Preschoolers use minimal statistical information about social groups to infer the preferences and group membership of individuals

Natalia Vélez (nvelez@stanford.edu), Hyowon Gweon (gweon@stanford.edu)
Department of Psychology, 450 Jane Stanford Way
Stanford, CA 94305 USA

Abstract

We don't learn about each person we meet from scratch: Our knowledge of social groups (e.g., cognitive scientists) shapes our expectations about new individuals (e.g., the reader). Here we explore how 4- and 5-year-old children and adults use minimal statistical evidence about social groups to support inductive inferences about individuals. Overall, we find that both children and adults readily infer the preferences and group membership of new individuals when they have appropriate evidence to support these inferences. However, our results also suggest that children and adults interpret this information in different ways. Adults' responses align closely with a Bayesian model that assumes that each group's preferences are independent of one another. By contrast, we find preliminary evidence that children's inferences about the preferences of new group members are sensitive to the composition (Experiment 1) and size (Experiment 2) of the opposing group. Our work provides insights into how people form structured, generalizable representations of social groups from sparse data.

Keywords: cognitive development; statistical reasoning; social groups

Introduction

Social groups give structure to an otherwise crowded social world by highlighting similarities and affiliations between people (Macrae & Bodenhausen, 2000; Rhodes & Baron, 2019). However, thinking of others in terms of their group membership can be a double-edged sword. On one hand, categorizing others into social groups can elicit a wide range of intergroup biases, even when the groups are entirely novel and emphatically arbitrary (Dunham, 2018). On the other hand, social groups can signal latent properties that support predictions about individuals. For example, one can use one's past experiences interacting with "cognitive scientists" to infer the interests and prior knowledge of new readers.

One important latent property of social groups is their preferences—what its members like and dislike. From an early age, children expect certain preferences to indicate social connection. For example, even preverbal infants expect that two people who like the same foods will interact positively with one another (Liberman et al., 2014) and that two people who interact positively with one another will like the same foods (Liberman et al., 2016). Infants' early-emerging abilities to reason about preferences may help solve evolutionarily ancient, life-or-death problems, such as selecting foods that are safe to eat.

However, humans can also learn to attach meaning to an incredible variety of seemingly arbitrary preferences. For example, we may feel a deep kinship towards someone who likes the same obscure movie as us (Vélez et al., 2019) and even create groups of people who like the same movie (e.g., fanclubs), even though movies have only existed for an eye-blink in our evolutionary history. How do seemingly arbitrary preferences come to have meaning?

A rich body of work has explored how humans transmit information about generalizable properties through language (Tessler & Goodman, 2019). Generic statements enable children to learn about novel preferences among social groups that they have never encountered (e.g., "Gazorps eat bubbas"). Preschool-aged children even enforce these regularities and negatively evaluate group members who don't share this preference (Roberts, Gelman, & Ho, 2017). Despite the power of generic statements, however, they may not always be available, especially for groups that are rather novel or clustered around arbitrary preferences.

Here, we explore a second, complementary learning mechanism: People can also build representations of social groups *de novo* by harnessing patterns of available data in the environment. Past work suggests that the foundations of this ability are present early in life. For example, suppose that a child sees an experimenter draw a sample of five red balls from a box containing mostly white balls. Because this sample is unlikely to have occurred randomly (Xu & Garcia, 2008), children use this evidence to make rich inferences about the experimenter's communicative intent (Gweon et al., 2010) or even about her preferences (Kushnir et al., 2010). Children can also generalize these preferences selectively to other individuals: If *two* experimenters draw the same unlikely sample, children infer that people within the experimenters' social group, but not people outside it, will have the same preference (Diesendruck et al., 2015).

Thus, both generic language and patterns of observation can support inferences from groups to properties. However, past work suggests that children don't readily draw inferences in the reverse direction—namely, inferring an individual's group membership based on shared properties (Gelman et al., 1986; Vélez et al., 2018). In one such study, preschool-aged children were told that boys have a substance called "andro" in their bodies and that girls have "estro" in their bodies. Children readily inferred that other boys would also have andro in their bodies, but not that children who have andro are boys (Gelman et al., 1986). Given that the two inferences have highly similar logical structures, the asymmetry in children's responses may seem rather puzzling.

To explain this asymmetry, prior work raised the possibility that children may lack sufficient prior knowledge about the relevance of a given property to group membership (Gelman et al., 1986). In the absence of this prior knowledge, children may (appropriately) hesitate to make this inference, especially given sparse evidence. However, this explanation leaves open questions about what makes certain properties more relevant than others, and how children come to acquire this knowledge.

Here, we explore the possibility that children's inferences are constrained by their beliefs about how properties are distributed across groups. In particular, inferring group membership based on shared properties may not be licensed if one believes that all groups are equally likely to have the property. More concretely, suppose that both boys and girls produce andro and estro. (This is, in fact, closer to the ground truth. Women actually secrete more androgen than they do estrogen; see Burger, 2002.) In this case, having andro would be consistent within each group, and it would thus generalize to other group members. However, it would not be diagnostic of group membership.

Critically however, even in the absence of specific prior knowledge about groups and properties, children may be able to flexibly infer group membership when they have appropriate evidence to support this inference. The present work explores how children and adults use minimal statistical information about social groups to infer the preferences and group membership of new individuals. Across two experiments, we compare the intuitions of adults and 4- and 5-year-old children to the predictions of a Bayesian computational model that formalizes these inferences. Computational and developmental approaches have been productively combined to study how children draw inferences about object categories, word labels, and even individual preferences (Xu & Tenenbaum, 2007; Lucas et al., 2014; Gweon et al., 2010). Our work provides an empirical and computational framework for understanding how children build structured, generalizable representations of social groups from sparse data.[1]

Experiment 1

This experiment examines how preschool-aged children and adults use statistical information about social groups to infer the preferences and group membership of new individuals. In contrast to prior work (Gelman et al., 1986; Vélez et al., 2018), we hypothesized that children may even be able to use individuals' preferences to infer their group membership when they have appropriate evidence to support this inference. We used a Bayesian beta-binomial model to formalize participants' inferences.

Methods

Participants 262 adults (average N = 44/condition) participated in this study online on Amazon Mechanical Turk. An additional 26 adults were excluded due to a technical error.

97 children (4- and 5-year-olds, average N = 16/condition) were recruited from local preschools and children's museums. An additional 6 children were excluded due to peer interference (2), experimenter error (2), or failure to comply with the task (2). In all experiments, children participated with the informed consent of a parent or legal guardian.

[1] Study materials can be found at: osf.io/ak83b/

Procedure Children watched a Keynote presentation about the preferences of novel agents called Gazorps (stimuli were adapted from Vélez et al., 2018; Vélez et al., 2019). Figure 1A–B shows representative stimuli. Gazorps could belong to one of two minimal groups—the "red team" or the "blue team"—as indicated by the color of their uniforms. The presentation depicted a classroom scene. The back of the classroom had a red and blue table where each team would sit; both tables were empty at the start of the task. The front of the classroom had two baskets, each containing a different novel fruit (bubbas and kikis).

Children watched Gazorps from each team choose their favorite fruit, one by one. Each Gazorp appeared in the center of the screen, approached the basket containing their favorite fruit, then returned to the center of the screen holding the fruit. In the next slide, the Gazorp would appear sitting at its team's table with its chosen fruit, and a new Gazorp would appear in the center of the screen. To keep children engaged, the experimenter sang a song that described the events on the screen; Figure 1A shows part of this process and accompanying lyrics. Children repeated this process until they had observed four members from each team.

Tasks In the critical test question, children were introduced to a new Gazorp (hence the "target"). Children were assigned to two between-subjects tasks, which differed in the inference that children were asked to make of the target (Figure 1B). In the *Infer Group* task, children learned what the target likes and were asked which team it belongs to. The target carried its favorite fruit and did not have a uniform. The experimenter explained that "this Gazorp lost its shirt, so we don't know what team this Gazorp is on." The experimenter then asked: "Is this Gazorp on the red team or the blue team?"

In the *Infer Preference* task, children learned which team the target belongs to and were asked what it likes. The target wore its team's uniform and did not have a fruit. The experimenter explained that "This Gazorp doesn't have a snack yet, so we don't know what this Gazorp likes to eat." The experimenter then asked: "Does this Gazorp like kiki or bubba?"

Conditions Within each task, we assigned children to three between-subjects conditions. These conditions differed in how preferences were distributed across the two teams and, thus, affected the evidence that children had available to support their inferences about the target. Figure 1C shows representative distributions in each condition.

In the *Unique & Consistent* condition, each group unanimously liked a different fruit. In the two remaining conditions, one group's preferences were unanimous, while the other group's preferences were heterogeneous. Critically, these conditions differed in the preferences (Infer Group) or group membership (Infer Preference) of the target. In the *Unique* condition, the target had a preference that was unique to the heterogeneous group (Infer Group) or belonged to the heterogeneous group (Infer Preference). In the *Consistent* condition, the target had a preference that was shared across both groups (Infer Group) or belonged to the unani-

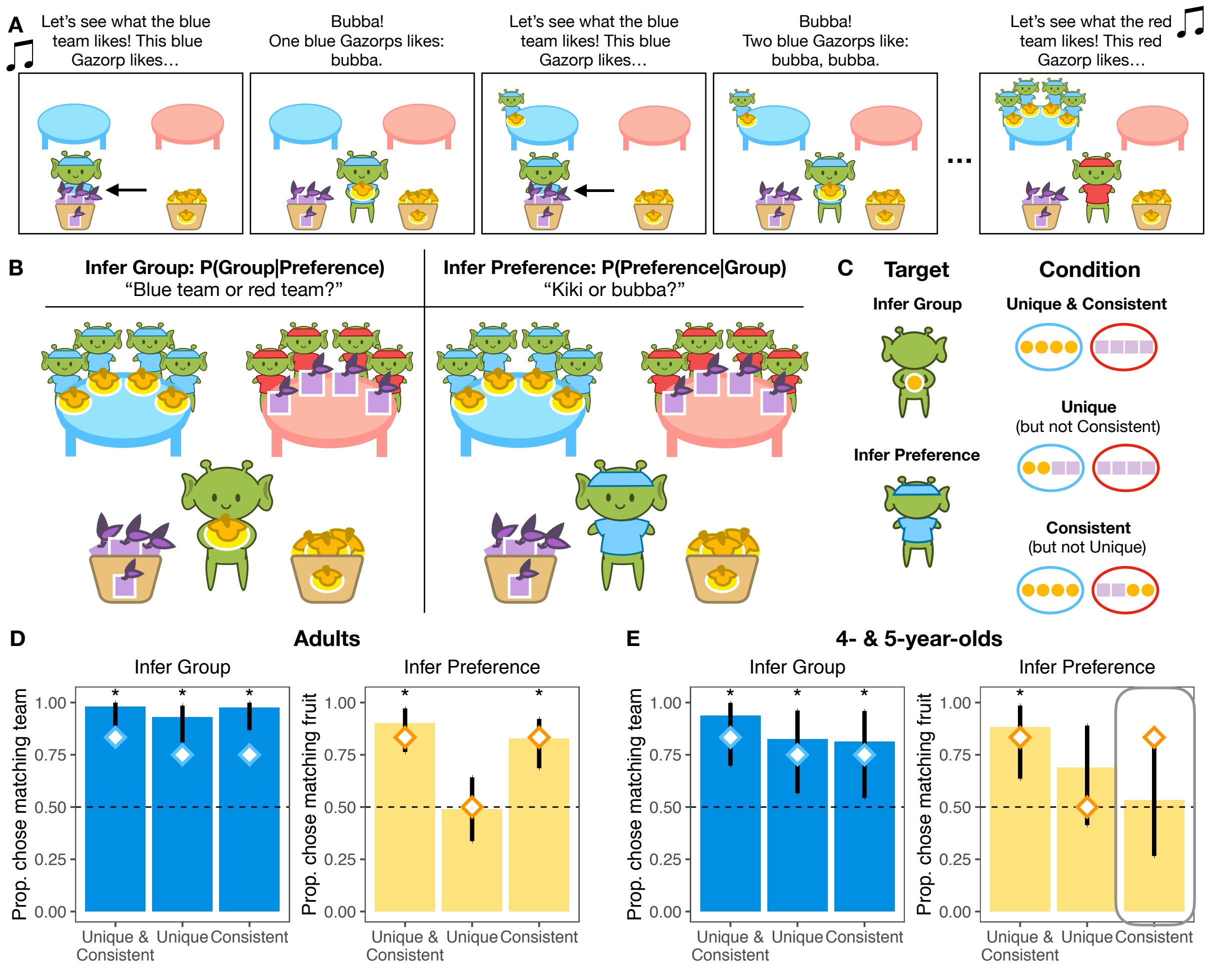

Figure 1: Experiment 1. (a) Procedure: Example slides are shown on the bottom, with the experimenter's script on top. (b) Test questions. (c) Conditions: Blue and red ovals represent each team's table during the final test question. Yellow circles represent bubbas, purple squares represent kikis. (d–e) Results: Proportion of (d) adults and (e) children who selected the matching team (Infer Group) or fruit (Infer Preference) by condition. Error bars denote 95% CI; $*$ shows significant within-condition contrasts (two-tailed binomial test, $p < 0.05$). Diamonds show the average model-predicted probability that the target belongs to the matching team (Infer Group) or likes the matching fruit (Infer Preference). The light gray square highlights a discrepancy between the model predictions and children's responses.

mous group (Infer Preference).

Counterbalancing Within each of these conditions, we designated one group as the reference group and one fruit as the reference item. For example, in the Consistent condition depicted in Figure 1B, liking bubba (yellow circles) is consistent within the blue team, but not unique to the blue team. Thus, the blue team is the reference group, and bubba is the reference item. We counterbalanced the reference group and reference item, as well as the locations of the teams and baskets, to create 16 counterbalanced variants of the stimuli. Each child saw a unique variant.

Adults completed a streamlined version of the task, where each team's preferences were shown in a single static image and all superficial aspects of the task were randomized.

Model

We formalized our predictions using a Bayesian beta-binomial model. Importantly, the model assumes that the red and blue teams' preferences are independent; thus, its observations of one team have no bearing on its beliefs about the preferences of the other team.

The model describes each team's choices (D) as a random sample from a binomial distribution. For example, the likelihood of the blue team's choices is:

$$P(D|\theta_b) = \binom{n_b}{k_b} \theta_b^k (1-\theta_b)^{n_b-k_b}$$

where n_b is the total number of Gazorps on the team, k_b is the number of Gazorps who chose the reference item, and θ_b is a latent variable that describes the blue team's underlying preferences. This variable is akin to the weight of a biased coin, and each team member's choice is akin to a coin flip.

The model infers each team's preferences from the observed choices. We describe the model's posterior beliefs about each team's preferences using a Beta distribution:

$$\theta_b|D \sim \text{Beta}(\alpha_0 + k_b, \beta_0 + n_b - k_b)$$

where $\alpha_0 = \beta_0 = 1$ reflect that the model has a uniform prior belief over the weights.

In the Infer Preference task, the model generates a new prediction about the target by sampling from the posterior. Figure 1D–E shows the mean of the posterior distribution for the target's team ($\mathbf{E}[\theta_b|D]$; orange diamonds).

In the Infer Group task, the model observes the target's choice (c) and entertains two competing hypotheses: The target could belong to the blue team ($z_b = 1$) or to the red team ($z_b = 0$). Because both teams are of the same size, these two hypotheses have the same prior probability ($P(z_b) = 0.5$). Thus, the model's beliefs about the target's group membership ($P(z_b|c,D)$) are driven by the likelihood of observing the target's preferences in each group. We describe the relationship between these beliefs using Bayes rule:

$$P(z_b|c,D) \propto P(z_b) \int_{\theta_z} P(c|\theta_z)P(\theta_z|D)d\theta_z$$

where $\theta_z = \theta_b$ if the target is on the blue team, and $\theta_z = \theta_r$ if the target is on the red team. Figure D–E shows the mean of this posterior distribution ($\mathbf{E}[z_b|c,D]$; blue diamonds).

Results

Adults Adults' responses closely followed the pattern of model predictions (Figure 1C). In the Infer Group task, adults readily inferred the target's group membership based on its preferences in the Unique & Consistent (47/48; two-tailed binomial test, 95% CI: $[.89, 1]$, $p < .001$), Unique (40/43; two-tailed binomial test, 95% CI: $[.81, .99]$, $p < .001$), and Consistent conditions (39/40; two-tailed binomial test, 95% CI: $[.87, 1]$, $p < .001$).

In the Infer Preference task, adults readily inferred the target's preferences based on its group membership in the Unique & Consistent (36/40; two-tailed binomial test, 95% CI: $[.76, .97]$, $p < .001$) and Consistent conditions (38/46; two-tailed binomial test, 95% CI: $[.69, .92]$, $p < .001$), but not in the Unique condition (22/45; two-tailed binomial test, 95% CI: $[.34, .64]$, $p = 1$).

Children In the Infer Group task, children readily inferred the target's group membership based on its preferences in the Unique & Consistent (15/16; two-tailed binomial test, 95% CI: $[.7, 1]$, $p < .001$), Unique (14/17, two-tailed binomial test, 95% CI: $[.56, .96]$, $p = 0.01$), and Consistent conditions (13/16, two-tailed binomial test, 95% CI: $[.54, .96]$, $p = 0.02$).

In the Infer Preference task, children in the Unique & Consistent condition expected the target to like the same item as its teammates (15/17; two-tailed binomial test, 95% CI: $[.64, .99]$, $p = .002$), and those in the Unique condition did not systematically expect the target to like the item that was uniquely (but not consistently) liked by its teammates (11/16, proportion = .69; two-tailed binomial test, 95% CI: $[.41, .89]$, $p = .2$). While these were in line with model predictions and similar to Exp.1a, we observed a discrepancy between children's responses and the model prediction in the Consistent condition (Figure 1B, highlighted; children did not systematically predict the target's preference (8/15, proportion = .53; two-tailed binomial test, 95% CI: $[.27, 79]$, $p = 1$). We return to this discrepancy below.

Overall, we find that both children and adults use statistical information to infer the group membership and preferences of new individuals. Our results suggest that children can use preferences to infer group membership when they have the appropriate evidence to support it, even in subtle cases where the target's preference is consistently, but not uniquely, associated with a particular group. However, children's inferences about the target's preferences deviated from the model predictions in the Consistent condition. In this case, children did not expect the target to share the same preference as its teammates, even though the target's team unanimously liked the same fruit.

Note that there were two conditions where the preferences of the target's team were unanimous: Unique & Consistent and Consistent. Unlike adults, however, children generalized the team's preferences in the former condition, but not the latter—and the only difference between the two conditions is the preferences of the opposing team. Critically, the model assumes each team's preferences were generated through independent random processes; thus, the opposing team has no bearing on the model's beliefs about the target.

One possible explanation for this discrepancy is that children's responses reflect not only their hypotheses about how preferences are distributed within each group, but also *overhypotheses* about how groups are generated (Kemp et al., 2007). For example, observing that the opposing team's preferences are heterogeneous may lend support to overhypotheses where preferences are heterogeneous within groups—and thus temper children's expectations about how consistent preferences are within the target's team. Adults, however, might have assumed that these two groups are independent; this could reflect their beliefs about the pragmatics of the task setup (i.e., that team colors must be relevant to preferences) or their broader beliefs about minimal novel groups. In Experiment 2, we present ongoing work that tests this possibility.

Experiment 2

This experiment tests whether adult's and children's inferences about the target's preferences are influenced by the size of the target's team and by the size of the opposing team. Data collection is ongoing; the results below are preliminary.

Methods

Participants 300 adults (N = 50/condition) participated in this study on Amazon Mechanical Turk. Data collection with

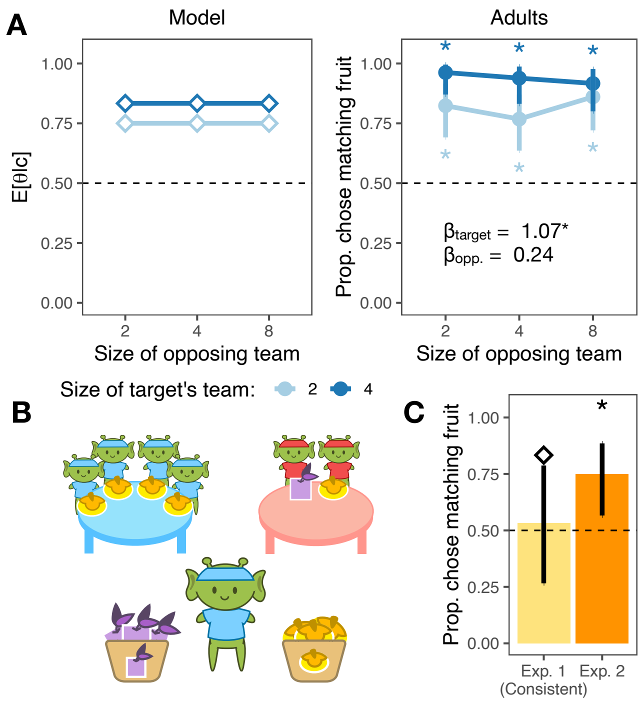

Figure 2: Experiment 2. (a) Model predictions. Diamonds show the mean probability of the target liking the matching fruit, as a function of the size of the target's team (line color) and the size of the opposing team (x-axis). (b–c) Results: (b) Proportion of adults who selected the matching fruit. (c) Proportion of children who selected the matching fruit (orange bar), compared to children's responses and model predictions in Experiment 1 (light yellow bar; Infer Preference task, Consistent condition). Error bars show 95% CI; $*$ denotes $p < 0.05$ (two-tailed binomial test).

children is ongoing; 32 4- and 5-year-old children have been recruited for the condition described here (see Procedure).

Procedure This task was identical to the Consistent condition of the Infer Preference task in Experiment 1. As in Experiment 1, participants inferred which fruit the target likes based on its group membership. The target's team unanimously liked one fruit, while the opposing team's preferences were evenly split between the two fruits.

Critically, we independently manipulated the size of the target's team (2 or 4 members) and of the opposing team (2, 4, or 8 members) to create 6 between-subjects conditions. We have completed data collection in adults in all conditions. Below, we report preliminary results on children's responses in a single condition, in which the target's team has 4 members and the opposing team has 2 members (Figure 2B).

Results

Figure 2A shows model predictions. The model assumes that the two teams' preferences are generated through independent processes. Thus, the model's beliefs are not affected by the size of the opposing team. As the size of the target's team increases, the model becomes increasingly confident that the target will like the matching fruit. Thus, the model predicts a main effect of the size of the target's team, no main effect of the opposing team, and no interactions between the two.

Adults' responses followed model predictions closely. We used a logistic regression to predict participants' responses (match or non-match) as a function of the size of the target's team and of the opposing team (response ˜ target_team $*$ opp_team). Consistent with model predictions, participants were more likely to generalize the preferences of the target's team to the target when the team had more members (main effect of target_team: $\beta = 1.07$, $SD = 0.45$, $z = 2.37$, $p = 0.02$). We observed no effect of the size of the opposing team ($\beta = 0.24, SD = 0.22$, $z = 1.05$, $p = .3$) and no interaction between team sizes ($\beta = -0.09$, $SD = 0.08$, $z = -1.12$, $p = .26$).

Conversely, based on Experiment 1, we predicted that children's inferences about the target's preferences would also be sensitive to the preferences of the opposing team. Children may expect the target to like the same item as its teammates when the opposing, heterogeneous team is smaller (and thus contributes less evidence). Our preliminary results suggest that this is indeed the case (30/32, two-tailed binomial test, 95% CI: $[.57, .89]$, $p = .007$; Figure 2C).

Our results provide preliminary evidence that children and adults interpret statistical information about social groups in rational, yet distinct, ways. Adults' responses in Experiments 1 and 2 are consistent with a model that considers the two groups' preferences as completely independent of one another. Conversely, children's inferences are not only sensitive to the preferences of the target's team, but also to the composition (Experiment 1) and size (Experiment 2) of the opposing team. These results suggest that children pool evidence across social groups to make inferences about individuals. Future work will continue to explore these differences and make future refinements to the model, in order to precisely characterize this pattern of developmental change.

General Discussion

The current work finds that preschool-aged children and adults can use minimal statistical information about how preferences are distributed across social groups to infer the preferences and group membership of new individuals. Overall, Experiment 1 found that 4- and 5-year-old children's intuitions aligned closely with adults' responses and with the predictions of a computational model. However, we did find instances where children's responses diverged. In particular, children's inferences about the preferences of new group members were influenced both by the composition (Experiment 1) and size (Experiment 2) of the opposing group.

These results naturally raise the question of why children and adults had different intuitions. One possibility is that such discrepancies reflect particular pragmatic assumptions that adults have in the context of an experiment. Participants were given information about two teams and two fruits to answer questions about a new individual; thus, even without prior knowledge of the specific groups and properties in this task, adults may have answered the questions with a stronger expectation that shared properties must be relevant to group membership. It is an open question whether these assumptions can be weakened when there is explicit information that group memberships is arbitrary and clearly unrelated to properties. Ongoing work tests this possibility.

We tested adults' and children's ability to make these inferences in a minimal, supportive context, where participants have equal opportunities to learn about all social groups and where the properties that they learn about are neutral. This approach provides a window into our ability to draw judgments about social groups and their members. It is especially impressive that, unlike previous work (Gelman et al., 1986; Vélez et al., 2018), even young children were able to make judgments in both directions. Rather than simply being unable to infer groups from preferences, children in our study drew flexible inferences about new individuals when they had access to relevant evidence. This simplified task is amenable to a computational modeling approach, allowing us to characterize children's inferences in precise, quantitative terms.

Moving forward, however, it is important to consider how the information that children actually receive from their social environment differs from the information they received in this task. One key difference is the child's point of view. Children learn about many categories in order to make sense of the world around them (Gelman & Roberts, 2017); children can learn about cars, dogs, and blickets without *being* cars, dogs, or blickets themselves. But social groups are categories that children can be a part of—and thus, the mere presence of social groups creates boundaries between us and them, ingroup and outgroup (Dunham, 2018).

These boundaries may impose constraints on what children can and want to learn from their social environment. Children do not sample information evenly from their social environment, but are instead biased to seek information that casts their ingroup in a positive light (Over et al., 2018). More importantly, however, children also do not have equal opportunities to learn about ingroup and outgroup members. For example, in U.S. contexts, neighborhoods tend to be segregated by race and income (Reardon & Owens, 2014); thus, many children are raised in communities that are largely composed of families like theirs. An emerging line of work suggests that children's cultural and social environment may influence their conceptual knowledge about social groups, such as whether individuals' preferences can deviate from the group's (Roberts et al., 2018) and who should and will be friends (Eason et al., 2019; Roberts, Williams, & Gelman, 2017).

Overall, these results are cause for both hope and concern. On one hand, these results suggest that children can build rich, generalizable representations of social groups from just a few observations and, critically, that these representations reflect the evidence they've observed. On the other hand, children in our task learned about fairly neutral properties of social groups that they're not a part of. In naturalistic contexts, these inferences might easily go awry, supporting the formation of biases and stereotypes. Our results are thus relevant to understanding how prejudice forms—and how it might be corrected with evidence.

Acknowledgments

This work was supported by an NIH F99/K00 D-SPAN Fellowship to NV and by a McDonnell Scholars Award and Jacobs Foundation Research Fellowship to HG.

References

Burger, H. G. (2002). Androgen production in women. *Fertility and sterility*, *77*, 3–5.

Diesendruck, G., Salzer, S., Kushnir, T., & Xu, F. (2015). When choices are not personal: The effect of statistical and social cues on children's inferences about the scope of preferences. *Journal of Cognition and Development*, *16*(2), 370–380.

Dunham, Y. (2018). Mere membership. *Trends in Cognitive Science*, *22*(9), 780–793.

Eason, A. E., Kaiser, C. R., & Sommerville, J. A. (2019). Underrepresentation and the perception of others' racial attitudes. *Social Psychological and Personality Science*, *10*(6), 757–767.

Gelman, S. A., Collman, P., & Maccoby, E. E. (1986). Inferring properties from categories versus inferring categories from properties: The case of gender. *Child Dev*, 396–404.

Gelman, S. A., & Roberts, S. O. (2017). How language shapes the cultural inheritance of categories. *PNAS*, *114*(30), 7900–7907.

Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *PNAS*, *107*(20), 9066–9071.

Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical bayesian models. *Developmental science*, *10*(3), 307–321.

Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological science*, *21*(8), 1134–1140.

Liberman, Z., Kinzler, K. D., & Woodward, A. L. (2014). Friends or foes: Infants use shared evaluations to infer others' social relationships. *J. Exp. Psychol. Gen.*, *143*(3), 966.

Liberman, Z., Woodward, A. L., Sullivan, K. R., & Kinzler, K. D. (2016). Early emerging system for reasoning about the social nature of food. *PNAS*, *113*(34), 9480–9485.

Lucas, C. G., Griffiths, T. L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T., ... Hu, J. (2014). The child as econometrician: A rational model of preference understanding in children. *PloS one*, *9*(3).

Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual review of psychology*, *51*(1), 93–120.

Over, H., Eggleston, A., Bell, J., & Dunham, Y. (2018). Young children seek out biased information about social groups. *Developmental Science*, *21*(3), e12580.

Reardon, S. F., & Owens, A. (2014). 60 years after brown: Trends and consequences of school segregation. *Annual Review of Sociology*, *40*, 199–218.

Rhodes, M., & Baron, A. (2019). The development of social categorization. *Annu. Rev. Dev. Psychol.*, *1*, 359–386.

Roberts, S. O., Gelman, S. A., & Ho, A. K. (2017). So it is, so it shall be: Group regularities license children's prescriptive judgments. *Cognitive Science*, *41*, 576–600.

Roberts, S. O., Guo, C., Ho, A. K., & Gelman, S. A. (2018). Children's descriptive-to-prescriptive tendency replicates (and varies) cross-culturally: Evidence from china. *Journal of Experimental Child Psychology*, *165*, 148–160.

Roberts, S. O., Williams, A. D., & Gelman, S. A. (2017). Children's and adults' predictions of black, white, and multiracial friendship patterns. *Journal of cognition and development*, *18*(2), 189–208.

Tessler, M. H., & Goodman, N. D. (2019). The language of generalization. *Psychological review*, *126*(3), 395.

Vélez, N., Bridgers, S., & Gweon, H. (2019). The rare preference effect: Statistical information influences social affiliation judgments. *Cognition*, *192*, 103994.

Vélez, N., Wu, Y., & Gweon, H. (2018). Consistent but not diagnostic: Preschooler's intuitions about shared preferences within social groups. In *Proceedings of the cognitive science society.*

Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *PNAS*, *105*(13), 5012–5015.

Xu, F., & Tenenbaum, J. B. (2007). Word learning as bayesian inference. *Psychological review*, *114*(2), 245.