1 **Running head: INFERENCES FROM GENERICS**

2
3
4 The unintended consequences of the things we say:
5 What generics communicate to children about unmentioned categories

6 Kelsey Moty and Marjorie Rhodes

7 New York University

8 Department of Psychology

9 6 Washington Place, New York, NY 10003, USA
10
11

12 **Corresponding author:** Kelsey Moty; *moty@nyu.edu*; (212) 998-7902; 6 Washington
13 Place, rm. 404; New York, NY 10003-6634. The authors declare no conflict of interest.

14 This research was supported by National Institute of Child Health and Human
15 Development R01HD087672 and National Science Foundation 1729540 (to M.R.) and
16 by a National Science Foundation Graduate Research Fellowship (to K.M.). The content
17 is solely the responsibility of the authors and does not necessarily represent the official
18 views of the National Institutes of Health or National Science Foundation.
19

Abstract

Adults frequently use generic language (e.g., "Boys play sports") to communicate information about social kinds to children. Whereas previous research speaks to how children often interpret information about the kinds described by generics, less is known about what generics may implicitly communicate about unmentioned kinds (e.g., the possibility that "Boys play sports" implies that girls do not). Studies 1 ($N = 287$ 4- to 6-year-olds, 56 adults) and 2 ($N = 84$ 4- to 6-year-olds) found that children as young as 4.5 years draw inferences about unmentioned categories from generic claims (but not matched specific statements)—and that the tendency to make these inferences strengthens with age. Study 3 ($N = 181$ 4- to 7.25-year-olds, 65 adults) provides evidence that pragmatic reasoning serves as a mechanism underlying these inferences. This paper concludes by discussing the role that generic language may play in inadvertently communicating social stereotypes to young children.

Keywords: generics, pragmatics, conceptual development, social kinds

The unintended consequences of the things we say:

What generics communicate to children about unmentioned categories

A child told that, “Girls wear pink,” “A boy doesn’t cry,” or “Scientists discover new things,” can learn (or mislearn) much about the world. These statements—in contrast to other forms of phrasing (e.g., “Some girls wear pink,” “Peter never cries,” “This scientist discovered something new”)—communicate *generic* information about what members of these groups are usually like, supposed to be like, or what makes them as a group unique in comparison to other groups (Carlson, 1977; Carlson & Pelletier, 1995). By age four, and in some cases earlier, children assume from statements like, “Boys love to wrestle,” that “*boys*” is an informative way of grouping people (Rhodes, Leslie, Bianchi, & Chalik, 2018), that *many* boys love to wrestle (e.g., Brandone, Gelman, & Hedglen, 2015; Gelman, Ware, & Kleinberg, 2010; Rhodes, Tworek, & Leslie, 2012), and that they do so *because* of something intrinsic to being a boy (Cimpian & Erickson 2012; Cimpian & Markman, 2009; 2011; Gelman et al., 2010; Rhodes et al., 2012). These assumptions are often wrong—adults use generics to describe rare properties, too (as in “Mosquitoes carry West Nile virus”, which is true of less than 1%; Cox, 2004; Cimpian, Brandone, & Gelman, 2010), as well as to refer to extrinsic causal relations (as in, “Girls wear pink,” which is the result of marketing, not anything intrinsic to *girls* or the color itself; Cimpian & Salomon, 2014).

Here we consider the possibility that children’s (sometimes inaccurate) interpretations of generics extend even further—to shape their beliefs about groups not mentioned at all. That is, upon hearing, “Girls are good at drawing,” children may assume that a new girl will be a good drawer. But will they also then assume that a new

boy will be bad? If so, these generic claims may inadvertently communicate social stereotypes to young children, a possibility important to consider given the frequency of generics in everyday parent-child conversations (Gelman, Goetz, Sarnecka, & Flukes, 2008; Graham, Nayer, & Gelman, 2011) and the ubiquity of generics across languages (Leslie, 2008).

Whether they do so might depend on children's beliefs about *why* a speaker chose these words (Grice, 1975; Clark, 1996; Horn, 1984). For example, if children (a) believe that the speaker knows a lot about both boys and girls, and therefore (b) *could* have knowledgeably made claims about both, then children could interpret the speaker's decision to say that "*Girls* are good at drawing" as intending to imply that boys are not (Jaswal & Neely, 2006; Koenig & Harris, 2005; Shafto, Goodman, & Frank 2012). This is particularly the case when the speaker could so easily have said something else like "Boys and girls are good at drawing" or "Kids are good at drawing." Similarly, if the speaker saw a particular girl doing a great drawing but chose to make a general claim about *girls* (e.g., saying "Girls are good..." rather than "This girl is good..."), then this decision could be interpreted as intent to imply that the particular girl's behavior is an example of more general trends.

This type of reasoning—referred to as *pragmatic reasoning* or *pragmatic inferencing*—can be difficult for young children. For instance, children before age seven struggle to realize that when someone says, "Jim ate some of the cookies," that this means that he likely did not eat *all* of them (Huang & Snedeker, 2009; examples with other quantifiers: Noveck, 2001; Papafragou & Musolino, 2003, etc.). Yet, this appears to be in part because children do not quite understand the meanings of "some" and "all",

and fail to consider that “all” was an alternative word that the speaker could have chosen (Barner, Brooks, & Bale, 2011). In simplified contexts that either (a) do not require knowing the meaning of these quantifiers or (b) make clearer the alternatives that speakers could have said, preschool-age children engage in sophisticated pragmatic reasoning (e.g., Stiller, Frank, & Goodman, 2015; Jara-Ettinger, Floyd, Huey, Tenenbaum, & Schulz, 2019). In contrast to quantifiers (like “all” or “some”), young children understand generics and can distinguish between generic claims and more specific ones by age two and a half (e.g., Graham, Gelman, & Clarke, 2016; Graham et al., 2011). Therefore, we hypothesized that children might draw sophisticated pragmatic inferences about the intended meanings of generics by quite early in childhood.

The aims of the present work were two-fold: to identify (a) what children infer, if anything, about unmentioned groups from generic claims (Study 1 and 2), and (b) whether children’s inferences about unmentioned groups rely on their assumptions about a speaker’s mental states (Study 3). To examine these questions, we used a novel social dichotomy between “zarpies” and “gorps” (marked by clothing color) so that participants would not be biased by pre-existing stereotypes. Doing so provides a strong test of what generics communicate about unmentioned kinds. We examined children ages four to six years because during this time, children’s pragmatic reasoning develops (e.g., Stiller et al., 2015; Huang & Snedeker, 2009), generics about social groups are commonly said to children (e.g., Gelman et al., 2004; Rhodes & Leslie, 2017), and children’s representations of social categories undergo many changes (Liberman, Woodward, & Kinzler, 2017).

Study 1

This study examined whether participants expect properties described as true of one of two groups (using generic or specific language) extend to other individuals from the mentioned or unmentioned group.

Pre-registered hypotheses. Consistent with previous research, children should more readily assume after hearing generics like, “Zarpies are good at baking pizzas,” that other zarpies are good at baking pizzas. Critically, if generics also communicate information about unmentioned groups (e.g., gorps), then children should respond that gorps are not. In contrast, children should be less sure about the status of gorps after hearing the specific claim, “This zarpie is good at baking pizzas.” We planned to examine age-related trends in children’s inferences about both groups to distinguish among three possibilities: (1) these inferences emerge early and robustly with no changes across age; (2) these inferences emerge at a young age but become more robust across early childhood; or (3) these inferences emerge later in development. An additional adult sample was tested for comparison.

Method

The study design and analytic code to test these hypotheses were pre-registered and can be found at <https://osf.io/c3fpv>.

Participants. Two-hundred eighty-seven children (57% girls), ages 4 to 6 years (112 4-year-olds, $M = 4.59$ years; 83 5-year-olds, $M = 5.40$ years; 92 6-year-olds, $M = 6.44$ years; note that we examined children in binned age groups for analyses that also included adults, but continuously in analyses that only examined children), and 56 adults (37% women, $M = 38.1$ years) participated. Of those parents who provided racial and ethnic demographic information for their children ($N = 248$), 56% of them identified

their child as White, 15% as Asian, 13% as Black/African-American, 2% as Middle Eastern or Northern African, and 15% as Mixed or Biracial; 22% of these participants were also Hispanic/Latino. We recruited children from local public preschools, elementary schools, and a children's museum in New York City. Adult participants were recruited via Amazon's Mechanical Turk. Adult participants predominantly identified as White (87%; other racial groups present in our sample each comprised less than 5% of the sample); 10% of adult participants identified as Hispanic/Latino. Approximately half of participants were assigned to one of two language conditions (*Generic* language versus *Specific* language).

We pre-registered a sample size of 60 participants per age bin (4s, 5s, 6s, and adults) with half in each language condition (planned total $N = 180$ children, 60 adults) based on the sample sizes used in comparable studies exploring pragmatic abilities in early childhood (e.g., Stiller et al., 2015; Horowitz & Frank, 2016). We planned to stop collecting data after gathering useable data from 180 children (with approximately 60 children per age group). We ended up with a larger sample of children, however, for two reasons. First, we ultimately decided to adopt more inclusive criteria than we specified in the pre-registration (because in retrospect we thought that the pre-specified criteria were too stringent for the youngest children; this allowed an additional 44 children to be retained for analyses). Second, we originally thought data from 57 participants were lost due to technical difficulties, but these data were ultimately recoverable. Analyses that precisely follow our pre-registered criteria and stopping rule (e.g., analyses of the first 180 children that were recruited who passed the pre-specified inclusion criteria) reveal similar patterns as reported here (see <https://osf.io/bk5x8/>).

Forty-three additional children and 4 adults participated in the study but were excluded for not completing the entire study ($N = 7$ children), incorrectly responding to attention checks ($N = 8$ children, 4 adults), interference from parents ($N = 2$ children), or technical difficulties with the testing software (e.g., software failed to record data in a manner that was unrecoverable; $N = 26$ children).

Procedure. Participants completed two tasks in a fixed order: (1) a novel-category inference task and (2) a context-dependent pragmatic ability task previously used with children even younger than those included here (Stillier, Goodman, & Frank, 2015). We presented these tasks on touchscreen tablets using testable.org. The context-dependent pragmatic ability task was included to make sure that this testing environment (using animations on a screen) could successfully elicit pragmatic reasoning in young children; performance on this task across all three studies was excellent (near ceiling) and will not be discussed further in the paper (see <https://osf.io/bk5x8/>). An example of the full procedure and materials can be found at: <https://nyu-cdsc.github.io/pidi-demo/>.

Novel-category inference task.

Learning phase. A recorded narrator introduced participants to a special town that had only two kinds of people: zarpies and gorps (see Figure 1 for summary of protocol). Group membership was perceptually marked by the clothing color (e.g., zarpies wore green and gorps wore yellow; group color was randomized). Besides wearing the same clothing color, members of both groups were otherwise diverse in respects to race/ethnicity, gender, and physical features. Next, to introduce the groups and make sure that children understood them as meaningful and distinct, the narrator

told participants four generic statements about each group (8 statements in total) in alternating order (e.g., “Zarpies like to climb tall fences,” “Gorps like to draw stars on their knees”). The learning phase included two knowledge checks—one where participants had to verbally provide the correct label for each group and one where they had to point to members of each group—to ensure participants learned the two groups. Participants were provided feedback on their responses regardless if they answered the question correctly and, overall, had a level of accuracy ($M = 91\%$; only 5 children answered both questions incorrectly).

Test phase. Following the learning phase, participants were presented with four test trials. For each trial, the narrator introduced participants to a new group of zarpies and gorps. The narrator then directed the participant’s attention to an individual (i.e., “Look at this zarpie!”) and described a property (e.g., “good at baking pizzas”) true of that individual. The language used to describe the property varied by condition: participants in the *Generic* condition heard about this property via a generic statement (e.g., “Zarpies are good at baking pizza”)—whereas participants in the *Specific* condition learned about this property via a specific statement (e.g., “This zarpie is good at baking pizza”). The language used to describe the target property across the 4 test trials is the only way in which the two conditions (*Generic* vs *Specific*) differed from one another.

After the target property was introduced, the individual moved to the top right corner of the screen to serve as a memory cue and the language used to introduce the target property was repeated once more (e.g., “I’ll move it here to remind you that zarpies are good at baking pizzas,” for the *Generic* condition). Finally, participants

responded either yes or no whether two additional individuals—a zarpie and a gorp drawn from a group of zarpies and gorps—also possessed the property (e.g., whether the new zarpie/gorp was also good at baking pizzas). Participants were asked about the individuals one at a time, and the order of the two questions (i.e., whether participants were asked about a zarpie or gorp first) was randomized across trials. Responses to these two questions across the four trials were used in our analyses.

Results

Our data was compiled, processed, and analyzed in R (Version 3.6.1; R Core Team, 2019) using the following packages: *tidyverse*, *geepack*, *lubridate*, *here*, *interactions*, *lme4*, *emmeans*, *cowplot*, and *ggpubr*. Our R code is available on GitHub (<https://github.com/nyu-cdsc/pidi>).

We analyzed participants' choices on the novel-category inference task using general estimating equations (GEE) because it can assess both within- and between-subject effects in binary data. GEE yields Wald χ^2 values as indicators of main effects and interactions. The models tested here used a binomial outcome distribution with a logit link function and an exchangeable correlation matrix (Ballinger, 2004). For analyses including adults, we treated age as a categorical variable (4s, 5s, 6s, and adults); for comparable analyses with children only, we use children's exact age (e.g., 4.53 years) treated as a continuous variable. We report the results of the Wald tests as well as estimated marginal means and slopes with 95% confidence intervals.

Property Extension. First, we examined the likelihood of extending the target property (i.e., responding “yes”) to other individuals as a function of the individual's group membership (from previously mentioned or unmentioned group), participant's

age, and language condition (*Generic* or *Specific*; preregistered on OSF <https://osf.io/c3fpv/>). Participants extended the target property more to members of the previously mentioned kind ($M = 69\%$, 95% CI [66%, 73%]) than the unmentioned kind ($M = 28\%$, 95% CI [25%, 31%], Group: $\chi^2(1) = 158$, $p < 0.001$), but the tendency to do so was stronger after hearing generic claims (Mentioned: $M_{generic} = 80\%$, 95% CIs [75%, 84%] vs $M_{specific} = 60\%$, [55%, 65%], Unmentioned: $M_{generic} = 18\%$, [14%, 22%] vs $M_{specific} = 37\%$, [32%, 42%]; Group x Condition: $\chi^2(1) = 44.1$, $p < 0.001$) and with age (Age x Group: $\chi^2(3) = 50.6$, $p < 0.001$; Age x Group x Condition: $\chi^2(3) = 11.9$, $p = 0.008$, see Figure 2A). All other effects had $ps > 0.10$.

To tease apart the interaction between age, condition, and group, we used separate GEEs to analyze the simple effects of condition and stimulus for each age group. At each age, participants responded that the individual from the unmentioned kind did not have the property more often if they heard the property described with generic than specific language (Condition x Group: 4s: $\chi^2(1) = 9.86$, $p = 0.002$, 5s: $\chi^2(1) = 9.05$, $p = 0.003$; 6s: $\chi^2(1) = 25.5$, $p < 0.001$; adults: $\chi^2(1) = 13.7$, $p < 0.001$).

These patterns of results held when examining only child participant data (age as continuous), including the three-way interaction between age, condition, and group membership ($\chi^2(1) = 9.88$, $p = 0.002$). With age, children were more likely to extend the target property to others from the mentioned group ($\beta = 0.62$, 95% CI [0.28, 0.95]) and less likely to members of the unmentioned group ($\beta = -0.68$, 95% CI [-1.10, -0.26]) after hearing generic language. Their rate of extending the target property after hearing specific language—to either the mentioned ($\beta = 0.12$, 95% CI [-0.16, 0.40]) or unmentioned group ($\beta = 0.0007$, 95% CI [-0.28, 0.28])—did not shift with age. Critically,

by 4.7 years of age (determined using the Johnson-Neyman technique, Johnson & Fay, 1950), children extended the target property to members of the unmentioned group less often after hearing generic language than specific language (see Figure 2B).

Property Inference. In the previous analyses, we examined the likelihood that participants extended the target property to mentioned and unmentioned kinds averaged across all trials. After submitting our pre-registration, we realized that an even stronger test of the hypothesis that participants infer that properties expressed via generics are not true of unmentioned kinds would be to examine participants' responses on a trial-by-trial basis. That is, for a given target property, do participants infer both (1) that the property is true (i.e., responded "yes") of members of the mentioned group but (2) not true (i.e., responded "no") of members of the unmentioned group? Participants who responded in this pattern were marked as making the inference and assigned a 1 for that trial; participants who responded in any other pattern did not make the expected inference and were assigned a 0.

When examining both the child and adult data, we found that participants were more likely to make the expected inference after hearing a generic statement ($M_{generic} = 69\%$, 95% CIs [66%, 73%] vs $M_{specific} = 41\%$, [38%, 45%]; Condition: $\chi^2(1) = 42.8$, $p < 0.001$), and the tendency to do so increased with age (Age: $\chi^2(3) = 43.3$, $p < 0.001$, Age x Condition: $\chi^2(3) = 8.89$, $p = 0.031$, see Figure 3A). Importantly, across all 4 age groups, participants who heard generic statements made the expected inference (i.e., responding "yes" to member of mentioned kind but "no" to member of unmentioned kind) more often than participants who heard specific statements (4s: $p = 0.002$; 5s: $p = 0.013$; 6s: $p < 0.001$; Adults: $p = 0.002$).

When we examined the child data alone (with age treated continuously), the two-way interaction between age and condition remained significant ($\chi^2(1) = 8.33, p = 0.004$). Using the Johnson-Neyman procedure (Johnson & Fay, 1950), we found that children older than 4.5 years of age made the expected inference significantly more frequently after hearing generic language than specific language (see Figure 3B), and the tendency to make inferences from generics increased with age (*Generic*: $\beta = 0.82$, 95% CI [0.45, 1.19]; *Specific*: $\beta = 0.14$, 95% CI [-0.14, 0.42]).

Together, these findings suggest that children as young as 4.5 years old infer meaning about unmentioned groups from generics.

Study 2

Study 2 was designed to replicate Study 1 with one key change. In Study 1's learning phase, children (in both language conditions) heard generics about both categories. While this was intended to help children learn the categories, it raised the possibility that children would make inferences about unmentioned categories only after first hearing a series of distinct generic claims about both groups. By removing this from the introduction, Study 2 tested whether children's inferences from generics rely on a cumulative effect after hearing many generics—or if these inferences persist even after hearing only a single generic with minimal introduction to each group.

Methods

We followed the methodological and analytic decisions pre-registered for Study 1. Although we found age related-changes in Study 1 (which we explore in more detail in Study 3), for this follow-up we did not intend to examine age-related changes and considered children as a single group.

Participants. Eighty-four children (62% girls), ages 4 to 6 years of age (31 4-year-olds, $M = 4.53$ years; 25 5-year-olds, $M = 5.57$ years; 28 6-year-olds, $M = 6.52$ years) participated in this study. Of those parents who provided racial and ethnic demographic information for their children ($N = 68$), 59% identified their child as White, 15% as Asian, 9% as Black/African-American, and 18% as Mixed or Biracial; 21% of these participants were also Hispanic/Latino. Children were recruited and assigned to a condition (*Generic* vs *Specific*) in the same manner as Study 1. Additionally, like Study 1, we determined our sample size based on comparable studies exploring the development of pragmatic reasoning (e.g., Stiller et al., 2015; Horowitz & Frank, 2016). Fifteen additional children participated but were excluded for technical difficulties with the testing software ($N = 10$), failure to complete the entire study ($N = 1$), and failure to correctly answer attention checks ($N = 4$).

Procedure. Procedures were identical to Study 1, except that in the novel-category inference task, we removed the portion of the learning phase in which children had heard generic language about both groups. That is, in Figure 1, we removed the portion depicted in panel 3.

Results

As in Study 1, we analyzed participants' choices on the novel-category inference task using GEE. We followed the same strategy of analyses as in Study 1, first examining how children extended the target property and then examining what inferences children made. Given the smaller sample size, we did not examine age-related changes in children's responses. We report the results of the Wald tests and estimated marginal means with 95% confidence intervals.

Property Extension. We found the same pattern of results as in Study 1:

Children were more likely to extend the target property to members of the previously mentioned group ($M = 0.68$, 95% CI [0.60, 0.75]) than the unmentioned group ($M = 0.33$, 95% CI [0.26, 0.40], Group: $\chi^2(1) = 30.8$, $p < 0.001$), but whether children did so depended on the language used to introduce the target property (Group x Condition: $\chi^2(1) = 7.41$, $p = 0.006$, see Figure 4A). Children who heard generics extended the target property to members of the mentioned group ($M = 75\%$, 95% CI [67%, 84%]) more frequently than to members of the unmentioned group ($M = 25\%$, 95% CI [16%, 34%], $p = 0.007$). However, children who heard specific language did not differentiate their property extensions to members of the mentioned ($M = 60\%$, 95% CI [50%, 69%]) and unmentioned groups ($M = 41\%$, 95% CI [31%, 50%], $p = 0.12$).

Property Inference. Next, we examined whether at the level of individual trials, children inferred that the target property was true of the mentioned group but not of the unmentioned group. We found that children more frequently made this inference after hearing the property conveyed via a generic ($M = 66\%$, 95% CI [55%, 75%]) versus a specific statement ($M = 40\%$, 95% CI [30%, 51%], $\chi^2(1) = 10.6$, $p = 0.001$, see Figure 4B), thus, replicating and extending the findings from Study 1 by demonstrating that children's inferences from generics do not depend on a cumulative effect from many generics but rather persist even with minimal introduction to both groups.

Study 3

Studies 1 and 2 demonstrated that children make inferences about unmentioned groups from generics. Study 3 provided a direct test of the mechanism underlying these inferences: do children generate these inferences by reasoning pragmatically about the

mental states of the speaker, or do they do so using some kind of heuristic that does not employ higher level reasoning about mental states (e.g., applying a rule that contrasting groups are opposite)? If pragmatic ability underlies these inferences, then children should only infer meaning about unmentioned groups when the generic is communicated by someone knowledgeable.

Pre-registered hypotheses. We predicted that participants who heard “Zarpies are good at baking pizzas,” for example, from a knowledgeable speaker (but not from an unknowledgeable one) will infer that gorps are not. Given Study 1’s findings, we hypothesized an age-related increase in children’s tendency to infer meaning about unmentioned groups from generics conveyed by knowledgeable speakers.

Method

The procedures and analytic code to test these hypotheses were pre-registered and can be found at <https://osf.io/ykwg4> and <https://osf.io/wfxrv>.

Participants. One-hundred-eighty-one children (50% girls), ages 4 to 7.25 years (59 4-year-olds, $M = 4.60$ years; 56 5-year-olds, $M = 5.47$ years; 66 6- to 7.25-year-olds, $M = 6.55$ years), and 65 adults ($M = 34.8$ years, 22% women) participated in this study. Of those parents who provided racial and ethnic demographic information for their children ($N = 161$), 56% of them identified their child as White, 20% as Asian, 7% as Black/African-American, 3% as Middle Eastern or Northern African, and 13% as Mixed or Biracial; 22% of these participants were also Hispanic/Latino. We did not gather demographic data for adult participants; however, participants were likely predominantly White, as was the case in Study 1 and in other studies (e.g., Huff & Tingley, 2015) that have examined the racial and ethnic makeup of participants recruited via Amazon

Mechanical Turk.

Children and adults were recruited in the same manner as in Study 1 and 2. As was the case in Study 1, we pre-registered a sample size of 60 participants per age bin (4s, 5s, 6s, and adults) with half in each condition (planned total $N = 180$ children, 60 adults) based on the sample sizes used in comparable studies exploring pragmatic ability in early childhood (e.g., Stiller et al., 2015; Horowitz & Frank, 2016). We planned to stop data collection after gathering useable data from 180 children (with approximately 60 children per age group). Approximately half of participants were assigned to one of two conditions (*knowledgeable* versus *unknowledgeable* speaker). One additional adult participated in the study but was excluded for incorrectly responding to attention checks. Eight children were excluded for failure to complete the entire task ($N = 2$), incorrectly responding to attention checks ($N = 5$), or technical difficulties ($N = 1$).

Procedure. Children participants completed three tasks in the following order: a warm-up task, a novel-category inference task, and a context-dependent pragmatic task. Adult participants only completed the warm-up task and novel-category inference task. We presented these tasks using an open-source application for running behavioral experiments (<https://github.com/nyu-cdsc/toku>) developed by the first author. An example of the procedure can be found at <https://nyu-cdsc.github.io/pidi-demo/>, and materials are available at <https://osf.io/m4jwf/>.

Novel-category inference task. This task was similar to the ones used in Study 1 and 2. We describe the differences in Study 3 from the previous versions below.

Learning phase. Similar to Study 1 and 2, a narrator introduced participants to a

special town that has of only two kinds of people (zarpies versus gorps). However, in Study 3, half the participants learned about the novel social dichotomy from a knowledgeable speaker (i.e., the speaker lived in the same neighborhood) and the other half learned about the dichotomy from an unknowledgeable speaker (i.e., the speaker was visiting the town for the first time; see <https://osf.io/8aqfc/> for scripts). Participants were then asked a set of knowledge checks to ensure that participants were (1) familiar with the distinction between zarpie and gorps and (2) aware of the speaker's knowledge or lack thereof of the two groups. Participants received feedback on their responses to these questions to ensure they learned the category boundaries and critical manipulation. Children were not excluded for incorrectly answering these questions and performed well overall ($M = 80\%$, 83% of children answered at least 3 of the 4 questions correctly; no child answered all 4 wrong).

Test phase. Participants saw four test trials similar to those used in Study 1 and 2, with two notable differences. One, regardless of condition (knowledgeable vs unknowledgeable speaker), the speaker saw a category member doing the target property and then introduced the property using a generic statement (e.g., speakers in both condition saw an individual zarpie baking pizzas and then stated "Zarpies are good at baking pizzas"). Thus, speakers in both conditions had access to the same evidence (of an individual zarpie baking a pizza) at the time they made their claim; what differed across condition was the *background knowledge* about zarpies and gorps that the speaker brought to the task. Halfway through, participants received a reminder about the speaker's background knowledge (i.e., whether this was their first time visiting the town). Two, in previous versions of the task, during the test phase, participants were

provided a visual cue about the target property and the speaker's statement was repeated. In Study 3, participants did not receive this visual aid nor was the generic statement repeated. Removing these cues helps rule out possible alternative explanations that the findings from the previous studies were the result of low-level perceptual mechanisms.

Results

Property Extension. Following the same analysis strategy described in Study 1, we first examined the likelihood of extending the target property (i.e., responding "yes") to other individuals as a function of the individual's group membership (from previously mentioned or unmentioned group), participant's age, and condition (whether the speaker was knowledgeable or unknowledgeable). Children and adults were more likely to extend the target property to members of the previously mentioned than unmentioned group (Group: $\chi^2(1) = 123, p < 0.001$). Critically, their tendency to do so depended on the speaker's knowledge of two groups (Group x Condition: $\chi^2(1) = 39.5, p < 0.001$) and the participant's age (Age x Group: $\chi^2(3) = 50.1, p < 0.001$; Age x Condition: $\chi^2(3) = 18.9, p < .001$; Age x Group x Condition: $\chi^2(3) = 7.56, p = 0.056$, see Figure 5A). Across all four age groups, children and adults were more likely to infer that the target property was *not* true of individuals from the unmentioned group after hearing generics from a knowledgeable speaker compared to when they heard the exact same generic statement made by an unknowledgeable speaker (Group x Condition: 4s: $\chi^2(1) = 8.23, p = 0.004$, 5s: $\chi^2(1) = 8.86, p = 0.003$; 6s: $\chi^2(1) = 12.2, p < 0.001$; adults: $\chi^2(1) = 20.6, p < 0.001$).

We found the same pattern of results when looking at children's responses only

(these analyses used children's exact age as continuous variable). Children inferred that the target property was true of individuals from the previously mentioned group but not true of individuals from the unmentioned group—but they did so only when the speaker was knowledgeable and the tendency to do increased with age (Age x Condition x Group: $\chi^2(1) = 5.14$, $p = 0.023$). With age, children were less likely to extend the target property to members of the unmentioned group after hearing generic language from a knowledgeable speaker ($\beta = -0.88$, 95% CI [-1.30, -0.45]), but not after hearing a generic from an unknowledgeable speaker ($\beta = 0.12$, 95% CI [-0.16, -0.39]). Critically, by 5.1 years of age (determined using the Johnson-Neyman technique, Johnson & Fay, 1950; see Figure 5B), children extended the target property to members of the unmentioned group less often after hearing a generic from a knowledgeable speaker than an unknowledgeable speaker.

Property Inference. Next, looking at the inferences that participants made for a given target property (i.e., infer that the mentioned group has the property but that the unmentioned does not), children and adults were more likely to make the expected inference when the generic was communicated by a knowledgeable speaker ($M = 0.70$, 95% CI [0.63, 0.76]) than an unknowledgeable speaker ($M = 0.39$, 95% CI [0.33, 0.45], Condition: $\chi^2(1) = 41.1$, $p < 0.001$, see Figure 6A). The tendency to make this kind of inference increased with age (Age: $\chi^2(3) = 24.1$, $p < 0.001$). While we did not find an interaction between age (when treated categorically) and condition with the adult data included in the model (Age x Condition: $\chi^2(3) = 4.81$, $p = 0.19$), we do see this interaction between age and condition when examining the child data alone (here, age treated continuously; Age x Condition: $\chi^2(1) = 8.16$, $p = 0.004$). By 4.7 years of age

(determined using the Johnson-Neyman procedure, see Figure 6B; Johnson & Fay, 1950), children were more likely to infer that properties conveyed via generics are not true of unmentioned groups when the generic was made by a knowledgeable speaker than when the same generic statement was said by an unknowledgeable speaker. The tendency to make inferences about unmentioned groups from generics only when conveyed by knowledgeable speakers increased with age ($\beta = 0.74$, 95% CI [0.32, 1.17]).

Together, these findings suggest children engage in pragmatic reasoning (as opposed to employing a general heuristic or rule) when making inferences about unmentioned groups, given that they only do so when the generic is communicated by a knowledgeable speaker.

General Discussion

In these three studies, children made inferences about unmentioned categories from generic statements. That is, upon hearing “Zarpies are good at baking pizzas”, they assumed not only that a new zarpie would be good at this, but also that a new gorp would not. Children only made these inferences when generics were said by knowledgeable speakers, suggesting that children reason about a speaker’s intended meaning to draw these inferences. Children’s tendency to make these inferences emerged early (around 4.5 years) and became more robust with age.

Beyond providing insight on children’s inferences from generics, these studies also provide evidence for a potential mechanism by which social stereotypes may be inadvertently communicated to young children. Parents may believe that generic statements—such as “Boys are good at sports” or “Girls play with dolls”—provide an

innocuous way to describe social groups. In fact, parents use generic language frequently when talking with their children (e.g., Gelman et al., 2008). But these studies suggest that young children take away not only the positive message from these statements (e.g., that boys are good at sports), but an additional negative message about other groups not explicitly mentioned (e.g., that girls are not). This can reinforce existing stereotypes or perhaps even create new ones. This is likely especially true when social categories have clear and salient contrasts, as was the case in the present experiments, as well as with gender (when presented as binary: boys versus girls).

This idea that generics may communicate stereotypes about unmentioned groups builds upon a growing body of research demonstrating the power of generics in transmitting social stereotypes about explicitly mentioned kinds to young children—by signaling how to carve up the social world into meaningful categories (Rhodes et al., 2018), promoting essentialist beliefs about social kinds (Rhodes et al., 2012), and consequently increasing social stereotyping and negative inter-group relations (e.g., Hammond & Cimpian, 2017; Leslie, 2017; Rhodes, Leslie, Saunders, Dunham, & Cimpian, 2017). Although we have primarily focused on children, generics also bias the beliefs of adults (as confirmed here, among other studies, e.g., Rhodes et al., 2012; Cimpian et al. 2010), with broad consequences for reasoning across domains, including how they interpret and value scientific research (DeJesus, Callanan, Solis, & Gelman, 2019).

The present studies suggest a number of interesting open questions that future research should explore. First, it would be informative to examine the developmental mechanisms underlying the age-related increases in the robustness of children's

inferences. For instance, these inferences could become more robust because of increased efficiency with pragmatic reasoning generally, or because of more general improvements in children's ability to make sense of our experimental task. Consistent with this latter possibility, our supplementary analyses, which included only children who met the more stringent inclusion criteria of passing all attention checks, found that children reliably made these inferences about six months earlier than found in analyses with the broader sample. Second, it is unknown whether the unmentioned group needs to be physically present to warrant the inference (as was the case in the present studies), or if children would also extend inferences about unmentioned kinds even when they are not perceptually available. The latter may be particularly difficult for young children, who already struggle to determine the relevant scope for pragmatic inferences (e.g., Barner et al., 2011). Third, these studies suggest that children make inferences about unmentioned kinds when the mentioned and unmentioned kind contrast along a social dichotomy (akin to a binary and discrete perspective of gender). Many social categories have multiple relevant contrasts, however, as is the case for race, or are presented as having less distinct or discrete boundaries. Future research should examine the types of inferences that children draw about unmentioned kinds in these cases, as well as how these inferences play out for real-world categories (e.g., gender-based inferences). In sum, the present studies highlight how implicit features of communication can shape the development of children's social understanding, paving the way for future work to examine the full extent of these effects.

Author Contributions.

K. Moty developed the study concept. All authors contributed to the study design. Testing and data collection were performed by K. Moty. K. Moty performed the data analysis and interpretation under the supervision of M. Rhodes. K. Moty drafted the manuscript, and M. Rhodes provided critical revisions. All authors approved the final version of the manuscript for submission.

Acknowledgements.

We thank Angela Sorenson, Alison De Leon, Jerolynn Liu, Kevin Bui, Samantha Aboody, Samuel Krystal, and Shelly Reizen for assistance with data collection. We also thank the staff and families at the Children's Museum of Manhattan, New York City District 2 preschools, and PS 132 for their support and participation.

References

- Ballinger, G. A. (2004). Using Generalized Estimating Equations for Longitudinal Data Analysis. *Organizational Research Methods*, 7(2), 127-150.
<https://doi.org/10.1177/1094428104263672>
- Barner, D., Brooks, N., & Bale, A. (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition*, 118(1), 84-93.
<https://doi.org/10.1016/j.cognition.2010.10.010>
- Brandone, A., Gelman, S. A., & Hedglen, J. (2015). Children's developing intuitions about the truth conditions and implications of novel generics vs. quantified statements. *Cognitive Science*, 39(4), 711-38. <https://doi.org/10.1111/cogs.12176>
- Carlson, G. N. (1977). A unified analysis of the English bare plural. *Linguistics and Philosophy*, 1(3), 413-457. <https://doi.org/10.1007/BF00353456>
- Carlson, G. N., & Pelletier, F. J. (1995). *The generic book*. Chicago, IL: University of Chicago Press.
- Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic Statements Require Little Evidence for Acceptance but Have Powerful Implications. *Cognitive Science*, 34(8), 1452–1482. doi:10.1111/j.1551-6709.2010.01126.x
- Cimpian, A., & Erickson, L. C. (2012). Remembering kinds: New evidence that categories are privileged in children's thinking. *Cognitive Psychology*, 64(3), 161–185. doi:10.1016/j.cogpsych.2011.11.002
- Cimpian, A., & Markman, E. M. (2009). Information learned from generic language becomes central to children's biological concepts: Evidence from their open-ended explanations. *Cognition*, 113(1), 14-25.

- 555 <http://dx.doi.org/10.1016/j.cognition.2009.07.004>
- 556 Cimpian, A., & Markman, E. M. (2011). The generic/nongeneric distinction influences
- 557 how children interpret new information about social others. *Child Development*,
- 558 82(2), 471-492. <https://doi.org/10.1111/j.1467-8624.2010.01525.x>
- 559 Cimpian, A., & Salomon, E. (2014). The inherence heuristic: An intuitive means of
- 560 making sense of the world, and a potential precursor to psychological
- 561 essentialism. *Behavioral and Brain Sciences*, 37(5), 461–480.
- 562 doi:10.1017/s0140525x13002197
- 563 Clark, H. (1996). *Using language*. Cambridge: Cambridge University Press.
- 564 <https://doi.org/10.1017/CBO9780511620539>
- 565 Cox, A. (2004, May 3). U.S. poised for epidemic West Nile year. Retrieved from
- 566 <http://edition.cnn.com/2004/HEALTH/05/03/wnv.outlook/>
- 567 DeJesus, J. M., Callanan, M. A., Solis, G., & Gelman, S. A. (2019). Generic language in
- 568 scientific communication. *Proceedings of the National Academy of Sciences*,
- 569 201817706. doi:10.1073/pnas.1817706116
- 570 Everett, D. L. (2009). Pirahã culture and grammar: A response to some criticisms.
- 571 *Language*, 85(2), 405–442. doi:10.1353/lan.0.0104
- 572 Gelman, S. A., Goetz, P. J., Sarnecka, B. S., & Flukes, J. (2008). Generic language in
- 573 parent-child conversations. *Language Learning and Development*, 4(1), 1-31.
- 574 doi:10.1080/15475440701542625
- 575 Gelman, S. A., Sánchez Tapia, I., & Leslie, S.-J. (2016). Memory for generic and
- 576 quantified sentences in Spanish-speaking children and adults. *Journal of Child*
- 577 *Language*, 43(6), 1231–1244. doi:10.1017/s0305000915000483

- 578 Gelman, S. A., Ware, E. A., & Kleinberg, F. (2010). Effects of generic language on
579 category content and structure. *Cognitive Psychology*, 61(3), 273-301.
580 <https://doi.org/10.1016/j.cogpsych.2010.06.001>
- 581 Graham, S. A., Gelman, S. A., & Clarke, J. (2016). Generics license 30-month-olds'
582 inferences about the atypical properties of novel kinds. *Developmental*
583 *Psychology*, 52(9), 1353–1362. doi:10.1037/dev0000183
- 584 Graham, S. A., Nayer, S. L., & Gelman, S. A. (2011). Two-year-olds use the
585 generic/nongeneric distinction to guide their inferences about novel kinds. *Child*
586 *Development*, 82(2), 493-507. <https://doi.org/10.1111/j.1467-8624.2010.01572.x>
- 587 Grice, H. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and*
588 *semantics* (Vol. 3, pp. 41–58). New York, NY: Academic Press.
- 589 Hammond, M. D., & Cimpian, A. (2017). Investigating the cognitive structure of
590 stereotypes: Generic beliefs about groups predict social judgments better than
591 statistical beliefs. *Journal of Experimental Psychology: General*, 146(5), 607–
592 614. doi:10.1037/xge0000297
- 593 Horn, L. (1984). Toward a new taxonomy for pragmatic inference: Q-based and r-based
594 implicature. In Schiffrin, D. (Ed.), *Meaning, form, and use in context: Linguistic*
595 *applications* (pp 11-42). Washington, D.C.: Georgetown University Press.
- 596 Horowitz, A. C., & Frank, M. C. (2016). Children's pragmatic inferences as a route for
597 learning about the world. *Child Development*, 87(3), 807-819.
598 <https://doi.org/10.1111/cdev.12527>
- 599 Huang, Y. & Snedeker, J. (2009). Semantic meaning and pragmatic interpretation in
600 five-year olds: Evidence from real time spoken language comprehension.

- 601 *Developmental Psychology*, 45(6), 1723-1739.
- 602 <http://dx.doi.org/10.1037/a0016704>
- 603 Huff, C., & Tingley, D. (2015). "Who are these people?" Evaluating the demographic
- 604 characteristics and political preferences of MTurk survey respondents. *Research*
- 605 & *Politics*, 2(3), 205316801560464. doi:10.1177/2053168015604648
- 606 Jara-Ettinger, J., Floyd, S., Huey, H., Tenenbaum, J. B., & Schulz, L. E. (2019). Social
- 607 pragmatics: Preschoolers rely on commonsense psychology to resolve referential
- 608 underspecification. *Child Development*. doi:10.1111/cdev.13290
- 609 Jaswal, V., & Neely, L. A. (2006). Adults don't always know best: Preschoolers use past
- 610 reliability over age when learning new words. *Psychological Science*, 17(9), 757-
- 611 758. <https://doi.org/10.1111/j.1467-9280.2006.01778.x>
- 612 Johnson, P. O., & Fay, L. C. (1950). The Johnson-Neyman technique, its theory and
- 613 application. *Psychometrika*, 15(4), 349–367. <https://doi.org/10.1007/bf02288864>
- 614 Leslie, S.-J. (2008). Generics: Cognition and acquisition. *Philosophical Review*, 117(1),
- 615 1-47. <https://doi.org/10.1215/00318108-2007-023>
- 616 Leslie, S.-J. (2017). The original sin of cognition: Fear, prejudice, and generalization.
- 617 *Journal of Philosophy*, 114(8), 393–421. doi:10.5840/jphil2017114828
- 618 Liberman, Z., Woodward, A. L., & Kinzler, K. D. (2017). The origins of social
- 619 categorization. *Trends in Cognitive Sciences*, 21(7), 556–568.
- 620 doi:10.1016/j.tics.2017.04.004
- 621 Mannheim, B., Gelman, S. A., Escalante, C., Huayhua, M., & Puma, R. (2011). A
- 622 developmental analysis of generic nouns in Southern Peruvian
- 623 Quechua. *Language Learning and Development*, 7(1), 1–23.

<https://doi.org/10.1080/15475441003635620>

Noveck, I. A. (2001). When children are more logical than adults: Experimental investigation of scalar implicatures. *Cognition*, 78(2), 165–188.

[https://doi.org/10.1016/s0010-0277\(00\)00114-1](https://doi.org/10.1016/s0010-0277(00)00114-1)

Papafragou, A., & Musolino, J. (2003). Scalar implicatures: Experiments at the semantics–pragmatics interface. *Cognition*, 86(3), 253–282.

[https://doi.org/10.1016/s0010-0277\(02\)00179-8](https://doi.org/10.1016/s0010-0277(02)00179-8)

Rhodes, M., Leslie, S.-J., Bianchi, L., & Chalikh, L. (2018). The role of generic language in the early development of social categorization. *Child Development*, 89(1), 148–155. <https://doi.org/10.1111/cdev.12714>

Rhodes, M., Leslie, S.-J., Saunders, K., Dunham, Y., & Cimpian, A. (2017). How does social essentialism affect the development of inter-group relations? *Developmental Science*, 21(1). doi:10.1111/desc.12509

Rhodes, M., Leslie, S.-J., & Tworek, C. M. (2012). Cultural transmission of social essentialism. *Proceedings of the National Academy of Sciences*, 109(34), 13526–13531. doi:10.1073/pnas.1208951109

Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others: The consequences of psychological reasoning for human learning. *Perspectives on Psychological Science*, 7(4), 341–351.

<https://doi.org/10.1177/1745691612448481>

Stiller, A. J., Goodman, N. D., & Frank, M. C. (2015) Ad-hoc implicature in preschool children. *Language Learning and Development*, 11(2), 176–190.

<https://doi.org/10.1080/15475441.2014.927328>

647 Tardif, T., Gelman, S. A., Fu, X., & Zhu, L. (2012). Acquisition of generic noun phrases
648 in Chinese: Learning about lions without an '-s'. *Journal of Child Language*,
649 39(1), 130-161. <https://doi.org/10.1017/s0305000910000735>

650

Figure 1*Protocol for Novel-Category Inference Task***Learning phase:**

1
Introduced to narrator who wants to tell participants about a special town.

2
Narrator described novel social dichotomy between zarpies and gorps. Participants asked to recall group names.

3
Participants heard a set of generic statements about both groups.

4
Participants asked to point to the zarpie and gorp to check knowledge of group distinction.

Test phase (4 trials):

5
Narrator drew a random individual from a sample of zarpies and gorps.

6
Narrator described property in generic ("Zarpies are good at baking pizzas") or specific form ("This zarpie is good at. . .").

7a
Narrator randomly selected a zarpie from sample. Participants responded yes/no if that zarpie is also good at the property.

7b
Narrator did the same for a random gorp. Order in which participants were asked about zarpie vs. gorp was randomized.

Note. Summary of the protocol for the novel-category inference task used in Study 1 (with sample stimuli). Zarpies are shown here in yellow and gorps in green. Study 2 used the exact same protocol, except participants did not hear the set of generic statements in Panel 3.

Figure 2.**Proportion of Trials Participants Extended Target Property in Study 1**

Notes. Proportion of trials that participants extended the target property to individuals of the mentioned and unmentioned groups by age and condition. (A) Error bars and bands reflect 95% CIs. Small shapes reflect individual averages, and large shapes and lines reflect group averages. (B) Table reflects average proportions and 95% CIs across age, condition, and group membership. (C) Johnson-Neyman plots. Lines reflect differences in simple slopes of condition predicting participants' extension of the target property as a function of age and group membership. Rate of property extension to the mentioned group became significantly different at 4.1 years, and to the unmentioned group at 4.7 years.

Figure 3

Proportion of Trials Participants Made Pragmatic Inference in Study 1

Notes. Proportion of trials that participants made the expected inference (i.e., responded “yes” individual of mentioned group has target property but “no”, individual of unmentioned group does not) by age and condition. (A) Error bars and bands reflect 95% CIs. Small shapes reflect individual averages, and large shapes and lines reflect group averages. (B) Table reflects average proportions and 95% CIs across age and condition. (C) Johnson-Neyman plots. Line reflect difference in simple slopes of condition predicting participants’ inferences as a function of age. Rate of making inferences across conditions becomes significantly different at 4.5 years.

Figure 4

Proportion Trails Participants Extended Target Property and Made Pragmatic Inference in Study 2

Notes. (A) Proportion of trials that children extended the target property to individuals of the mentioned and unmentioned groups by condition. Error bars reflect 95% CIs. Small shapes reflect individual averages, and large shapes and lines reflect group averages. (B) Proportion of trials that children extended the target property to individuals of the mentioned and unmentioned groups by condition.

Figure 5.
Proportion of Trials Participants Extended Target Property in Study 3

Notes. Proportion of trials that participants extended the target property to individuals of the mentioned and unmentioned groups by age and condition. (A) Error bars and bands reflect 95% CIs. Small shapes reflect individual averages, and large shapes and lines reflect group averages. (B) Table reflects average proportions and 95% CIs across age, condition, and group membership. (C) Johnson-Neyman plots. Lines reflect differences in simple slopes of condition predicting participants' extension of the target property as a function of age and group membership. Rate of property extension to the mentioned group was significantly different across all observed ages, and became significantly different at 5.1 years for the unmentioned group.

Figure 6

Proportion of Trials Participants Made Pragmatic Inference in Study 3

Notes. (A) Proportion of trials that participants made the expected inference (i.e., responded “yes” individual of mentioned group has target property but “no”, individual of unmentioned group does not) by age and condition. Error bars and bands reflect 95% CIs. Small shapes reflect individual averages, and large shapes and lines reflect group averages. (B) Table reflects average proportions and 95% CIs across age and condition. (C) Johnson-Neyman plot. Line reflects difference in slopes between knowledgeable and unknowledgeable speaker conditions across age. Slopes between the conditions become significantly different at 4.7 years-old.