prevalent, yet the generic is intuitively false. But again there are ‘positive counterinstances’ – people who are not right-handed are instead *left-handed*; elementary school teachers who are not female are *male*. Intuitively, it feels as though we would be ‘overlooking’ these hardcover books, these left-handed individuals, and these male school teachers if we accepted the generics in question. Thus, I proposed that generics are likely to be rejected even if the property is highly prevalent, if there are these ‘positive counterinstances’. (For a more detailed presentation of these aspects of the proposal, see Leslie, 2007, 2008, in press a; for a discussion of social prejudice in light of this account, see Leslie, in press b.)

In this earlier work, I relied on my own evaluations of the generics in question, and on evaluations that have been made in the literature by linguists and philosophers for a number of years now (e.g. Carlson, 1977; Carlson & Pelletier, 1995; Lawler, 1973). That is, I took for granted that “ticks carry Lyme disease” can be true even if few ticks carry the disease; this was one key point I sought to explain. Recently, several experiments have been

conducted to test whether naïve participants also share the theorists' intuitions that, e.g., a generic such as "ticks carry Lyme disease" can be true even though very few ticks carry the disease. Yet one might wonder what the value of conducting such empirical tests could possibly be: haven't the theorists already established the conclusion that the empirical work will surely support? That is, don't we already *know* that "ticks carry Lyme disease" can be true despite the low percentage of infected ticks?

Empirical work in this regard is valuable, even if it is unlikely to overturn orthodoxy. For one, even if it is *unlikely* to overturn orthodoxy, one can never be certain. Suppose, for example, that false beliefs about the percentage of ticks that carry Lyme disease abound in the population: most people believe that approximately 80% of ticks are infected. Suppose further that people only reliably accept generics when they believe the prevalence of the property to be high. In this scenario, the generic sentence "ticks carry Lyme disease" is reliably asserted, but only because of false beliefs about prevalence. A theorist may grow up hearing "ticks carry Lyme disease" – may read the sentence repeatedly in the media – and so hear the sentence as true. Later, while conducting work on generics, the theorist may learn that in fact very few ticks carry Lyme disease, but since her belief in the truth of "ticks carry Lyme disease" has been built up over many years, she may still remain confident that the sentence is true, and so write that "ticks carry Lyme disease" can be true despite few ticks carrying the disease. However, all this scenario would reflect is that our beliefs in the truth of generics are not easily overturned, despite being in the general case driven solely by underlying beliefs about prevalence. If it is only the theorists who accept generics at low prevalence levels, this should give us pause for thought – and whether such a situation obtains cannot be determined on the basis of theorist intuition.

In this case, however, results from recent empirical studies comport with the observations made by linguists and philosophers. For example, Amanda Brandone, Andrei Cimpian, Susan Gelman and I found that adults are significantly more likely to accept sentences such as "birds lay eggs" (or "lions have manes", or "pigs give milk their young") than they are to accept sentences such as "birds are female" (or "lions are male", or "pigs are female"), despite appreciating that there are comparable numbers of female birds (or male lions, or female pigs) and egg-laying birds (or maned lions, or nursing pigs) (Brandone, Cimpian, Leslie & Gelman, in press). That is, for a range of examples, participants understand that only one sex of the kind has a particular characteristic property, yet they are far more likely to accept a generic that predicates the characteristic property than they are to accept a generic that predicates belonging to that sex.

Sandeep Prasada, Sangeet Khemlani, Sam Glucksberg and I further examined participants' judgments of a range of generics, including generics

that predicate non-prevalent but strikingly dangerous properties (e.g. “ticks carry Lyme disease”, “sharks attack swimmers”, “mosquitoes carry malaria”), and generics that predicate prevalent properties, but have ‘positive counterinstances’ (e.g., “books are paperbacks”, “Canadians are right-handed”, and “teachers are female”). We also elicited prevalence estimates from participants – that is, participants were asked to judge what percentage of the kind has the property in question. As predicted, participants gave much higher prevalence estimates for the *positive counterinstance* items than for the *strikingly dangerous* items, yet were much more likely to endorse the *strikingly dangerous* generics than the *positive counterinstance* generics (Prasada, Khemlani, Leslie & Glucksberg, submitted).

It is sometimes suggested that generics such as “ticks carry Lyme disease” are accepted because, even though few ticks carry it, ticks are the *only* creatures that do so. That is, even though few ticks carry Lyme disease, Lyme disease carriers are overwhelmingly ticks. Thus, perhaps we do not need to appeal to content-based factors such as whether the property is dangerous or not; perhaps it is simply that *if* something carries Lyme disease, *then* it is almost certainly a tick. Such a measure would be the ‘inverse’ of prevalence: prevalence can be modeled as the conditional probability of carrying Lyme disease *given that* one is a tick, while this measure would be the conditional probability of being a tick *given that* one carries Lyme disease. In the psychology literature, this measure is known as *cue validity*.

While there are some examples, actual and hypothetical, that tell against this possibility (see, e.g., Leslie, 2008), it is difficult to rule out definitively while relying solely on our intuitions as pure theorists of generics. However, it is easy to evaluate this hypothesis in the laboratory. To this end, Prasada et al. also elicited cue validity ratings in our study, and so we were able to control for cue validity across their different categories of generics. That is, the items in the *strikingly dangerous* category had the same average cue validity ratings as the items in the *positive counterinstance* category, yet the above results obtained nonetheless. If the *strikingly dangerous* generics were accepted because they had high cue validity ratings, this would not have been possible; the *positive counterinstance* generics had the same cue validity ratings, yet they were not accepted – despite being associated with much higher prevalence estimates. Cue validity cannot explain why these striking generics were accepted at low prevalence levels.

The category of *strikingly dangerous* generics in the Prasada et al. study was constructed so as to contain generics that predicate dangerous yet non-prevalent properties. That it was possible to do so – that is, to construct a range of generics that predicate strikingly dangerous properties and are accepted despite having low associated prevalence estimates and mid-range cue validity estimates – lends indirect support to the hypothesis in Leslie

(2007, 2008). However, the support is only indirect, since it does not involve directly testing the impact of a property's being dangerous on people's judgments. Recent work conducted by Andrei Cimpian, Amanda Brandone, and Susan Gelman, however, examined precisely this question (Cimpian, Brandone & Gelman, 2010). In their experiment, participants were told about a novel animal kind, e.g. 'lorches'. They were then told that a certain percentage of lorches had a particular property, e.g., purple feathers, then they were given some further information about the feathers. Crucially, in some cases, this further information described a dangerous property— e.g., the feathers were poisonous to the touch – whereas in other cases it was anodyne. The participants were then asked to evaluate a generic attributing the property to the kind. The experimenters found that, at low prevalence levels, participants were significantly more likely to accept the generic if the property in question was dangerous, thus lending support to the hypothesis that our judgments of generics are sensitive to factors such as whether the property being predicated is strikingly dangerous (Leslie, 2007, 2008).

One point to note is that in the Cimpian et al. study, the non-dangerous properties were still inherent, biologically-based properties (e.g., purple feathers), and so may have been understood to be characteristic properties of the kind.² That is, the rate of acceptance of these non-dangerous but inherent properties may already have been high relative to more 'accidental' properties. To examine this, Cimpian et al. adapted their paradigm to include properties that would likely be interpreted as accidental, non-characteristic properties – e.g. muddy feathers, broken legs. In the case of such properties, participants almost never accepted the generic at low prevalence levels; instead they only accepted the generic when a strong majority of the kind possessed the property, which is consistent again with the predictions made in Leslie (2007, 2008). Susan Gelman and Paul Bloom (2007) also found converging evidence of the impact of the inherent and biological vs. the adopted and accidental on adults' judgments of generics.

Recent empirical work thus suggests that adults' judgments of generics are sensitive to rich, content-based factors such as whether the property in question is strikingly dangerous, whether it is an inherent characteristic property of the kind, or whether it is merely accidentally possessed by members of the kind. The impact of these factors does not appear to be reducible to abstract, quantitative measures such as prevalence and cue validity.

The Psychological Significance of Generics

The question one faces at this juncture is whether the foregoing observations and findings simply amount to the following: there is a class of sentences, evaluations of which are sensitive to a number of surprising factors. This

would be of some interest to some theorists, but may not be terribly gripping at the end of the day. Is this as far as the study of generics can take us, or might they hold interest from the perspective of the nature of cognition more generally?

Taking a step back for a moment, let us note that one of the most important tasks a learner faces is how to arrive at general information that applies to situations and instances that they have not yet encountered – how to move from the *particular* to the *general*. If one is unable to generalize information in this sense, one's knowledge will be forever bound to the here-and-now – every new situation must be encountered without the support of prior knowledge. As it happens, if one touches a hot stove and burns one's hand, one will be careful about touching that stove *in the future*, and one will also be careful about touching *other stoves*. This sort of learning involves generalization in the sense I am using it here – in the first instance, generalization from one time to future times; and in the second, from one instance of a kind to other instances of a kind, namely other stoves. There is a sense in which the capacity to generalize is the backbone of all learning – without it, one could only catalog discrete and seemingly unrelated events.

It is perhaps unsurprising, then, that infants in the first year of life are able to form general judgments (e.g. Baldwin, Markman & Melartin, 1993; Graham, Kilbreath & Welder, 2001). They are, for example, able to form expectations concerning the properties of novel instances of a kind or category on the basis of experience with other instances. Thus they expect that, if one item makes a rattling noise when it is shaken, perceptually similar items will do so too. Infants are not simply trapped in the here-and-now – they can extract general information from a situation and apply it to a novel one.

There must therefore be something about the infants' cognitive capacities that allows them to do this, something in their cognitive makeup that permits them to form and exploit general judgments – a cognitive mechanism that forms general judgments. We might then call the generalizations that are formed by this mechanism *cognitively fundamental generalizations*. This mechanism would represent our most natural and default way of forming general judgments, and would produce the sorts of general judgments that we make from our earliest days.

An interesting question, then, is whether language later provides us with a way of giving voice to these generalizations that we have been making all along. It is not *guaranteed* that natural language must provide us with such a vehicle, but it is reasonable to suppose it might (and would be quite surprising if it did not). We can thus ask the question: of the many ways we have of linguistically expressing generalizations, which form articulates our cognitively fundamental generalizations? That is, if an adult forms a general

judgment using this fundamental, default mechanism of generalization and wishes to communicate this belief using natural language, what linguistic form will be used? The answer I propose here and elsewhere is: the generic form.

One observation that can be made in support of this hypothesis is that, while there is a word in English “all” that signals a universal generalization, a word “some” that signals an existential one, and so on for all the quantifier words – both determiners and adverbs – there is no word “gen” in English that signals a generic generalization. That is, we do not say “gen tigers are striped” in the way we say “all tigers are striped” or “most tigers are striped.”³ This is not an isolated fact about English either: there is no known language that has a dedicated, articulated generic operator – no word “gen” that uniquely signals that a generalization is a generic one. While languages differ dramatically in how generic claims are syntactically manifested, one does not find phonological forms that are dedicated to the expression of generic claims – generics are always unmarked in this way (Dahl, 1985). This would seem a puzzling fact: why should this be the case? And even if a language with a “gen” in it were to be found, we would still be left with the fact that the overwhelming majority of languages do not contain such a linguistic item – a linguistic quasi-universal that would still cry out for explanation.

Such an explanation is forthcoming if generics are language’s way of letting us give voice to cognitively fundamental generalizations (Leslie, 2007, 2008). If one wishes to interact efficiently with a system, and the system has a basic, default way of proceeding or performing a task, then one need only issue an explicit instruction to the system if one wishes it to *deviate* from this default way of proceeding. As an intuitive illustration, suppose that one is dealing with a recalcitrant teenager who by default never tidies his room. If on a particular occasion one does *not* wish for said teenager to tidy his room, one need not say anything – it would be a waste of breath, as it were, to say “please do not tidy your room today”, since if one says nothing, this is precisely what will happen. One need only speak up if one wishes the teenager to *deviate* from his default; then one must ask that the room be cleaned. Quantifiers, I propose, are articulated precisely because quantified generalizations are non-default generalizations. For the cognitive system to form a universal generalization it must be *told* to form a universal generalization, as it were – and this requires that there be a word in natural language like “all”. Such non-default generalizations must be explicitly *marked* – they are like the instruction to the teenage to clean his room. However, generic sentences contain no such marking – there is no item in them that specifies what *sort* of generalization is here being expressed. As a result, the cognitive system supplies its own default means of generalizing.

To explicitly instruct a system to operate according to its default is inefficient and unnecessary – akin to instructing the teenager *not* to tidy his room – hence the universal or near universal absence of a word “gen”.

The hypothesis that generics articulate cognitively default generalizations, then, allows us to explain this otherwise puzzling cross-linguistic fact. Further, it issues in a number of empirical predictions. For example, if processing quantified statements requires the cognitive system to deviate from its generic default, we would expect that sometimes this default would be inadequately suppressed, and quantified statements would be processed as though they were generics. This might show up in a number of ways – in evaluation tasks, in memory tasks, in reasoning tasks. We might further predict that this tendency would be found in adults, but that it would be most pronounced in young children – since in general, children are less likely than adults to be able to successfully inhibit default processes. We would also correspondingly predict that generics would prove quite easy for young children to acquire – easier to acquire than quantifiers.

In case these predictions seem in any way obvious, consider the following contrasting point of view. Quantified statements depend solely on clean, quantitative information, of the sort that is theoretically highly tractable. Generics, on the other hand, appear to be sensitive to all sorts of messy, complicated, content-based factors, as the discussion in the previous section indicates. Surely it would be easier to acquire and process the former than the latter! Further, the fact that there is no word “gen” should make the learner’s task even harder – she is not even presented with an articulated word whose meaning she can try to figure out. So we have something elegant, tractable, and articulated like “all” on the one hand, and the opaque, elusive and mysterious generic on the other. It would be quite natural to suppose that a child would find the former easier to acquire than the latter, and that the latter would impose greater processing demands so that if anything children and adults alike should default *from* the generic *to* the quantified, rather than vice-versa.

Generics in Acquisition

Toddlers begin to produce generics at approximately 30 months of age, which is the first time that they reliably have the requisite background syntactic capabilities (e.g. sufficiently many words per utterance, plurality, etc; Gelman, 2010; Gelman, Goetz, Sarnecka & Flukes, 2008). By the time children are between 3 to 4 years of age, they produce generics as frequently as adults do – which is not something that is true for all parts of language by a long shot (Gelman et al., 2008). Thirty-month-old toddlers are further able to draw different inferences from generics than from specific statements. That is, if they are shown some members of a novel animal kind performing

an action such as drinking milk and told either “blicks drink milk” or “these blicks drink milk”, they are more likely to generalize the action of drinking milk to novel blicks if they hear generic language (Graham, Nayer & Gelman, 2010). Thus, 30-month-olds understand that generics express generalizations, in contrast to specific statements, and further are able to deploy this fact in drawing inferences.

Two-year-olds also understand that generics tolerate exceptions. For example, Susan Gelman and Lakshmi Raman (2003) showed young children pictures of familiar animals that were exceptions to a generic generalization – the children might have been shown, e.g., a picture of penguins. The children were then either asked “do these birds fly?” or “do birds fly?” If young children understand that generics tolerate exceptions, then they should answer ‘no’ to the specific question but ‘yes’ to the generic one, and the experimenters found that toddlers as young as 30 months were able to do just this. Several other studies have also found that preschoolers understand that generics tolerate exceptions in this way (e.g., Chambers, Graham & Turner, 2008; Gelman & Bloom, 2007; Gelman, Star & Flukes, 2002).

What of preschoolers’ comprehension of ‘troublesome’ generics – the sorts of generics that have caused difficulty for semanticists studying these sentences? Amanda Brandone, Andrei Cimpian, Susan Gelman and I have begun to research this issue. The question we posed was: if a preschooler knows that, e.g., only girl birds lay eggs, how will she evaluate “birds lay eggs” and “birds are girls”? If a child has the relevant background knowledge, then she will at least implicitly believe that there are as many girl birds as there are egg-laying birds. Thus, if preschoolers’ judgments, unlike adults’, are driven by considerations of prevalence, then such a child should be equally likely to endorse the two generics. If, however, preschool children are sensitive to the same factors as adults, then they should be more likely to accept “birds lay eggs” than “birds are girls”. We used a range of items that included methods of reproduction, salient identifying physical features, and means of nurturing the young, and we only considered data from trials in which the preschooler in question knew that only one gender of the animal kind in question had the property. We found that preschoolers, like adults, were far more likely to endorse generics that attribute a characteristic property to the kind than they were to endorse generics that predicated belonging to a given gender.⁴ This suggests that even for preschoolers, generics are not simple expressions of beliefs about prevalence (Brandone et al., in press).

While a more complete investigation into preschoolers’ understanding of generics is needed (and indeed is underway), the available evidence suggests that, like adults, they understand that generics a) tolerate exceptions, and b) do not depend solely on considerations of prevalence. The

data so far also indicate that generics are acquired easily and early. We might ask, then, how the acquisition of generics compares to the acquisition of quantifiers. Consider, for example, the quantifier “most”, which one might suppose to be the closest cousin of generics, since it too could be said to ‘tolerate exceptions’. Is “most” acquired in a comparable time frame?

The answer to this question appears to be a resounding ‘no’. Preschoolers appear to find “most” very difficult to process, even on simple tasks. For example, suppose there are six crayons, some of which are in a box, others of which are not. The child is then asked “are most of the crayons in the box?” Papafragou and Schwarz (2005/2006) found that preschool children found the task exceedingly difficult; there was not even a consistent pattern to their responses. Further, even 6- to 8-year-olds – who are quite old in terms of language acquisition – differed from adults on the task. Papafragou and Schwarz’s finding has subsequently been replicated (e.g., Barner, Chow & Yang, 2009; Leslie & Gelman, 2012, experiment 2). (Some researchers have suggested that preschoolers have *some* understanding of “most”, which manifests itself if the task is made even easier (Haldberda, Taing & Lidz, 2008) – however, even if preschoolers have some *competence* with “most”, it is clear that it is minimally very difficult for them to process.)

In the case of the quantifiers “all” and “some”, however, even young preschoolers are able to pass a comparable task – that is, they can say whether, e.g., some/all of a small set of crayons before them are in a box or not (e.g., Barner et al., 2009). There is a sense, though, in which such scenarios do not require children to process full-blooded *generalizations*. They are not asked to make generalizations about crayons in general -- as a kind or class -- but only to form some beliefs about a set of crayons directly before them. The judgment that *all of these crayons are in the box* does not involve any expectation as to whether a newly introduced crayon will be in the box or not. It simply does not apply to any crayon that is not one of *these crayons*. The judgment here is not an open-ended, kind-wide, projectable generalization. (Note that generic generalizations can *only* be of this open-ended sort – there is no way of even formulating a generic that pertains only to *these 6 crayons*.) Developmental studies of the acquisition and processing of quantifiers, however, have mostly considered only these sorts of scenario: ones in which children are presented with a small set of items and then asked to evaluate quantified statements pertaining to *those particular items*. Let us then pose the question: How do preschoolers fare when asked to consider open-ended, kind-wide, projectable quantified generalizations?

Defaulting to the Generic

In a relatively early study, Michelle Hollander, Susan Gelman, and Jon Star (2002) tested preschooler’s comprehension of kind-wide quantified

statements. In their study, 3- and 4-year-olds and adults were asked a number of yes/no questions. The questions were either existentially quantified (e.g. “are some fires hot?”), universally quantified (e.g., “are all fires hot?”), or generic (e.g. “are fires hot?”). Importantly, for all three question types, their participants were asked to evaluate open-ended, kind-wide generalizations – as opposed to being asked questions about a specific and limited set of items before them. The questions covered a range of properties – some were had by all or almost all members of the kind, others by some but not all members of the kind, and others by no members of the kind.

In the case of responses to the generic questions, the experimenters found no developmental differences in the responses; that is, 3-year-olds, 4-year-olds and adults gave statistically indistinguishable responses to the questions when they occurred in generic form. In the case of the quantified questions, however, they found clear developmental changes, with all three age groups giving different patterns of responses. What was most intriguing about the data, though, were the response patterns provided by the 3-year-olds: this youngest group of participants responded in the same way regardless of whether the question was universal, existential, or generic. This pattern of responses was the same as the pattern of responses provided by the 4-year-olds and the adults to the generic questions. Thus an intriguing possibility emerged: might the 3-year-olds be interpreting the quantified statements as though they were generics? Such a possibility is precisely what one would expect to find if generics articulate cognitively default generalizations. Processing and evaluating kind-wide quantified statements would require inhibiting the generic default and engaging in a more taxing cognitive process – people, especially young children, may not always do so successfully, and so would be susceptible to treating quantified statements as though they were generics.

It should be noted that, even if the 3-year-olds found the quantified statements difficult to process, there are many, many possible alternative response patterns that they might have produced. They could, for example, simply have said ‘yes’ to every quantified question, or ‘no’, or given responses with no clear pattern to them. (Such responses are frequently encountered in developmental work – if children do not understand a task or a question, they often respond at random.) But this is not what happened here; instead they gave the same responses to the quantified questions as they gave to the generic questions – which in turn were indistinguishable from the responses adults gave to those generic questions. Further, the 4-year-old children also differed from adults in their responses to the quantified – but not the generic – questions. In particular, their response patterns were intermediate between the 3-year-olds’ generic pattern, and the adults’ pattern. Thus, it is possible that, in some trials, the 4-year-olds were also interpreting

quantified statements as generics.

After they had asked the children these kind-wide questions, the experimenters asked the children universal and existential questions that pertained to specific sets of items before them (e.g. “are all of these crayons in the box?”). Here, even the youngest children were able to answer correctly – thus, they did not lack a basic competence with “all” and “some”, but rather appeared to be unable to properly process them when confronted with an open-ended, kind-wide generalization.

Hollander et al.’s findings were recently replicated and extended to Mandarin-speaking preschoolers (Tardif, Gelman, Fu & Zhu, 2010). The results from the Mandarin-speaking population were even more dramatic: even Mandarin-speaking 4-year-olds gave indistinguishable responses to universals, existentials, and generics. Again, the researchers found no developmental differences in answers to the generic questions – the Mandarin-speaking preschoolers responded just like adults to the generic questions, but also gave this generic response pattern to the quantified questions.

The results from Hollander et al. (2002) and Tardif et al. (2010) suggest that, even if young children are able to understand “all” and “some” when used in the context of a small set of items (e.g., “all of these crayons”), they nonetheless have difficulty processing kind-wide quantified statements, and may fall back on their interpretation of the relevant generic instead. However, in the Hollander et al. and Tardif et al. studies, children had to rely on their background knowledge of the kind to evaluate the kind-wide statements, whereas they had the all the relevant instances before them (e.g. the 6 crayons) when asked to evaluate quantified statements concerning a small set of items. To what extent is this asymmetry important? Certainly, relying on memory increases the cognitive demands associated with performing a task, and so perhaps the children only defaulted to the generic because of the difficulty of the task. However, it would constitute an even more dramatic illustration of the power of the generic default if they treated quantified statements as generics even when they do not have to rely on memory to complete the task.

Susan Gelman and I (Leslie & Gelman, 2012, experiment 4) presented 4-year-olds and adults with pictures of a novel animal kind, introduced with a novel label e.g., “these are gorps.” In each case, there were 8 individual members of the novel kind, and we ensured that the child understood that each of the 8 was indeed a member of the kind. The 8 depicted instances were identical, except for a specific target feature – e.g. curly hair – that was had by only 6 of the 8. Each child was then asked a question about the novel animals: the question was either a generic (“do gorps have curly hair”), a ‘specific’ universal (“do all *of these* gorps have

curly hair?”), or a kind-wide universal (“do all gorps have curly hair?”). The 4-year-olds, like the adult, tended to agree to the generic question, saying ‘yes’ to it on 83% of trials. They also performed well with the specific universal, correctly denying that *all of these gorps* have curly hair on 79% of trials. However, when asked the kind-wide universal question, the 4-year-olds responded ‘yes’ on 51% of trials, which was significantly higher than the rate at which they accepted the specific universal (21% of trials). The children had the pictures of the gorps in front of them during questioning, and so had the falsifying counterexamples – the hairless gorps – before them when evaluating the kind-wide universal. They clearly understood that not all *of these gorps* have curly hair, yet they were not consistently able to correctly reject the kind-wide universal, even under these circumstances. Again, these otherwise puzzling data are readily explained by the hypothesis that preschoolers sometimes substitute their evaluation of a generic for their evaluation of a kind-wide universal. The preschoolers replied that *all gorps have curly hair* because they did not consistently evaluate the universal as a universal, but rather on some trials assimilated it the generic *gorps have curly hair*. We are currently investigating this tendency at other prevalence levels, with different age groups, and with other quantifiers (in particular, “some”).

Thus, several experiments have turned up findings that are precisely what one would expect if generics do indeed express cognitively default generalizations, in contrast to quantifiers. The data would be otherwise quite puzzling and surprising – e.g., why are 4-year-olds able to use hairless gorps to consistently reject “all of these gorps have curly hair” but not to consistently reject “all gorps have curly hair”? However, the data reviewed thus far all concern preschool children. On the generics-as-defaults hypothesis, one would expect this tendency to be most pronounced in young children, but one would also expect to find evidence of it in adults. Thus, we might pose the question: do adults ever interpret quantified statements as generics?

Sangeet Khemlani, Sam Gluckberg and I investigated precisely this question (Leslie, Khemlani & Glucksberg, 2011). In particular, we were interested to learn whether adults would ever accept universals such as “all ducks lay eggs”, despite knowing that, e.g., male ducks do not lay eggs. We found that adults have a robust tendency to do precisely this. For example, if asked to first evaluate universals such as “all ducks lay eggs” and then later to evaluate statements such as “male ducks lay eggs”, participants who correctly rejected the latter were just as likely to have *accepted* the universal as they are to have *rejected* it. One might wonder what would happen if the ordering was reversed – if participants *first* rejected “male ducks lay eggs” (as the vast majority of them do), and were then subsequently asked to evaluate the universal “all ducks lay eggs”. Reversing the order of the blocks

in this way reduced the tendency to accept the universal, but it did not eliminate it: on almost 20% of trials, adults still went on to accept “all ducks lay eggs” despite having just rejected “male ducks lay eggs” (Leslie et al., 2011, experiment 3). (This percentage of trials was significantly greater than zero, and significantly greater than the percentage of trials on which participants gave other incorrect response patterns.) Again, these data are difficult to explain, unless one supposes that the adult participants were *not really evaluating* the universals at all – but were instead evaluating the corresponding generic “ducks lay eggs”.⁵ Meredith Meyer, Susan Gelman, and Sarah Stilwell tested adults’ evaluations of universally quantified statements and also found evidence that they were interpreting them as though they were generics. Importantly, they also found that the tendency to do so becomes stronger when adults are under time pressure (Meyer, Gelman & Stilwell, 2010).

Converging evidence is also to be found in studies on adult reasoning. Several studies have found that adults make specific reasoning errors when dealing with kind-wide universal statements – errors that would be readily explained if they were treating those universals as generics. For example, Steven Sloman (1993, 1998) investigated adults’ evaluations of arguments that involve the quantifier “all”, and found that their evaluations did not conform to the logic of universal quantification. His participants judged that, e.g., argument A is stronger than argument B, despite judging that reptiles are indeed animals:

- (A) All animals use norepinephrine as a neurotransmitter; therefore all mammals use norepinephrine as a neurotransmitter
- (B) All animals use norepinephrine as a neurotransmitter; therefore all reptiles use norepinephrine as a neurotransmitter

Such an assessment would be readily understandable if Sloman’s participants were in fact interpreting these universals as generics, since reptiles maybe be naturally thought to be more likely than mammals to be exceptions to the generic “animals use norepinephrine as a neurotransmitter”.

Similarly, Martin Jönsson and James Hampton found that adults consistently judged universals such as “all ravens are black” to be more likely to be true than universals such as “all young jungle ravens are black” – despite the fact that young jungle ravens are a subset of ravens, and so if the latter universal is false then so is the former (Jönsson & Hampton, 2006). If one replaces the universals with generics, however, then this judgment is very reasonable: since generics tolerate exceptions, it is possible for “ravens are black” to be true even if young jungle ravens are not. Since the latter are an unknown quantity, one could reasonably be more confident in the more inclusive statement. Thus, again, if Jönsson and Hampton’s participants were interpreting the universals as generics, the findings would be readily

explicable.

Further evidence for the generics-as-default hypothesis also emerged from a recent memory study. Since much of the foregoing evidence derives from evaluation tasks, Susan Gelman and I wished to use a measure that would tap into a distinct set of behavioral responses. To do so, we tested preschoolers' and adults' recall of quantified statements and generics (Leslie & Gelman, 2012). We predicted that 3-year-olds and 4-year-olds and adults would all tend to recall quantified statements as generics. We indeed found a marked tendency in all three age groups to recall statements quantified with "all" and "most" as generics. Adults and 4-year-olds generally recalled "some" statements correctly, however 3-year-olds tended to recall even these existentials as generics. We then examined 3-year-olds' recall of "no" statements – for example "no butterflies taste with their mouths". One possibility was that 3-year-olds may have simply been 'dropping the first word' – namely the quantifier – and so employing a 'dumb' strategy which would not lend any particular support to the generics-as-defaults hypothesis. If this was the case, then the 3-year-olds should recall "no butterflies taste with their mouths" as "butterflies taste with their mouths" – i.e. "no" statements would be recalled as positive generics. However, 3-year-olds almost never recalled "no" statements as positive generics – instead, "no" statements were frequently recalled as negated generics, e.g. "butterflies don't taste with their mouths". Thus the youngest children recalled quantified statements as generics even when doing so required introducing a negation into the sentence, thereby adding to the growing body of evidence in favor of the hypothesis that generics express cognitively default generalizations.

Conclusion

The available empirical evidence would seem to support the hypothesis that generics articulate cognitively fundamental, default generalizations, while quantifiers articulate more taxing and sophisticated ones. The hypothesis pulls together and explains a range of otherwise puzzling results from a variety of experimental paradigms across diverse age groups. If we combine this idea with the observations concerning the non-quantitative nature of generic generalizations, we are led to the following conclusion: our most basic way of forming general judgments is not fundamentally driven by considerations of prevalence or cue validity as many psychologists have supposed (e.g. Rosch, 1978), but rather is sensitive to rich, content-based factors, such as whether the property being generalized is dangerous, or characteristic of the kind in question.

References

Baldwin, D. A.; Markman, E.; Melartin, R. (1993). Infants' ability to draw inferences

- about nonobvious object properties: Evidence from exploratory play. *Child Development* 64: 711–28.
- Barner, D.; Chow, K.; Yang, S. (2009). Finding one's meaning: A test of the relation between quantifiers and integers in language development. *Cognitive Psychology* 58: 195-219.
- Barwise, J.; Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy* 4: 159-219.
- Brandone, A.; Cimpian, A.; Leslie, S. J.; Gelman, S. A. (in press). Do lions have manes? For children, generics are about kinds not quantities. *Child Development*.
- Carlson, G. (1977). *Reference to kinds in English*. Doctoral dissertation, University of Massachusetts, Amherst.
- Carlson, G.; Pelletier, F.J. (1995). *The Generic Book*. Chicago: University of Chicago Press.
- Chambers, C.G.; Graham, S.A.; Turner, J.N. (2008). When hearsay trumps evidence: How generic language guides preschoolers' inferences about unfamiliar things. *Language and Cognitive Processes* 23: 749-766.
- Cimpian, A.; Brandone, A. C.; Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science* 34: 1452–1482.
- Dahl, O. (1985). *Tense and aspect systems*. Oxford: Blackwell.
- Gelman, S. A. (2010). Generics as a window onto young children's concepts. In F. J. Pelletier (ed.). *Kinds, things, and stuff: The cognitive side of generics and mass terms*. (New Directions in Cognitive Science v. 12.) New York: Oxford University Press.
- Gelman, S. A.; Bloom, P. (2007). Developmental changes in the understanding of generics. *Cognition* 105: 166-183.
- Gelman, S.A.; Goetz, P.J.; Sarnecka, B.S.; Flukes, J. (2008). Generic language in parent-child conversations. *Language Learning and Development* 4: 1-31.
- Gelman, S. A.; Star, J.; Flukes, J. (2002). Children's use of generics in inductive inferences. *Journal of Cognition and Development* 3: 179-199.
- Graham, S. A.; Kilbreath, C. S.; Welder, A. N. (2001). Words and shape similarity guide 13-month-olds' inferences about nonobvious object properties. In *Proceedings of the Twenty Third Annual Conference of the Cognitive Science Society*: 352–57. Hillsdale, NJ: Erlbaum.
- Graham, S. A.; Nayer, S. L.; Gelman, S. A. (2011). Two-year-olds use the generic/non-generic distinction to guide their inferences about novel kinds. *Child Development* 82: 493-507.
- Gelman, S.A.; Raman, L. (2003). Preschool children use linguistic form class and pragmatic cues to interpret generics. *Child Development* 74: 308-325.
- Halberda, J.; Taing, L.; Lidz, J. (2008). The development of “most” comprehension and its potential dependence on counting-ability in preschoolers. *Language Learning and Development*, 4: 99-121.

- Hollander, M. A.; Gelman, S. A.; Star, J. (2002). Children's interpretation of generic noun phrases. *Developmental Psychology* 38: 883-894.
- Jönsson, M. L.; Hampton, J. A. (2006). The inverse conjunction fallacy. *Journal of Memory and Language* 55: 317-334.
- Lawler, J. (1973). "Studies in English generics." *University of Michigan Papers in Linguistics* 1.
- Leslie, S. J. (2007). Generics and the structure of the mind. *Philosophical Perspectives* 21 : 375-405.
- Leslie, S. J. (2008). Generics: Cognition and acquisition. *The Philosophical Review* 117: 1-49.
- Leslie, S. J. (in press a) *Generics and Generalization*. Oxford: Oxford University Press.
- Leslie, S. J. (in press b). The original sin of cognition: Fear, prejudice, and generalization. *Journal of Philosophy*.
- Leslie, S. J.; Gelman, S. A. (2012). Quantified statements are recalled as generics: Evidence from preschool children and adults. *Cognitive Psychology* 64: 186-214.
- Leslie, S. J.; Khemlani, S.; Glucksberg, S. (2011). All ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Language* 65: 15-31.
- Meyer, M.; Gelman, S. A.; Stilwell, S. M. (2011). Generics are a cognitive default: Evidence from sentence processing. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*. Boston, MA: Cognitive Science Society.
- Papafragou, A.; Schwarz, N. (2005/6). Most wanted. *Language Acquisition* 13: 207-251.
- Prasada, S.; Khemlani, S.; Leslie, S. J.; Glucksberg, S. (submitted). Conceptual distinctions amongst generics.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.) *Cognition and categorization*: 27-48. Hillsdale, NJ: Lawrence Earlbaum Associates.
- Sloman, S. A. (1993). Feature-based induction. *Cognitive Psychology* 25: 231-280.
- Sloman, S. A. (1998). Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology* 35: 1-33.
- Tardif, T.; Gelman, S. A.; Fu, X.; Zhu, L. (2011). Acquisition of generic noun phrases in Chinese: Learning about lions without an "-s". *Journal of Child Language* 30: 1-32.

¹ This paper will focus on bare plural generics, however English also allows for indefinite singular (e.g., "a tiger is striped") and definite singular generics (e.g., "the tiger is striped").

² Cimpian et al. also found that participants were more likely to accept the generic at low prevalence levels if they were told that the property in question was distinctive – i.e. that no other animals on the island had the property. One possibility is that this finding reflects an effect of cue validity on people's judgments. (Note that the possibility that cue validity affects judgments is different from the hypothesis that cue validity is *the explanation* – to the exclusion of the impact of danger -- of judgments of low-prevalence generics such as "ticks carry Lyme disease" – the latter is

inconsistent with the evidence that is currently available, while the former remains an empirical possibility.) That is, all else being equal, people may be *more likely* to accept a generic if cue validity is high. Another possibility is that this effect is limited to salient physical characteristics of the kind – the sorts of features that are extremely helpful to us in perceptually identifying whether something is a member of a given kind or not, even if the features are only had by some members – and that the effect of distinctiveness operates by way of making it more likely that the property in question is seen as *characteristic of the kind*. That is, rather than there being a general effect of cue validity on people’s judgments across the board, there may be a more limited effect of distinctiveness on identifying physical characteristics. This is an open empirical question at this point in time. (Again, this is a question that would be exceedingly difficult to decide by considering theorist’s intuitions about a small number of examples, but that can be readily resolved in the laboratory.)

³ It is sometimes proposed in conversation that we *do* have such a word, namely the adverb “generally”. However, inserting “generally” into a generic produces a change in meaning. For example, “books are generally paperbacks”, “Canadians are generally right-handed”, and “elementary school teachers are generally female” would all appear to be perfectly acceptable statements, while their generic counterparts are reliably rejected.

⁴ “Male” and “female” are not words that preschoolers are generally familiar with, hence the use of “boys” and “girls”. However, we wanted to be sure that the children did not think we were asking whether birds were *human* girls, so we conducted a follow-up study where we asked them to evaluate generics such as “mommy bears are girls”. Our participants readily endorsed such items. We were also concerned that there was a miss-match between the types of main verbs in the two sentences – the gender attribution items were all of the form “are boys/girls”, where as the characteristic items either contained action verbs or the verb “to have” (e.g. “lions have manes”). In a follow-up study, we converted all of the predicates to the form “are animals that ...” – e.g. “lions are animals that have manes” vs. “lions are animals that are boys”. The same results were obtained as in the original study.

⁵ One possible alternative explanation of these data could be that participants were interpreting the universal quantifier as ranging over subkinds of ducks, rather than individual ducks. That is, perhaps participants understood “all ducks lay eggs” to mean that mallard ducks lay eggs, and mandarin ducks lay eggs, and so on for each subkind of duck. There are several reasons to believe this was not the case. As a direct measure, we asked participants to paraphrase the statements they had been asked to evaluate; if a subkind interpretation could account for the high (close to 50%) rate of acceptance of the universals, one would expect that a reasonable percentage of the paraphrases would contain some indication of this. Crucially, participants who accepted a given universal should be more likely to use subkinds in their paraphrases than participants who rejected that universal. Instead, vanishingly few paraphrases had any reference to subkinds whatsoever, and there was no difference to be found along this dimension between participants who accepted vs. rejected a given universal (Leslie et al., 2011, experiment 2B). (The same applied to another possible explanation: namely that participants were understanding the universal to be domain restricted, e.g. to only be about female ducks. The paraphrase task found no evidence of this whatsoever.) Secondly, we ran a version of the study where participants were presented with population information designed to make a subkind interpretation very

unnatural. For example, participants were told: “There are 4.16 million ducks in the world. Evaluate the following statement: All ducks lay eggs”. The use of population information should make a subkind interpretation (or any kind of domain restricted interpretation) very odd, and yet we still found a robust tendency to accept the universal statements (Leslie et al., 2011, experiment 2A). Thirdly, Susan Gelman and I tested whether adults would be more likely to accept universal statements about novel kinds if a) not all individuals have the target property, but b) the kinds were presented to them as having salient subkinds which were such that each subkind could be said to generically have the property. That is, we manipulated the availability of a subkind interpretation of a universal, but did not find any increase in the (extremely low) rate at which adults accepted the universal as a result (Leslie & Gelman, 2012, experiment 4). Thus, there does not appear to be any evidence that participants are disposed to employ a subkind interpretation of these universals, and so this is very unlikely to be the explanation of the findings. Instead, we propose that, like young children, adults have some tendency to default to the generic when evaluating quantified statements.