28

Generics and Experimental Philosophy

ADAM LERNER AND SARAH-JANE LESLIE

Generics are statements that express generalizations about (members of) kinds, such as "dogs are mammals," "a tiger is striped," "the dinosaur is extinct," "kangaroos have pouches," and "mosquitoes carry the West Nile virus." Explicitly quantified statements such as "all dogs are mammals," "most tigers are striped," and "some mosquitoes carry the West Nile virus" also express generalizations about kinds, but unlike generics they carry information about how many members of the kind have the property in question. The truth conditions for explicitly quantified statements are well understood. For example, "all dogs are mammals" is true just in case the set of dogs is included in the set of mammals. Because the truth of an explicitly quantified statement depends solely on content-neutral information of the sort that can be represented in set-theoretic terms, theorists have had great success analyzing explicitly quantified statements using the standard tools of formal semantics (e.g., Barwise and Cooper 1981).

Theorists have had far less success analyzing the truth conditions of generics. Unlike explicitly quantified statements, the truth conditions of generics do not depend in any straightforward way on facts about how many members of the kind have the property in question. For example, while some generic statements (e.g., "dogs are mammals") seem to behave like universal generalizations (e.g., "all dogs are mammals") in requiring for their truth that every member of the kind have the property in question, many generics (e.g., "dogs have four legs") can remain true even if there are clear exceptions to the generalization. Other generics (e.g., "cars have radios") seem to be true because *most* members of the kind have the property in question, but it is neither necessary nor sufficient for the truth of generics more generally that most of the members of the kind have the property in question. For example, "mosquitoes carry the West Nile virus" is true even though less than 1% of mosquitoes carry the West Nile virus, and "books are paperbacks" is false even though the majority of books are paperbacks.

Philosophers of language have offered a number of theories intended to account for this puzzling pattern of data. In Section 28.1, we survey several semantic accounts of generics – however,

since our focus is on generics and experimental philosophy, this overview will be very brief and thus will not do full justice to these accounts. (For more detailed discussion, see Carlson and Pelletier (1995), Leslie (2008, 2012a), and others cited therein.) In Section 28.2, we briefly review empirical work that bears on these semantic accounts (for more detailed discussion of this work, see Leslie 2012b).

While generics constitute an interesting linguistic phenomenon worthy of study in their own right, the study of generics also has wide-ranging implications for questions beyond the philosophy of language, including questions in social psychology and cognitive science more generally. In Section 28.3, we review empirical work on the relationship between generics and cognition. The general theme of this section is that existing empirical work strongly supports the *generics-as-defaults* hypothesis: the hypothesis that generics reflect a cognitively default, fundamental mode of generalizing in humans. In Section 28.4, we review evidence that generic language and generic modes of generalizing play an important role in stereotyping and prejudice.

28.1   What Makes a Generic True?

28.1.1   Content-Neutral Theories

According to *normalcy-based* theories, the truth of a generic depends fundamentally on whether it is *normal* for members of the relevant kind to have the property in question. What it means for it to be normal for a member of a kind K to have a property F varies across accounts, but these theories share the core idea that a generic is true just in case and because every K is such that, in every world in which things go as normally as possible for that K, that K has F. (See, e.g., Asher and Moreau 1995; Pelletier and Asher 1997; for a somewhat different implementation of a normalcy account, see Nickel 2008.) If these theories are right, then "tigers are striped" is true because of a certain modal fact: every tiger whose life goes normally as possible has stripes, and any tiger without stripes would have stripes if its life went as normally as possible. Since it is clearly not normal for every duck to lay eggs—consider male ducks, for example—in order to account for generics such as "ducks lay eggs" this type of account typically posits domain restriction, so that "ducks lay eggs" is semantically equivalent to "female ducks lay eggs."

According to *probability-based* theories, a generic is true in virtue of a fact about the actual world: the probability that an arbitrary member of the kind will have the property in question. Cohen (1996) offers such an account. He introduces the notion of a set of *alternatives* to a property – for example, the alternatives to *laying eggs* would include *giving live birth,* and any other forms of reproduction; the alternatives to *being paperback* would be *being hardcover,* and so on. Similarly, he introduces the notion of a set of alternatives to a kind – for example, the alternatives to *being a dog* would include *being a cat*, *being a rabbit,* and so forth. (These sets of alternatives for Cohen are defined so as to also include the original property/kind itself.) The account then posits that a generic is true just in case *either* the probability that an arbitrary member of the kind that has a property in the set of alternatives will have the property in question is greater than .5; *or else* the probability that an arbitrary member of the kind that has a property in the set of alternatives will have the property in question is greater than the probability that an arbitrary member of an alternative kind, which has a property in the set of alternative properties, will have the property in question. For instance, "ducks lay eggs" is true on Cohen's theory because the probability that a randomly chosen duck which engages in some or other form of female reproduction will lay eggs is greater than .5, whereas "mosquitoes carry the West Nile virus" is true because it is more likely that a randomly chosen mosquito will carry the West Nile virus than that a

405

randomly chosen insect of any kind will carry the West Nile virus. (For further details and refinement, see Cohen 1996, 2004.)

28.1.2   Content-Based Theories

Although content-neutral views differ dramatically in their details, they all share the assumption that the truth conditions of a generic do not vary with the content of the generic.

Content-based theories reject this assumption.

Leslie (2007, 2008) offers a content-based account according to which the truth conditions of a generic depend on its content in multiple ways. Consider, for example, the generic "mosquitoes carry West Nile virus." This generic is robustly accepted, despite the fact that very few mosquitoes have the property in question. Consider also: "sharks attack bathers," "pitbulls maul children," "ticks carry Lyme disease," and so on. These generics are also widely accepted, though again very few members of the kind have the property in question. A natural hypothesis is that these generics are accepted because the property in question is dangerous, threatening, appalling – the sort of property that one would wish to be forewarned about. Thus, Leslie suggests that when we are faced with such a property, we are inclined to accept generics concerning it, even if very few members of the kind possess it (for details and refinement, see Leslie 2007, 2008, in press). Leslie argues that content is relevant for generics in other ways too. For example, she argues that generics such as "ducks lay eggs" and "peacocks have beautiful tails" are found acceptable because they specify *characteristic properties* of the kinds in question – the method by which the kind reproduces; a distinctive and identifying physical feature of the kind (for details and refinement see Leslie 2007, 2008; Lerner and Leslie 2013). Only when the property in question is neither characteristic nor strikingly dangerous do we require that the majority of the kind have the property in question (hence the fact that, e.g., "cars have radios" appears to require majority satisfaction). Finally, Leslie accounts for generics such as "books are paperbacks," "people are right-handed," and "elementary school teachers are female" (which would seem false even though the majority of the kind have the property) by proposing that generics are rejected if the members of the kind that lack the attributed property have an equally salient, concrete, positive property in its stead. While the mosquitoes that do not carry West Nile virus simply fail to carry it, elementary school teachers who are not female are male, non-paperback books are hardcover, and non-right-handed people are left-handed.

28.2   The Cognitive Psychology of Generics

A great deal of ink has been spilled debating the relative merits of the foregoing accounts (see, e.g., Carlson 1977; Carlson and Pelletier 1995; Cohen 1996; Pelletier and Asher 1997; Schubert and Pelletier 1987; Nickel 2008; Leslie 2007, 2008). Despite the fact that generics are first and foremost a feature of natural language, these debates have been conducted almost exclusively from the armchair. However, there are strong reasons to think that empirical work on generics should play an important role in evaluating these accounts. For example, if adults are competent speakers of English, then their intuitions about whether an English generic is true should be regulated by whether they believe its truth conditions are satisfied. In this way, empirical work on what regulates people's acceptance of generics can provide evidence in favor of one account of truth-conditions over another. While empirical work of this kind is still in its early days, the studies conducted to date generally provide support for content-based theories over content-neutral theories. For example, recent findings from distinct experimental paradigms tell against the idea that generics such as "ducks lay eggs" are understood to involve domain restriction to

female ducks (Cimpian, Brandone, and Gelman 2010; Khemlani, Leslie, and Glucksberg 2009; see Leslie 2012b for discussion). If people do not understand "ducks lay eggs" as restricted to female ducks, then it is also difficult to maintain that this generic is accepted because it is normal for the members of the kind in question to lay eggs, as it is clearly not normal for male ducks to lay eggs. Nonetheless, Brandone *et al.* (2012) found that adults readily accept generics like "ducks lay eggs" despite knowing that only one sex has the property. Prasada *et al.* (2013) also provide further evidence against normalcy-based theories. In their study, participants regularly accepted generics of the form *Ks are F* while denying that *all normal Ks are F* or that *Ks are supposed to be F*. They also found that people accepted troublesome generics like "mosquitoes carry the West Nile virus" despite recognizing that only a minority of the kind in question have the attributed property.

Cimpian, Brandone, and Gelman (2010) found direct experimental support for the idea that people are more likely to accept generics at low prevalence levels if the property in question is strikingly dangerous. In their experiment, participants were told about a novel animal kind, for example, "lorches". They were then told that a certain percentage of lorches had a particular property – for example, purple feathers – and were given some further information about the feathers. Crucially, in some cases, this further information described a dangerous property – for example, that the feathers were poisonous to the touch – whereas in other cases it was neutral. The participants were then asked to evaluate a generic attributing the property to the kind, and the experimenters found that participants were significantly more likely to accept the generic when the property had been described as dangerous.

28.3   The Generics-as-Defaults Hypothesis

As the foregoing discussion indicates, empirical work on generics has the potential to illuminate philosophical and linguistic work on the topic. However, empirical work on generics may also have profound implications for questions in cognitive science. In particular, there is reason to believe that empirical work on generics can shed new light on foundational topics in cognitive science, including learning, reasoning, and stereotyping.

A great deal of learning depends crucially on the ability to generalize: to form general conclusions on the basis of particular cases. When a small child chases after a bee and is stung as a result, she learns to avoid bees in the future. This learning is able to take place because the child moves from a particular instance of a kind exhibiting a property – a bee causing pain – to a general conclusion linking that property to other instances of the kind. It is this general conclusion about the kind and its properties that allows the child to adjust her future behavior in light of her past experience.

Learning of this sort begins at an early age: even infants who are less than a year old can form general conclusions and adjust their behavior in light of them (see, e.g., Baldwin, Markman and Melartin 1993; Graham, Kilbreath and Welder 2001). This suggests that these generalizations are the product of an early-developing mode of generalizing. Let us then call these general conclusions *cognitively fundamental* generalizations. A set of critical questions for cognitive scientists concerns the nature of cognitively fundamental generalizations: what is their structure, how are they formed, and what role do they play in cognition more generally?

Empirical work on generics could provide insight into these questions if it turned out that we use generics to express these cognitively fundamental generalizations in natural language. The hypothesis that generics do in fact articulate these cognitively default generalizations has come to be known as the *generics-as-defaults hypothesis*.

407

28.3.1 Initial Motivations for the View

An initial reason to accept the generics-as-defaults hypothesis is that it solves two puzzles about generic language. The first is that few, if any, languages contain a dedicated, articulated generic operator (Dahl 1985; Krifka *et al.* 1995). In English, for instance, there is no word "gen" that marks generic generalizations in the way that explicit quantifiers like "all," "most," and "some" mark quantified statements. There are no English sentences of the form "gen tigers are striped," whereas sentences like "all tigers are striped" and "some tigers are striped" are commonplace. Indeed, part of the way to identify a generic sentence is by the fact that it contains no explicit operator: it is the absence rather than the presence of a word that signals the need for a generic interpretation. A ready explanation for this fact is that explicit operators are required only in order to signal the need for the cognitive system to deviate from the default interpretation of the sentence. In general, this is an efficient way for an interactive system to operate: if the system has a default mode of proceeding, then it only needs to be explicitly instructed if it is to *deviate* from that mode. In the absence of instruction, the system proceeds according to its default. (Compare the need to, e.g., *explicitly tell* an unhelpful teenager to perform a chore; if one does not wish him to perform the chore, one need say nothing at all.) If the default interpretation of a generalization is generic, then it would be inefficient for a natural language to include a word instructing the cognitive system to interpret a sentence generically: this is just what the cognitive system would have done in the absence of such a word (Leslie 2007, 2008).

A second, related puzzle readily explained by the generics-as-default hypothesis is that, despite the fact that semanticists have had far more trouble analyzing generic statements than quantified statements, children master generic language at an earlier age than they master quantified language. Toddlers as young as 30 months understand that generics express generalizations about kinds (Graham, Nayer, and Gelman 2011), and that these generalizations tolerate exceptions (Gelman and Raman 2003; Chambers, Graham, and Turner 2008; Gelman and Bloom 2007; Gelman, Star, and Flukes 2002). Preschoolers understand that these generalizations do not reduce to claims about how many members of the kind have the property in question (Brandone *et al.* 2012). For example, Brandone *et al.* (2012) found that preschoolers accept generics like "ducks lay eggs" but reject "ducks are girls" despite knowing that there are at least as many female ducks as there are ducks that lay eggs.

Children have much more difficulty mastering explicitly quantified statements. For example, when presented with a set of six crayons – some of which are in a box and some of which are not – preschool children have an exceptionally difficult time determining whether "most" of the crayons are in the box (Papafragou and Schwarz 2005/2006). Moreover, even children between the ages of six and eight failed to perform as well as adults on such a task. (See also Barner, Chow, and Yang 2009; Leslie and Gelman 2012, Experiment 2.)

Although preschool children have great difficulty processing the quantifier "most," they have more success with other quantifiers: although they cannot reliably determine whether "most" of the crayons in front of them are in a box, they can reliably determine whether "all" or "some" of the crayons in front of them are in the box (Barner, Chow, and Yang 2009). Nevertheless, preschoolers' competence with quantifiers like "all" and "some" appears to be limited to contexts in which the generalization concerns only a finite set of objects immediately in front of them; when asked to consider open-ended, kind-wide generalizations, preschoolers' competence with quantifiers wavers.

In an early study, Hollander, Gelman, and Star (2002) found that 3-year-old children processed universal generalizations as if they were generic generalizations. For instance, when asked "Do all books have color pictures?," three-year-olds responded "yes," just as often as they did when they were asked "Do books have colored pictures?" This effect seems to be due to three-year-olds

assimilating universal generalizations to generics rather than the other way around: three-year-olds' responses to the universally formulated questions differed from the more mature responses of the four-year-olds and adults, whereas the three-year-olds' responses to the generically formulated questions were statistically indistinguishable from the adults' responses.[1] These results provide especially strong support for the claim that children acquire competence with generic language earlier than they acquire competence with quantificational language.

28.3.2   Predictions

In addition to explaining these puzzles about generic language and its acquisition, the generics-as-defaults hypothesis generates a number of empirical predictions, many of which have been experimentally tested and confirmed. These predictions concern the formation of generic representations, the evaluation of general statements, and inferences about individual members of kinds.

28.3.2.1   The Formation of Generic Representations

If the generics-as-defaults hypothesis is correct, then two main predictions follow about the *formation* of generic representations. First, when people attempt to form quantified representations, they should occasionally "default to the generic" and form a generic representation instead. Second, when the cognitive system spontaneously draws a general conclusion from particular facts, that representation should be generic in form.

Consider the first prediction. If the generics-as-defaults hypothesis is correct, then it should be easier to process generic statements than quantified statements. If it is easier to process generic statements than quantified statements, then there should be cases where people mistakenly process quantified statements as generics. And if this is so, there should be cases where people mistakenly form generic representations on the basis of quantified statements.

One way to test this prediction is to ask people to memorize a series of quantified statements and to see whether they later recall some of those quantified statements as generic statements. Leslie and Gelman (2012) found that both adults and children who were asked to memorize a series of both generic statements and quantified statements more often misremembered quantified statements as generic statements than they misremembered generic statements as quantified statements. In a follow-up experiment, Leslie and Gelman found that even when the statements participants were asked to memorize were all quantified, the majority of participants mistakenly recalled at least one of those statements as a generic.

Sutherland, Cimpian, Leslie, and Gelman (2015) replicated Leslie and Gelman (2012)'s first finding among adults, and they further showed that the tendency to misremember quantified statements as generics increased when the quantified statements attributed *characteristic* properties (e.g., *taste with their feet*) to members of the kind instead of *accidental* properties (e.g., *have broken legs*). Sutherland *et al*. argue that these results show that people misremember quantified statements as generics not merely because it is easier to process a generic statement than a quantified statement, but because people have a default tendency to spontaneously form a generic belief that a kind has a particular characteristic property after learning that members of the kind exhibit that property. When people recall quantified statements as generics, they sometimes do this not only because they inadvertently processed the quantified statements as generics, but because they have a default tendency to spontaneously form generic representations on the basis of quantified information. These results provide strong support for the second prediction: that when the cognitive system spontaneously draws general conclusions from particular facts, those general conclusions are generic in form.

409

28.3.2.2   The Evaluation of General Statements

If the generics-as-defaults hypothesis is correct, then three main predictions follow about the *evaluation* of general statements. First, we would expect there to be cases where people "default to the generic" and mistakenly evaluate quantified statements as though they were generics. Secondly, if generics reflect cognitively default generalizations, then we would expect people to accurately evaluate generics more quickly than quantified statements. Lastly, we would expect the accurate evaluation of quantified statements to recruit more cognitive resources than the accurate evaluation of generics.

Consider the first prediction. If the default mode of processing a general statement is to process it as a generic, then accurately evaluating quantified statements will require inhibiting the cognitive system from relying on its default mode of processing. If this is so, then we would expect there to be circumstances where people are unable to inhibit this default mode and for that reason go on to evaluate quantified statements as if they were generics. These circumstances might include those in which people lack the general ability to reliably inhibit default processes (e.g., because they are children) as well as those in which people have the ability to inhibit default processes but cannot exercise it (e.g., because of time pressure or exhaustion).

The work of Hollander, Gelman, and Star (2002) provides a case where people – in this case, children – treat quantified statements as if they were generics because they lack the ability to override their default mode of processing.[2] Other studies provide cases of adults who have the ability to override their default mode of processing but fail to exercise that ability on particular occasions. In these studies, adults regularly endorse universal generalizations that are false when those generalizations would be true in generic form. Leslie, Khemlani, and Glucksberg (2011) call this phenomenon "the *generic overgeneralization* (GOG) effect, since it involves overgeneralizing from the truth of a generic to the truth of the corresponding universal statement." (17)

Leslie *et al.* found that adults exhibit the GOG effect when the generalizations in question are *characteristic* in the sense described in Section 28.1.2. When the characteristic property attributed to the kind was possessed by most of the kind (e.g., tigers having stripes), participants accepted the universal generalization 78% of the time. When the characteristic property was a property possessed by only a minority of the kind (e.g., lions having manes), participants accepted the false universal generalization 51% of the time. (Leslie, Khemlani, and Glucksberg 2011, Experiment 1)

Leslie *et al.* argue that these results are best explained by the generics-as-default hypothesis. However, several alternative explanations are available for why participants might accept false universal generalizations. One explanation is that participants simply do not realize that these characteristic universal generalizations are subject to counterexamples.[3] To test this hypothesis, Leslie *et al.* redesigned their experiment in a way that would allow them to see whether participants knew that the *minority* characteristic universals were subject to counterexamples. As in their earlier experiment, participants evaluated whether universal generalizations were true or false, but they also took a "knowledge test" that included items such as "male ducks lay eggs" and "female lions have manes." If participants rejected these items, it would show that they were aware that the minority characteristic universals were subject to counterexamples.

Despite performing well on the knowledge test, participants continued to regularly commit the GOG error. When participants evaluated the minority characteristic universals before taking the knowledge test, they accepted the universal and then went on to agree that it had counterexamples on 40% of trials. When participants took the knowledge test first – and were thus reminded that the universals had counterexamples before evaluating the universals – those who correctly agreed that one sex lacked the property went on to accept the universals on 19% of trials. These results show that the GOG effect is not due to participants' ignorance; even reminding participants of counterexamples before asking them to evaluate the universals did not completely eliminate the effect (Leslie, Khemlani, and Glucksberg 2011, Experiment 3).

410

Meyer, Gelman, and Stilwell (2011) provide additional support for the generics-as-defaults explanation of the GOG effect. Meyer *et al.* presented participants with a series of universal and generic generalizations on a computer screen. Participants were asked to press one of two keys to indicate whether they believed these generalizations were true or false. Half of the participants were asked to respond as quickly as possible, while the other half were told they could take as long as they would like. Those who were told they could take their time responding evaluated false majority characteristic universals such as "All dogs have four legs" with a reasonable level of accuracy. However, among those who were asked to answer as quickly as possible, the mean level of accuracy was much lower: participants were much more likely to commit the GOG error when they were rushed (Meyer, Gelman, and Stilwell, 916).

Importantly, no analogous effect was found among majority characteristic *generics*: participants had a high level of accuracy in evaluating these generic generalizations regardless of whether they were rushed (Meyer, Gelman, and Stilwell, 916). This pattern of results suggests that people are more likely to accept false majority characteristic universals when rushed, *not* because people are generally less accurate in evaluating generalizations under time pressure, but because people are less accurate in evaluating *universal* generalizations under time pressure. A compelling explanation for why people are less accurate in evaluating *universal* generalizations under time pressure is that correctly processing universal generalizations requires the cognitive system to inhibit its default generic interpretation of the generalization, and inhibiting a cognitive default takes time.

This explanation for why people are less accurate in evaluating universal characteristic generalizations than the corresponding generics when under time pressure assumes that it takes longer to correctly evaluate a universal generalization than to correctly evaluate a generic. This assumption is just the second prediction mentioned at the beginning of this section. This prediction is supported by additional findings by Meyer and colleagues. First, Meyer *et al.* found that participants under time pressure took longer to correctly evaluate false majority characteristic universals (e.g., "all dogs have four legs") than to correctly evaluate the true corresponding generics (e.g., "dogs have four legs"). Furthermore, they found that participants under time pressure took longer to correctly evaluate false majority characteristic universals than to correctly evaluate false "irrelevant-scope" universals that, unlike majority characteristic universals, have false generic counterparts (e.g., "All squirrels have beaks") (Meyer, Gelman, and Stilwell 2011, 917). This pattern of results suggests that the reason it takes more time to correctly evaluate a false majority characteristic universal than to correctly evaluate its true generic counterpart is not that correctly evaluating a false generalization generally takes more time than correctly evaluating a true generalization, but that correctly evaluating a false majority characteristic universal requires overcoming a default inclination to evaluate the universal as if it were a true generic.

The last prediction mentioned at the beginning of this section is that accurately evaluating quantified statements should require more cognitive resources than accurately evaluating generic statements. It has been argued that one way people evaluate quantified statements is by searching their semantic memory for counterexamples (Meyer, Gelman, and Stilwell, 2011; Lerner and Leslie, 2013). If accurately evaluating quantified statements requires more cognitive resources than accurately evaluating generic statements, and accurately evaluating quantified statements consists in retrieving counterexamples, then retrieving counterexamples should require more cognitive resources than accurately evaluating generics. There is some support for this claim. While the accurate evaluation of generics does not seem to depend in any special way on working memory, several studies show that counterexample retrieval is highly dependent on working memory. For instance, De Neys, Schaeken, and d'Ydewalle (2005) found evidence that people with higher working memory capacity were able to retrieve more counterexamples to

411

conditionals in a limited amount of time than those with lower working memory capacity. They also found that burdening people's working memory with a secondary task – continuously tapping out an unusual pattern with their fingers – significantly decreased the number of counterexamples they could retrieve in the allotted time. These findings lend support to the claim that counterexample retrieval is a cognitively demanding, non-default process that must compete with and overcome people's default inclination to evaluate false universal statements as true generics.

28.3.2.3   Inferences about Individual Members of Kinds

If the generics-as-defaults hypothesis is correct, then the acceptance of a generic "Xs are Y" should be closely tied to a person's default *inferences*: a person's tendency to infer that an arbitrary member of kind X has property Y. Empirical work by Cimpian, Brandone, and Gelman (2010) and Khemlani, Leslie, and Glucksberg (2012) confirms this prediction. Cimpian *et al.* found that people regularly inferred from statements of the form "Xs are Y" that almost *every* X is Y, despite believing that "Xs are Y" could be true even if only a minority of Xs were Y. Likewise, Khemlani *et al.* found that, even controlling for people's beliefs about the proportion of Xs that were Y, people's tendency to accept a generic of the form "Xs are Y" was a significant predictor of people's willingness to infer that an arbitrary X was Y.

28.4   Generics and Social Cognition

One prediction of the generics-as-defaults hypothesis is that hearing generics like "Ks are F" should play an important role in the formation of people's beliefs about various kinds, including *social* kinds. Unfortunately, recent work suggests that generic default reasoning underpins various forms of prejudice and stereotyping.

An "essentialist" belief about a kind is a belief that members of that kind share an internal nature or essence that causally grounds their common, enduring properties (e.g., Gelman 2003). People form essentialist beliefs about not only biological kinds, but also social kinds, and these beliefs about social kinds are tightly linked with various forms of prejudice. For example, people who hold essentialist beliefs about a social group are more likely to hold prejudiced attitudes toward members of that group (e.g., Haslam, Rothschild, and Ernst 2000, 2002), and experimental evidence suggests that essentialist beliefs may play a role in *causing* people to hold these prejudiced attitudes (Keller 2005). Furthermore, holding essentialist beliefs about a particular group is also linked to stereotyping that group more generally (e.g., Williams and Eberhardt 2008; Bastian and Haslam 2006; Prentice and Miller, 2006, 2007; Yzerbyt, Corneille, and Estrada 2001).

Recent work suggests that generic language encourages people to form essentialist beliefs. Gelman, Ware, and Kleinberg (2010) created three picture books about a novel animal kind, one of which was full of generic sentences (e.g., "Zarpies hate ice cream"), another which was full of specific sentences featuring a label for the kind (e.g., "this Zarpie hates ice cream"), and another which was full of specific sentences with no label for the kind (e.g., "this hates ice cream"). Gelman and colleagues found that both children and adults who read the picture book that included generic language were more likely to form essentialist beliefs about the novel animal kind than those who read the other picture books. However, other work suggests children form essentialist beliefs about animal kinds even in the absence of generic language (Gelman 2003), so generic language may do no more than speed up the formation of essentialist beliefs that children would have formed anyway.

412

Further work by Rhodes, Leslie, and Tworek (2012) suggests that generic language also causes people to form essentialist beliefs they would not have otherwise formed. Unlike Gelman *et al.*, Rhodes *et al.* made picture books featuring a novel *social* kind ("Zarpies") rather than a novel animal kind. In these picture books, Zarpies were depicted as a group of imaginary people who could not be assimilated to any existing social group. Thus, we would expect participants not to form essentialist beliefs about Zarpies absent key input.

As expected, those who read picture books without generic language showed little tendency to form essentialist beliefs about Zarpies. However, those who read picture books with generic language exhibited a significantly greater tendency to form essentialist beliefs about Zarpies. This tendency toward essentialization was found among four-year-olds both when they read the picture book twice immediately before being tested for essentialist beliefs, and when they read the book four times in the week leading up to testing. This suggests that generics encourage children to form essentialist beliefs quickly, and that these essentialist beliefs persist through time.

In a further experiment, Rhodes *et al.* also found that, by leading parents to form essentialist beliefs about Zarpies, they could more than double the number of generic utterances those parents used when describing Zarpies to their children. These results suggest that generics likely play a key role in the transmission of essentialist beliefs from parents to children.

The generics-as-defaults hypothesis can explain why hearing generics might lead people to form general beliefs about a kind. But why might hearing generics lead people to form *essentialist* beliefs? Cimpian and Markman (2009, 2011) presented children with novel facts about a familiar, essentialized kind (e.g., "butterflies have dust on their wings") or a specific member of that kind (e.g., "this butterfly has dust on her wings"), and then asked them to explain why these claims were true. While participants provided essentialist, kind-based explanations of generic facts (e.g., "they need the dust so they can fly"), they provided accidental, circumstance-based explanations of specific facts (e.g., "she flew through a dusty room"). This suggests that generics lead people to form essentialist beliefs via interpreting generics as *characteristic* generics – that is, via interpreting the generic as holding in virtue of something about the nature or essence of that kind.

28.5   Conclusion

Although generics are puzzling from the perspective of formal semantics, the empirical study of generics has already proven quite fruitful. Nevertheless, empirical work on generics is still in its early days. Much more work remains to be done before we have a complete psychological account of how people acquire and exercise their competence with generics. Future issues to address include, among other things, the relationship between the processing of generic statements and quantified statements, the revision of generic beliefs, and the role of generic beliefs in reasoning more generally.

Notes

1 For similar effects among Mandarin-speaking children and Quechua-speaking children, see Tardif *et al.* (2012) and Mannheim *et al.* (2011).
2 An alternative hypothesis is that children have the ability to override their default mode of processing, but they fail to exercise it because they do not understand the quantifiers "all" and "some." This hypothesis is undercut by the fact that, like the children in Barner, Chow, and Yang's (2009) study, Hollander *et al.*'s participants were able to identify both whether a determinate set of crayons were *all* in a box and whether *some* of those crayons were in the box. The fact that Hollander

et al.'s participants clearly understood the meaning of the quantifiers "all" and "some" when quantifying over a determinate set of concrete items suggests that their tendency to assimilate open-ended, kind-wide quantified statements to generics is not due to their general lack of competence with quantifier terms.

3 Leslie *et al.* consider and reject several other explanations (e.g., that participants implicitly restrict the domain of the universal to a subkind over which it holds) in the course of the paper.

References

Asher, Nicholas, and Michael Morreau. 1995. What Some Generic Sentences Mean. In *The Generic Book*, edited by Gregory N. Carlson and Francis Jeffry Pelletier, 53–74. Chicago: Chicago University Press.

Baldwin, Dare A., Ellen M. Markman, and Riikka L. Melartin. 1993. Infants' Ability to Draw Inferences about Nonobvious Object Properties: Evidence from Exploratory Play. *Child Development* 64: 711–728. DOI: 10.1111/j.1467-8624.1993.tb02938.x.

Barner, David, Katherine Chow, and Shu-Ju Yang. 2009. Finding One's Meaning: A Test of the Relation Between Quantifiers and Integers in Language Development. *Cognitive Psychology* 58: 195–219. DOI: 10.1016/j.cogpsych.2008.07.001.

Barwise, Jon, and Robin Cooper. 1981. Generalized Quantifiers and Natural Language. *Linguistics and Philosophy* 4: 159–219. 10.1007/BF00350139.

Bastian, Brock, and Nick Haslam. 2006. Psychological Essentialism and Stereotype Endorsement. *Journal of Experimental Social Psychology* 42: 228–235. doi:10.1016/j.jesp.2005.03.003.

Brandone, Amanda C., Andrei Cimpian, Sarah-Jane Leslie, and Susan A. Gelman. 2012. Do Lions Have Manes? For Children, Generics Are about Kinds Rather than Quantities. *Child Development* 83: 423–433. DOI: 10.1111/j.1467-8624.2011.01708.x.

Carlson, Gregory N. 1977. Reference to Kinds in English. PhD diss., University of Massachusetts.

Carlson, Gregory N., and Francis Jeffrey Pelletier, eds. 1995. *The Generic Book*. Chicago: Chicago University Press.

Chambers, Craig G., Susan A. Graham, and Juanita N. Turner. 2008. When Hearsay Trumps Evidence: How Generic Language Guides Preschoolers' Inferences about Unfamiliar Things. *Language and Cognitive Processes* 23: 749–766. DOI: 10.1080/01690960701786111.

Cimpian, Andrei, and Ellen M. Markman. 2009. Information Learned from Generic Language Becomes Central to Children's Biological Concepts: Evidence from Their Open-ended Explanations. *Cognition* 113: 14–25. DOI: 10.1016/j.cognition.2009.07.004.

Cimpian, Andrei, and Ellen M. Markman. 2011. The Generic/Nongeneric Distinction Influences How Children Interpret New Information about Social Others. *Child Development* 82: 471–492. DOI: 10.1111/j.1467-8624.2010.01525.x.

Cimpian, Andrei, Amanda C. Brandone, and Susan A. Gelman. 2010. Generic Statements Require Little Evidence for Acceptance But Have Powerful Implications. *Cognitive Science* 34: 1452–1482. DOI: 10.1111/j.1551-6709.2010.01126.x.

Cohen, Ariel. 1996. Think Generic: The Meaning and Use of Generic Sentences. PhD diss., Carnegie Mellon University.

Cohen, Ariel. 2004. Generics and Mental Representations. *Linguistics and Philosophy* 27: 529–556. DOI: 10.1023/B:LING.0000033851.25870.3e.

Dahl, Östen. 1985. *Tense and Aspect Systems*. Oxford: Blackwell.

De Neys, Wim, Walter Schaeken, and Géry d'Ydewalle. 2003. Inference Suppression and Semantic Memory Retrieval: Every Counterexample Counts. *Memory & Cognition* 31: 581–595.

De Neys, Wim, Walter Schaeken, and Géry d'Ydewalle. 2005. Working Memory and Counterexample for Causal Conditionals. *Thinking & Reasoning* 11: 123–150. DOI: 10.1080/13546780442000105.

Gelman, Susan A. 2003. *The Essential Child: Origins of Essentialism in Everyday Thought*. New York: Oxford University Press. DOI:10.1093/acprof:oso/9780195154061.001.0001.

Gelman, Susan A., and Paul Bloom. 2007. Developmental Changes in the Understanding of Generics. *Cognition* 105: 166–183. DOI: 10.1016/j.cognition.2006.09.009.

Gelman, Susan A., and Lakshmi Raman. 2003. Preschool Children Use Linguistic Form Class and Pragmatic Cues to Interpret Generics. *Child Development* 74: 308–325. DOI: 10.1111/1467-8624.00537.

Gelman, Susan A., Jon R. Star, and Jonathan E. Flukes. 2002. Children's Use of Generics in Inductive Inferences. *Journal of Cognition and Development* 3: 179–199. DOI: 10.1207/S15327647JCD0302_3.

Gelman, Susan A., Elizabeth A. Ware, and Felicia Kleinberg. 2010. Effects of Generic Language on Category Content and Structure. *Cognitive Psychology* 61: 273–301. DOI: 10.1016/j.cogpsych.2010.06.001.

Graham, Susan A., Cari S. Kilbreath, and Andrea N. Welder. 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,* edited by Johanna D. Moore and Keith Stenning, 352–357. Austin, TX: Cognitive Science Society.

Graham, Susan A., Samantha L. Nayer, and Susan A. Gelman. 2011. Two-year-olds Use the Generic/Nongeneric Distinction to Guide Their Inferences about Novel Kinds. *Child Development* 82: 493–507. DOI: 10.1111/j.1467-8624.2010.01572.x.

Haslam, Nick, Louis Rothschild, and Donald Ernst. 2000. Essentialist Beliefs about Social Categories. *British Journal of Social Psychology* 39: 113–127. DOI: 10.1348/014466600164363.

Haslam, Nick, Louis Rothschild, and Donald Ernst. 2002. Are Essentialist Beliefs Associated with Prejudice? *British Journal of Social Psychology* 41: 87–100. DOI: 10.1348/014466602165072.

Hollander, Michelle A., Gelman, Susan A., and Jon Star. 2002. Children's Interpretation of Generic Noun Phrases. *Developmental Psychology* 38: 883–894. DOI: 10.1037//0012-1649.38.6.883

Keller, Johannes. 2005. In Genes We Trust: The Biological Component of Psychological Essentialism and Its Relationship to Mechanisms of Motivated Social Cognition. *Journal of Personality and Social Psychology* 88: 686–702. DOI: 10.1037/0022-3514.88.4.686.

Khemlani, Sangeet, Sarah-Jane Leslie, and Sam Glucksberg. 2009. Generics, Prevalence, and Default Inferences. In *Proceedings of the 31st Annual Cognitive Science Society,* edited by Niels A. Taatgen and Hedderik van Rijn, 443–448. Austin, TX: Cognitive Science Society.

Khemlani, Sangeet, Sarah-Jane Leslie, and Sam Glucksberg. 2012. Inferences about Members of Kinds: The Generics Hypothesis. *Journal of Language and Cognitive Processes* 27: 887–900. DOI: 10.1080/01690965.2011.601900.

Krifka, Manfred, Pelletier, Francis Jeffry Gregory N. Carlson, Alice ter Meulen, Gennaro Chierchia, and Godehard Link. 1995. Genericity: An Introduction. In *The Generic Book*, edited by Gregory N. Carlson and Francis Jeffry Pelletier, 1–25. Chicago: Chicago University Press.

Lerner, Adam, and Sarah-Jane Leslie. 2013. Generics, Generalism, and Reflective Equilibrium. *Philosophical Perspectives* 27: 366–403. DOI: 10.1111/phpe.12034

Leslie, Sarah-Jane. 2007. Generics and the Structure of the Mind. *Philosophical Perspectives* 21: 375–403. DOI: 10.1111/j.1520-8583.2007.00138.x.

Leslie, Sarah-Jane. 2008. Generics: Cognition and Acquisition. *Philosophical Review* 117: 1–47. DOI: 10.1215/ 00318108-2007-023.

Leslie, Sarah-Jane. 2012a. Generics. In *The Routledge Companion to Philosophy of Language*, edited by Gillian Russell and Delia Graff Fara, 355–367. New York: Routledge.

Leslie, Sarah-Jane. 2012b. Generics Articulate Default Generalizations. *Recherches Linguistiques de Vincennes* 41: 25–45. DOI: 10.4000/rlv.2048.

Leslie, Sarah-Jane. In press. The Original Sin of Cognition: Fear, Prejudice, and Generalization. *Journal of Philosophy*.

Leslie, Sarah-Jane., and Susan A. Gelman. 2012. Quantified Statements Are Recalled as Generics: Evidence from Preschool Children and Adults. *Cognitive Psychology* 64: 186–214. DOI: 10.1016/j.cogpsych.2011.12.001.

Leslie, Sarah-Jane, Khemlani, Sangeet, and Sam Glucksberg. 2011. Do All Ducks Lay Eggs? The Generic Overgeneralization Effect. *Journal of Memory and Language* 65: 15–31. DOI:10.1016/j.jml.2010.12.005.

Mannheim, Bruce, Susan A. Gelman, Carmen Escalante, Margarita Huayhua, and Rosalina Puma. 2011. A Developmental Analysis of Generic Nouns in Southern Peruvian Quechua. *Language Learning and Development* 7: 1–23. DOI: 10.1080/15475441003635620.

Meyer, Meredith, Susan A. Gelman, and Sarah M. Stilwell. 2011. Generics Are a Cognitive Default: Evidence from Sentence Processing. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society,*

edited by Laura Carlson, Christoph Hölscher, and Thomas F. Shipley, 913–918. Austin, TX: Cognitive Science Society.

Nickel, Bernhard. 2008. Generics and the Way of Normality. *Linguistics and Philosophy* 31: 629–648.

Papafragou, Anna, and Naomi Schwarz. 2005/2006. *Most* wanted. *Language Acquisition* 13*:* 207–251. DOI:10.1207/s15327817la1303_3.

Pelletier, Francis Jeffry., and Nicholas Asher. 1997. Generics and Defaults. In *Handbook of Logic and Language*, edited by Johan van Benthem and Alice ter Meulen, 1125–1179. Cambridge, MA: MIT Press.

Prasada, Sandeep, Sangeet Khemlani, Sarah-Jane Leslie, and Sam Glucksberg. 2013. Conceptual Distinctions Amongst Generics. *Cognition* 126: 405–422. DOI: 10.1016/j.cognition.2012.11.010.

Prentice, Deborah A., and Dale T. Miller. 2006. Essentializing Differences between Women and Men. *Psychological Science* 17: 129–135. DOI: 10.1111/j.1467-9280.2006.01675.x.

Prentice, Deborah A., and Dale T. Miller. 2007. Psychological Essentialism of Human Categories. *Current Directions in Psychological Science* 16: 202–206. DOI: 10.1111/j.1467-8721.2007.00504.x.

Rhodes, Marjorie, Sarah-Jane Leslie, and Christina M. Tworek. 2012. Cultural Transmission of Social Essentialism. *Proceedings of the National Academy of Sciences (PNAS)* 109: 13526–13531. DOI: 10.1073/pnas.1208951109.

Schubert, Lenhart K., and Francis Jeffry Pelletier. 1987. Problems in Representing the Logical Form of Generics, Bare Plurals, and Mass Terms. In *New Directions in Semantics*, edited by Ernest Lepore, 387–453. London: Academic Press.

Sutherland, Shelbie L., Andrei Cimpian, Sarah-Jane Leslie, and Susan A. Gelman. 2015. Memory Errors Reveal a Bias to Spontaneously Generalize to Categories. *Cognitive Science* 39: 1021–1046. DOI: 10.1111/cogs.12189.

Tardif, Twila, Susan A. Gelman, Xiaolan Fu, and Liqi Zhu, L. 2012. Acquisition of Generic Noun Phrases in Chinese: Learning about Lions without an "-S." *Journal of Child Language* 39: 130–161. DOI: 10.1017/S0305000910000735.

Williams, Melissa J., and Jennifer L. Eberhardt. 2008. Biological Conceptions of Race and the Motivation to Cross Racial Boundaries. *Journal of Personality and Social Psychology* 94: 11033–1047.

Yzerbyt, Vincent, Olivier Corneille, and Claudia Estrada. 2001. The Interplay of Subjective Essentialism and Entitativity in the Formation of Stereotypes. *Personality and Social Psychology Review* 5: 141–155. DOI: 10.1207/S15327957PSPR0502_5.

416