

Rational learners and parochial norms

Scott Partington^a, Shaun Nichols^b, Tamar Kushnir^{c,*}

^a Department of History and Philosophy of Science, University of Cambridge, Free School Lane, Cambridge, CB2 3RH, United Kingdom

^b Department of Philosophy, Cornell University, Goldwin Smith Hall, Ithaca, NY 14853, United States of America

^c Department of Psychology & Neuroscience, Duke University, 417 Chapel Drive, Box 90086, Durham, NC 27708, United States of America

ARTICLE INFO

Keywords:

Norms psychology
Social learning
Statistical learning
Cognitive development

ABSTRACT

Parochial norms are narrow in social scope, meaning they apply to certain groups but not to others. Accounts of norm acquisition typically invoke *tribal biases*: from an early age, people assume a group's behavioral regularities are prescribed and bounded by mere group membership. However, another possibility is *rational learning*: given the available evidence, people infer the social scope of norms in statistically appropriate ways. With this paper, we introduce a rational learning account of parochial norm acquisition and test a unique prediction that it makes. In one study with adults ($N = 480$) and one study with children ages 5- to 8-years-old ($N = 120$), participants viewed violations of a novel rule sampled from one of two unfamiliar social groups. We found that adults judgments of social scope – whether the rule applied only to the sampled group (parochial scope), or other groups (inclusive scope) – were appropriately sensitive to the relevant features of their statistical evidence (Study 1). In children (Study 2) we found an age difference: 7- to 8-year-olds used statistical evidence to infer that norms were parochial or inclusive, whereas 5- to 6-year olds were overall inclusive regardless of statistical evidence. A Bayesian analysis shows a possible inclusivity bias: adults and children inferred inclusive rules more frequently than predicted by a naïve Bayesian model with unbiased priors. This work highlights that tribalist biases in social cognition are not necessary to explain the acquisition of parochial norms.

1. Introduction

Parochial norms are narrow in social scope, meaning they apply to certain groups but not others. Turnbull (1972) reports a vivid example: the Mbuti thought it was wrong to steal from each other, but they thought it was fine to steal from people in villages. Drawing on cases of this nature, anthropologists have long emphasized that many moral rules are parochial (e.g., Read, 1955, p. 255; Snare, 1980, p. 364). Likewise, contemporary psychological research often highlights the importance of the ‘moral circle’ or ‘tribe’ in judgment and decision-making (e.g., Graham, Waytz, Meindl, Iyer, & Young, 2017; Greene, 2013). Interpreted richly, this literature points to a stark conclusion: “tribal bias is a natural and nearly ineradicable feature of human cognition” (Clark, Liu, Winegard, & Ditto, 2019, p. 587). For those who favor inclusive norm systems, this interpretation highlights the troubling aspects of parochialism. But mundane cases of parochial norms also abound. Consider norms tied to institutional roles (Noyes, Dunham, Keil, & Ritchie, 2021). Only cashiers are allowed to take money from the register, only construction workers are allowed to operate jackhammers, and so on. Indeed, across history and cultures,

most rules of behavior apply to less than everyone. In this broad sense, parochialism is the norm.

With this paper, we investigate parochial norm acquisition using the methods familiar to computational cognitive science (cf. Marr, 1982). In particular, we take the first steps toward developing a *rational learning* account of this process, according to which principles of rational statistical inference play a central role (cf., Gopnik & Wellman, 2012; Xu & Kushnir, 2012, 2013; Perfors, Tenenbaum, Griffiths, & Xu, 2011; Xu, 2019). Our account stands in contrast to prevailing views from cultural evolutionary, social, and developmental psychology which emphasize the role of group-based biases in norm acquisition (cf. Chalik & Rhodes, 2020; Chudek & Henrich, 2011; Roberts, Gelman and Ho, 2017). We present experimental evidence and computational analyses that show principles of rational statistical inference can explain the acquisition of parochial rules. Thus, our work makes clear that a deep-rooted tribal bias is not necessary to explain why people are especially prone to acquiring parochial norms. Given evidence consistent with parochial rule, learners who follow simple principles of rational statistical inference will acquire a parochial rule.

The rest of the paper is structured as follows. First, we specify the

* Corresponding author.

E-mail addresses: sp933@cam.ac.uk (S. Partington), sbn44@cornell.edu (S. Nichols), tamar.kushnir@duke.edu (T. Kushnir).

explanatory targets for this project (1.1. ‘The acquisition of parochial norms’), survey the existing accounts which emphasize the role of group-based biases (1.2. ‘Tribal biases and parochial norms’) and detail the proposed rational learning account (1.3. ‘Rational learning and parochial norms’). Next, we test a unique prediction of the rational learning account across an experiment with adults (2. ‘Study 1’) and children ages 5- to 8-years-old (3. ‘Study 2’). The adults and older children displayed the predicted sensitivity to statistical evidence. Younger children showed some sensitivity to evidence in their open responses, but their judgments were *inclusive* – they generalized to others outside the sampled group regardless. With these findings in hand, we compare the experimental results against predictions from a naïve Bayesian model (4. ‘Computational analyses’). The computational analyses show that both adults and children displayed an *inclusive bias*: given the evidence presented in our studies, human learners were more likely to infer a wide-scope social norm than an unbiased Bayesian agent. Lastly, we discuss the implications of our findings for research in social learning and cognitive development, along with the limitations of the current research (5. ‘Discussion’).

1.1. The acquisition of parochial norms

A complete account of parochial norm acquisition would detail why naïve learners acquire narrow-scope, parochial rules rather than wide-scope, inclusive rules. As we use the term here, parochial norms are narrow in *social scope* in particular, meaning they apply to some social groups but not others. One key point of interest is a learning theoretic question: given the available evidence, how do people come to have parochial rules rather than inclusive ones?

One possibility is that testimonial instruction can bridge this gap: perhaps, people learn the social scope of norms simply because they are told the social scope of each candidate norm. However, there are several reasons for thinking that learning from explicit instruction about the exact scope of norms cannot be a complete explanation. First, outside of legal contexts, the full scope of a normative rule is rarely articulated. Yet, as proponents of moral nativism like Mikhail (2011) often emphasize, people naturally arrive at fine-grained and sophisticated distinction about the scope of moral rules (e.g., the duty to rescue, prohibitions on unreasonable risks, the doctrine of double effect, etc.). In the vast majority of cases, these sorts of distinctions are not explicitly taught to children. Analysis of parent-child conversation suggests that young children rarely receive explicit instruction about moral rules, and, when they do, it typically takes the form of injunctions (e.g., ‘no!’ or ‘don’t!’) (Wright & Bartsch, 2008). Hence, learning from explicit instruction cannot be a complete explanation of how people acquire parochial norms (Dwyer, Huebner, & Hauser, 2009, p. 6).

Nevertheless, children as young as 3-years-old acquire a subtle understanding of normative rules (Mikhail, 2007, p. 144; Pellizzoni, Siegal, & Surian, 2010), even in one-shot settings (Rakoczy, Warneken, & Tomasello, 2008). Moreover, children often accomplish this without the aid of explicit negative evidence, or evidence about when the rule does *not* apply. On this point there is a natural analogy with language acquisition. Just as children can learn the meaning of the word ‘dog’ without enumeration of all the ‘not dog’ objects in the world (Xu & Tenenbaum, 2007), children can also learn to whom the rule applies without enumeration of all the individuals to whom it doesn’t.

These core phenomena suggest that an adequate explanation of the acquisition of parochial norms should meet two minimal criteria:

SPARSE EVIDENCE: the learning principle explains how norm learners can infer social scope from one- or few-shot scenarios.

NO NEGATIVE EVIDENCE: the learning principle explains how norm learners infer social scope without the aid of explicit negative evidence.

By “explicit negative evidence” we mean explicit instruction about the subjects to whom a candidate rule does not apply. For example, “adults can stay up late” features only positive evidence about who can

stay up late whereas “adults can stay up late *and children cannot stay up late*” features positive and negative evidence.

Extant accounts posit a principle that accounts for both SPARSE EVIDENCE and NO NEGATIVE EVIDENCE. In recent years, a number of findings have been understood as convergent evidence for what we will label the *automatic group bias* account. According to this account, the norm learner’s mere representations of a social group foster an automatic tendency to consider the group’s behavioral regularities as both prescribed and bounded by group membership (cf., Roberts et al., 2017, p. 593; Chalik & Rhodes, 2020, p. 80). In the following section, we survey the motivations and evidence for this hypothesis, and we detail how an automatic group bias can explain the core phenomena that we laid out above.

1.2. Tribal biases and parochial norms

A common way to motivate the automatic group bias account starts with an evolutionary argument that norm acquisition evolved to facilitate within-group coordination (see, e.g., Boyd & Richerson, 2009; Rand & Nowak, 2013). Chudek and Henrich (2011, p. 219) propose that such pressures resulted in a coevolutionary selection for an ‘ethnic psychology,’ whereby social group members preferentially interact with and learn from others who share arbitrary ethnic markers. Chudek and Henrich write (2011, p. 219–220): “...the increasing fitness-relevance of coordination creates genetic selection pressures for skill at recognizing and representing the most common behaviors, beliefs, or strategies in one’s community and for dispositions to adopt them or even internalize them as proximate motivations or heuristics.”

On this account, the cognitive consequence is a host of biases that direct learning along ethnic lines, many of which are early emerging and persist across the lifespan (e.g., Efferson, Lalive, & Fehr, 2008; Kinzler, Dupoux, & Spelke, 2007; Kinzler, Shutts, DeJesus, & Spelke, 2009). Perhaps the most striking example are minimal group biases (Tajfel, 1970; Diehl, 1990; Dunham, 2018). In studies involving the minimal groups paradigm, children and adults display ingroup favoritism even when randomly assigned to previously unfamiliar social groups based on arbitrary cues (e.g., a label, ‘the red group’). For example, children are more willing to trust the testimony of ingroup members even in the absence of any prior knowledge about the group (MacDonald, Schug, Chase, & Barth, 2013).

Richly interpreted, such minimal group biases are thought to result from “a more abstract set of coalitional principles and expectations concerning group membership” (Dunham, 2018, p. 788). On this interpretation, minimal group biases not only reflect an *ingroup* bias, but also reflect a *mere group* bias resulting from a deeply rooted, and perhaps innate, set of expectations about group behavior. Here, the claim is that a learner’s intuitions about relations of trust, coordination, cooperation, and reciprocity are intrinsically tied to her representations of the social groups themselves (Yamagishi & Mifune, 2016)—a social learner cannot represent a social group without assuming that certain relations apply between members of the group. Indeed, such ‘coalitional heuristics’ are often marshaled to explain children’s third-party judgments about group behaviors (cf. Chalik and Rhodes, 2018), even as early as infancy (Bian, Sloane, & Baillargeon, 2018; Jin & Baillargeon, 2017).

Focusing on the specific case of norm learning, in a series of recent studies Roberts and colleagues (2017, 2017b; 2018, 2019) have used a novel groups paradigm to examine the link between such coalitional heuristics and learners’ mere representations of social categories. In this research, participants make third-party evaluative judgments about two previously unfamiliar social groups (i.e., animated characters labelled ‘Hibbles’ and ‘Glerks’). The typical finding is that children use generic descriptions of group behavior to infer how individual members ought to behave. For example, if told that Hibbles eat red berries and Glerks eat purple berries, the majority of children ages 3- to 12-years-old will say that it is ‘not okay’ for Glerks to eat red berries; the younger children in particular seem especially prone to make this response. In explaining

these results, Roberts and colleagues suggest that “group regularities may exert influence by rather automatically fostering a negative evaluative stance” (Roberts et al., 2017, p. 593).

Likewise, Chalik and Rhodes (2020, p. 80) consider these studies and others involving novel groups to indicate “children assume all norms (not just moral norms) are bounded by some type of category... once they learn that something falls under the scope of a norm... they assume there is a boundary on to whom it applies.” On this account, the critical importance of behavioral regularities is reinforced by the implications of generic claims about social groups (cf., Gelman et al., 2010). For example, Chalik and Rhodes (2020, p. 84) suggest “hearing the sentence ‘Italians eat pasta’ can lead a child to believe that people from Italy are an objectively different kind of people from people born in other places.” Chalik and Rhodes (2020) also emphasize the importance of “known norms” that carry evaluative content (e.g., helping vs. harming) and show that children make sophisticated inferences about coalitional relations between group members (e.g., who is obligated to help who). They favor an explanation for these findings in terms of domain-general category biases in learning, such that social learning relies on the expectation that social categories prescribe the behaviors of category members (Chalik & Rhodes, 2020, p. 80). As such, a tendency toward acquiring parochial norms is thought to be rooted in the learner’s mere representations of social groups.

To summarize, the automatic group bias hypothesis receives support from converging evidence across several research traditions: cultural evolution, social psychology, and cognitive development. In particular, the learning accounts put forth by Roberts and colleagues (2017, 2018, 2019) and Chalik and Rhodes (2013, 2014, 2018, 2020) suggest a learning principle which meets both the SPARSE EVIDENCE and NO NEGATIVE EVIDENCE challenges. If the acquisition of parochial rules is automatically fostered by linking a group’s representation and its behavioral regularities, then parochial norms can be acquired in one- or few-shot scenarios without the aid of explicit negative evidence.

Research in support of the automatic group bias account thoroughly demonstrates the importance of social group representations in norm acquisition. In particular, it reveals how learners of all ages are keenly attuned to signifiers of group membership, such as group labels or generic descriptions of group behaviors. Further, this research shows that an important function of social group representations is helping to guide predictions about how group members are likely to behave (Birnbaum, Deeb, Segall, Ben-Eliyahu, & Diesendruck, 2010; Chalik & Rhodes, 2020). We agree that these are key features that any alternative account of parochial norm acquisition will need to preserve. However, next we will introduce an alternative account of how learners relate their social group representations with the available evidence of social scope. In what follows, we motivate a rational learning account, detail its contents, and outline when and why it makes unique predictions from accounts which emphasize an automatic group bias.

1.3. Rational learning and parochial norms

A common motivation for rational learning accounts (Anderson, 1990) is that learning often involves making inferences that go well beyond our apparently limited evidence: “we build rich causal models, make strong generalizations, and construct powerful abstractions” despite our input data being “sparse, noisy, and ambiguous—in every way far too limited.” (Tenenbaum et al., 2011, p. 1279). Statistical inference under uncertainty is our most viable formal framework for explaining how this can be accomplished (Edelman & Shahbazi, 2011), and probabilistic accounts in particular have several normatively attractive features (see, e.g., Pettigrew, 2016). Indeed, in recent years, applying the formal framework of rational statistical inference to model learning under uncertainty has been a productive and influential approach to research in cognitive development (for reviews, see Xu, 2019; Xu & Kushnir, 2013; Perfors et al., 2011; Tenenbaum et al., 2011). Already, the rational learning framework has been applied to several

aspects of norm acquisition (Nichols, 2021; see, e.g., Ayars & Nichols, 2017, 2020; Nichols, Kumar, Lopez, Ayars, & Chan, 2016). Here we apply the framework to the case of parochial norm acquisition—in particular, to explain how learners can infer social scope on the basis of sparse and solely positive evidence. How do learners infer that a norm applies to some groups but not to others? To provide a computational-level analysis of this process (cf. Marr, 1982), we will need to specify the learner’s hypothesis space, the available evidence, and how the learner relates the evidence to the candidate hypotheses.

1.3.1. The hypothesis space

In the case of acquiring a parochial norm, the target is the belief that a candidate norm is narrow-scope, rather than wide-scope. Here, it is helpful to consider social scope in terms of a nested subset structure (Fig. 1), whereby a parochial norm applies to a smaller set of possible subjects than an inclusive norm. We can characterize a learner’s set of candidate hypotheses in a corresponding manner: $h_{\text{Parochial}}$ is the hypothesis that a given norm applies to a subset of possible subjects, whereas $h_{\text{Inclusive}}$ is the hypothesis that it applies to the whole population. In models using the Bayesian formalism, such hypotheses are interpreted as structured, symbolic representations to which learners assign different levels of certainty given their available evidence (Goodman et al., 2008). Indeed, like feature learning in general (cf. Austerweil & Griffiths, 2013), acquiring distinctions about the social scope of norms requires that learners relate evidence to richly structured representations. Hence, the Bayesian formalism is well-suited for modeling this sort of learning problem.

Given the importance of social group representations in norm acquisition, we reason that social category boundaries often suffice to create a hypothesis space that mirrors this subset structure. For example, when learning about a population of novel groups, learners consider $h_{\text{Parochial}}$ as applying to only Glerks, whereas $h_{\text{Inclusive}}$ applies to both Hibbles and Glerks. Since learners must consider a limited number of hypotheses in order for the statistical inference to remain tractable (Van Rooij, 2008), being attuned to informative signifiers of group membership remains critical. However, in contrast to the automatic group bias account, the rational learning account posits that simply noting the existence of two groups (for example, by labeling them) is not enough to elicit a parochial bias. For this, we need to detail the evidence available to learners and how learners relate the evidence to the candidate hypotheses.

1.3.2. The available evidence

There are many sources of evidence available to norm learners. One source of evidence is generic claims about group characteristics. As much recent work from Rhodes and colleagues demonstrates, children’s social learning is sensitive to generic formulations in language (e.g., of the form, ‘Italians eat pasta’) such that generic language about groups promotes greater generalizations in learning than non-generic language (see Chalik & Rhodes, 2020 for a review). Along these lines, children could hear about norms as general properties of groups (e.g. “Hibbles are not allowed to eat red berries”) or simply hear descriptions of group behavior in general terms (e.g. “Hibbles don’t eat red berries”) and infer that the description signals a norm (Roberts et al., 2017).

But children do not only receive generic evidence: instances of norm violations are sometimes labelled explicitly (e.g., “This is not okay,” “Don’t!”, etc.) and learners sometimes have to use this sort of evidence to infer the scope of the norm. Moreover, generic claims about groups are not always quantifiable in a way that is useful for a Bayesian learner, since they apply both in cases where the probability of the property being true of a given group member is high, and in cases where the probability is low (Prasada, Khemlani, Leslie, & Glucksberg, 2013). Even so, the connection between generics and Bayesian inference is a topic of ongoing research. For example, Tessler, Bridgers, and Tenenbaum (2020) attempt to quantitatively assess how much confirmation (in terms of single instances of evidence) generic claims confer onto

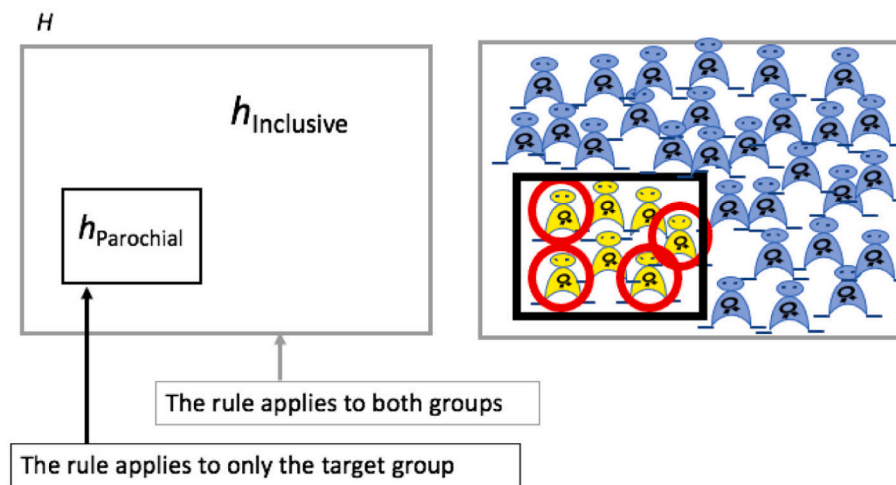


Fig. 1. Schematic depiction of a rational learner's hypothesis space for parochial norm acquisition (left) and how the hypotheses correspond to norm learning in the novel groups paradigm (right), where the red circles denote evidence of a norm violation. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

candidate hypotheses Likewise, preliminary studies show that adults and children learning new causal properties integrate generic language with statistical evidence, suggesting that even when language is present, it informs but does not overwhelm other sources of evidence in inductive generalization (Finiasz, Sheth, Karami, Gelman, & Kushnir, 2022; Kushnir & Gelman, 2016).

In the present research, we focus on evidence in the form of explicitly labelled rule violations (e.g., “This is against the rule”) sampled from a set of available examples. Such explicit rule violations are perhaps the most clear and direct form of evidence that bears on the social scope of a candidate norm. Thus, our work focuses on how learners acquire unequivocally prescriptive norms, rather than how descriptive norms acquire a prescriptive force (cf. Roberts et al., 2017).

1.3.3. The learning principle

In such cases, the key question is how learners overcome the problems of SPARSE EVIDENCE and NO NEGATIVE EVIDENCE. If the only rule violations apply exclusively to ribbon wearing Hibbles, yet some Glerks are also wearing ribbons, then, intuitively, we also find it likely that the Glerks are allowed to wear ribbons (Fig. 1), despite the fact we received no explicit negative evidence (i.e., of the form “This is not against the rule”).

The size principle can explain how people can infer social scope from sparse and not explicitly negative evidence. In a Bayesian framework, the size principle describes how the smallest hypothesis that is consistent with the observed evidence is exponentially preferred with each additional piece of evidence, all else equal (Tenenbaum & Griffiths, 2001). Mathematically, this is the case because the hypothesis corresponding to the smallest set of possible evidence receives greater relative likelihood for a given piece of observed evidence than its competitors. Intuitively, the size principle captures the sense of a ‘suspicious coincidence’ (cf. Lewis & Frank, 2018; Navarro & Perfors, 2010; Perfors et al., 2011).

Concretely, if all observed violations apply to only one group, how much the evidence supports a parochial inference will depend on the relative size of the group to which the violations apply. When all violations come from a small minority group, the size principle dictates that a parochial norm is much more likely than when sample violations come from a large majority group.

Indeed, one way to think about the sort of statistical sensitivity that is described by the size principle is in terms of being sensitive to the population from which the available evidence is drawn. In the example in Fig. 1, the learner assesses the proportion of Hibbles to Glerks, and then

assesses whether the examples are drawn from the minority group or the majority group. A number of studies have shown that children can make appropriately population-sensitive inferences in complex social learning tasks, ranging from learning about word meanings to learning about psychological causes such as an agent's preferences (see, e.g., Xu & Tenenbaum, 2007; Kushnir, Xu, & Wellman, 2010; Gweon and Schulz, 2011; Lucas et al., 2014; Wellman, Kushnir, Xu and Brink, 2016; Diesendruck, Salzer, Kushnir, & Xu, 2015; Vélez, Wu, & Gweon, 2018; Riggs, 2019; Riggs & Long, 2020). Especially relevant are findings from Heck, Kushnir, and Kinzler (2021); Heck, Bas, and Kinzler (2022) who in recent work show that children's learning about power in social hierarchies can be explained in terms of rational, population-sensitive statistical inferences.

As an example of population-sensitive learning, take Kushnir et al., 2010. These studies showed that when an agent selects a blue flower from a box containing some blue flowers and mostly red circles, 3- and 4-year-old children could use this evidence to infer the agent had a preference for blue flowers. Riggs (2019) used a similar logic to demonstrate that children's inferences about social scope can be sensitive to assumed population features. In these studies, children ages 3- to 7-years-old observed characters from populations of varying geographical scope (e.g., from the same or different city, state, or country) perform the same or different actions (e.g., kneeling vs. a push-up). The key finding was that children made statistically appropriate inferences about whether a norm that applied to a larger population (e.g., characters from the same country) also applied to a sub-population (e.g., characters from a city within that country). A natural interpretation of this result is that children's inferences about the social scope of norms are appropriately sensitive to assumed population sizes. In the current research, we provide a more direct test of this idea by explicitly manipulating population size. Furthermore, we also provide a formal definition of population-sensitivity (Section 4).

What, if anything, can prior studies tell us about children's population-sensitivity in norm learning? On the whole, inferences in the novel groups paradigm (e.g. Roberts et al., 2017) may be statistically appropriate (see Partington, Nichols, & Kushnir, 2021 for a meta-analysis). However, since the extant work does not directly manipulate the population-features of the evidence available, whether learners can acquire distinctions of social scope in a population-sensitive manner has yet to be subjected to severe testing (Mayo, 2018) and thus remains an open question. The central aim of the present research is addressing this question.

1.4. The present research

In the present research,¹ we examine participants' rule learning in a scenario involving two novel groups who live together on an island. As with previous research (cf. Roberts et al., 2017), using a novel groups paradigm allows us to examine participants' learning under plausibly minimal extraneous influences: what matters for learning are the learner's rule representations, social group representations, and the available evidence in the form of instances of violations. Using a novel rule learning paradigm (cf., Nichols, Kumar, Lopez, Ayars, & Chan, 2016), we hold behavioral regularity constant across the novel groups while varying the size of the target group from which sample rule violations are exclusively observed. The participant's task is inferring whether the novel rule is parochial (i.e., applies only to the target group) or inclusive (i.e., applies to the target group and the non-sampled group), based on the available evidence.

With these elements in place, we can examine whether adults and children can use the available statistical evidence of social scope in a population-sensitive manner. If the rational learning account is correct, then more participants should infer the rule is parochial when sample violations are observed in smaller sub-populations. This prediction is due to the size principle: all else equal, the smaller the hypothesis consistent with the evidence, the more it should be favored by naïve, rational learners. Moreover, this prediction is unique to the rational learning account. Since group behavioral regularities are held constant throughout, an automatic group bias cannot account for differences in the rate of parochial norm acquisition.

2. Study 1

2.1. Method

2.1.1. Participants

Adult participants ($n = 480$; 32% female; $m_{\text{Age}} = 35.5$ years, $sd = 10.6$) were recruited from Amazon MTurk to complete a survey for modest compensation. An additional 77 participants were excluded from analyses for failing to complete the survey. All participants included in analyses completed the entire survey and passed a series of comprehension checks about the novel groups.

2.1.2. Procedure

We presented participants with a scenario in which two groups (labelled "Hibbles" and "Glerks") live together on an island (in all but the 100% condition, see below). Participants were randomly assigned to one of four conditions. Across conditions, the size of the target group relative to the island's total population was varied; it was either 20%, 50%, 80%, or 100% of approximately 100 total inhabitants. In all conditions, a fixed proportion of each group (approximately 35%) was shown wearing a morally neutral item of clothing (e.g., Trial 1: ribbons, Trial 2: hats). To provide a concrete example: for Trial 1, participants in the 20% condition would view an island inhabited by 20 Glerks and 80 Hibbles (100 individuals in total), with 7 of the Glerks wearing a ribbon and 28 of the Hibbles wearing a ribbon (35% of the individuals in each group). See Fig. 2 for a visual reference.

Next, participants were told the island has rules and that their task would be to figure out one of the island's rules. Participants then watched a video highlighting a sample of 4 members of the target group (e.g., "Hibbles") wearing the clothing item (e.g., a ribbon) as violations ("This is against the rule."). Afterwards, participants were asked to infer if other individuals on the island were also violating the rule. The initial set of dependent measures were all 'rule judgments': for each individual, participants made a forced-choice judgment of whether "This is against the rule" or "This is not against the rule." We solicited these rule

judgments about all possible group member/clothing item combinations (order counterbalanced): another target group member with a ribbon, a target group member without a ribbon, a non-sampled group member with a ribbon, and a non-sampled group member without a ribbon. In a fifth and final rule judgment, participants made the same judgment about a visitor to the island who was wearing a ribbon. The visitor was called a "Zorg" and its body was purple and spiky.

The second dependent measure was an open response item in which participants articulated their understanding of the rule ("What is the rule?"). Participants also provided confidence ratings of this understanding (a 7-point scale, "How confident are you that you know the rule?" with 1 = Not at all, 7 = Very).

Participants then repeated the entire procedure in Trial 2, which was identical to Trial 1 except for the individuals wearing a different clothing item (e.g., hats). Thus, in Trial 2, four sample violations were drawn exclusively from the same target group as Trial 1, and participants responded to the corresponding dependent measures (rule judgments, open response, confidence rating) once again.

2.1.3. Coding

We scored participants' judgments as "This is against the rule" = 1 and "This is not against the rule" = 0. Since the evidence provided to participants underdetermines the content of the rule, we decide to score all items in the same manner. For example, say that each sample violation involves a Hibble wearing a hat. The rule may apply to wearing hats, it may apply to Hibbles, or it may apply to Hibbles wearing hats. Given this evidence, participants may also infer that it is against the rule for Glerks to not wear hats. To distinguish between the various sorts of rules that participants could acquire, we need to score each possibility in a consistent way.

For the open response question, responses were coded for whether participants articulated a parochial rule (e.g., "Glerks cannot wear hats") or an inclusive rule (e.g., "No hats allowed"). Responses that did not articulate a rule (e.g., "I just guessed") were not counted. Responses that articulated a rule were scored as either 'Parochial' = 0 or 'Inclusive' = 1 by a coder who did not know from which condition the responses originated. Inter-rater agreement for the open response coding was high (percent agreement: 95% (229/240), Cohen's $\kappa = 0.93$) and remaining disagreements were resolved in discussion between raters.

2.2. Results

2.2.1. Rule judgments

For each individual, we ran a logistic regression model with *rule judgment score* as the dependent variable and *condition* (dummy coded, with "20%" as reference group), *item* (dummy coded with "hats" as reference group), and *trial* (1 or 2) as the independent variable. Participants' rule judgments varied as a function of the relative population proportions when applying the rule to individuals from the non-sampled group ($\beta_{\text{condition}} = 0.640$, $SE = 0.070$, $p < .001$) and the visitor ($\beta_{\text{condition}} = 0.427$, $SE = 0.067$, $p < .001$). However, participants did not apply the rule to the target group at significantly different rates across conditions ($\beta_{\text{condition}} = -0.173$, $SE = 0.101$, $p = .09$).

Post-hoc analyses confirmed this overall pattern of judgments was consistent predictions of the rational learning account (see Fig. 3). In each condition, participants applied the rule most frequently to members of the target group ($m_{\text{Target}} = 0.90, 0.87, 0.92, 0.83$ in the 20%, 50%, 80%, and 100% conditions, respectively), whereas participants were most likely to judge the rule applied narrowly when sample violations came from a 20% minority ($m_{\text{Non-sampled}} = 0.28$, $m_{\text{Visitor}} = 0.42$), followed by the 50% condition ($m_{\text{Non-sampled}} = 0.39$, $m_{\text{Visitor}} = 0.48$) followed by the 80% condition ($m_{\text{Non-sampled}} = 0.48$, $m_{\text{Visitor}} = 0.53$). Finally, when the population contained only one group (the 100% condition), participants most frequently applied the rule to all novel individuals ($m_{\text{Non-sampled}} = 0.71$, $m_{\text{Visitor}} = 0.72$).

We further analyzed participants' judgments using a trial-level

¹ See OSF




Fig. 2. Example of novel groups and novel rule learning paradigm.




Fig. 3. Mean rule judgment score for each candidate group member (Target = blue, Non-sampled = yellow, Visitor = purple), error bars represent standard error. Higher bars indicate that participants more frequently judged the rule applied to the candidate individual. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

scoring system that considers the totality of the pattern of response across all three questions on each trial. Trials where the participant applied the rule to the target individual and not the other individual received a *trial score* = 1 (or “Parochial”). Trials where the participant applied the rule to the target individual and either the non-sampled individual or the visitor received a *trial score* = 2 (“More inclusive”). Trials where the participant applied the rule to all individuals received a *trial score* = 3 (“Most inclusive”). We excluded from analysis any responses that did not fit this pattern (and thus was uninterpretable, 115/960 or 12% of all trials). Table 1 shows the numbers and percentages of responses that fell into each category in each condition.

We analyzed these data using an ordinal regression model with *trial score* as the dependent variable and *condition* (dummy coded, with “20%” as the reference group), *item* (dummy coded, with “hats” as the reference group) and *trial* (1 or 2) as independent variables. We found that responses in the 20% condition were most parochial compared to the 50% condition ($b = 0.406, SE = 0.18, p = .002$), the 80% condition ($b = 0.611, SE = 0.18, p < .001$), and the 100% condition ($b = 1.82, SE$

Table 1

Summary statistics for rule judgments across the three individuals (Target, Non-sampled, Visitor) in Experiment 1.

Trial score	Condition			
	20%	50%	80%	100%
<i>Parochial</i> (Rule applies to Target & no other individual)	119/222 (54%)	96/211 (45%)	93/217 (43%)	37/195 (19%)
<i>More inclusive</i> (Rule applies to Target & 1 other individual)	49/222 (22%)	38/211 (18%)	29/217 (13%)	19/195 (10%)
<i>Most inclusive</i> (Rule applies to all individuals)	54/222 (24%)	77/211 (37%)	95/217 (44%)	139/195 (71%)

= 0.20, $p < .001$).

2.2.2. Open response scores

Next, we ran a logistic regression model with Open Response score as the dependent variable and *condition* (dummy coded, with “20%” as reference group) as the independent variable. As with rule judgments, open response scores also varied as a function of condition ($\beta_{condition} = 0.726, SE = 0.077, p < .001$). Once again, post-hoc analyses show participants most frequently articulated a parochial rule in the 20% condition ($m = 0.30$) and the 50% condition ($m = 0.33$), followed by the 80% condition ($m = 0.50$), and lastly the 100% condition ($m = 0.78$). Thus, participants' articulation of the rule provides further evidence that their social scope learning was population-sensitive.

2.2.3. Confidence ratings

Confidence ratings differed significantly between the 20% condition ($m = 4.38, SD = 1.74$) and the 100% condition ($m = 5.06, SD = 1.66$) (20% vs. 100%: $t(239) = -3.1, p = .006, d = 0.40$). There was no significant difference between ratings in the 20% condition, the 50% condition ($m = 4.77, SD = 1.51$), and the 80% condition ($m = 4.71, SD = 1.55$), nor was there a significant difference in ratings between the 50%, 80%, and 100% conditions.

2.3. Discussion

Study 1 provides evidence that adults can learn the social scope of novel norms in a sample-sensitive and population-sensitive appropriate manner. As predicted by the rational learning account, the absolute rate of parochial judgments varied as a function of the group size from which sample violations were observed. Parochial judgments were most frequent when sample violations were drawn from a 20% target group and decreased in frequency as the size of the target group increased. Participants' open responses provide further support for this trend: participants were most likely to articulate a parochial rule when sample violations were drawn from the 20% or 50% target groups, whereas participants were more likely to articulate an inclusive rule when sample violations were drawn from the 80% and 100% target groups. Altogether, evidence from Study 1 provides support for the rational learning account, at least as it applies to adults' norm learning.

3. Study 2

In Study 2, we examined whether children's norm learning also conforms to the predictions of the rational learning account. Much of the extant research on children's norm learning has focused on the inferences that children make from generic evidence (e.g., evidence like "Hibbles eat purple berries"). It is worth reiterating that a key difference between this work and the present research is that we provide children with explicitly labelled rule violations (i.e., "This is against the rule"). Research using generic evidence has typically found that as age increases children are less prone to making parochial inferences about the social scope of a novel norm (see, e.g., Roberts et al., 2011; see Partington et al., 2021, p. 2756 for a meta-analysis). It is an empirical question whether or not the same pattern also holds with explicitly labelled rule violations, the focus of the present research. Next, we seek evidence to address this question by examining norm learning in children ages 5- to 8-years-old.

3.1. Method

3.1.1. Participants

Children ages 5- to 8-years-old ($n = 120$; $m_{\text{age}} = 6.96$ years, $sd = 1.16$ years) were recruited to participate in a study conducted via Zoom online video calling. Prior work shows effects of sampling on inferences about object preferences (Garvin & Woodward, 2015; Kushnir et al., 2010), friendship preferences (Eason, Doctor, Chang, Kushnir, & Somerville, 2018; Heck et al., 2021), and group preferences (Diesendruck et al., 2015) in children as young as 3.5-years-old. Thus, we expected to find continuity across ages in our sample. Sample size was determined according to a priori power analyses by setting the smallest effect size of interest (Lakens, 2022) according to the observed result across the 20% and 80% conditions in Study 1 (see OSF for an R script of this power analysis). According to lab protocol, we aimed to balance this sample across ages by collecting $n = 30$ children from each of the following age groups: 5-year-olds, 6-year-olds, 7-year-olds, 8-year-olds. Children received \$5 e-gift cards for compensation. An additional 8 children were excluded from the analyses due to distractions or technology-related issues during testing. All children included in the analyses completed the entire survey and correctly identified the two novel groups in a brief series of comprehension checks.

3.1.2. Procedure

We presented children with an identical scenario as Study 1, where children were randomly assigned to one of two conditions. In the 20% condition, four rule violations were sampled from a 20% minority. In the 80% condition, four rule violations were sampled from an 80% majority. In Trial 1 children watched a video that highlighted four members of the target group, all wearing a clothing item, and were told "This is against the rule" (examples of these sampling videos can be viewed on OSF). The

members of the target group were highlighted in sequential order. After each Target individual was highlighted, the experimenter paused the video, moved their cursor to the highlighted individual, and said "This is against the rule." Afterwards, we asked children whether other individuals on the island were also breaking the rule (order counter-balanced): another target group member with a ribbon, a target group member without a ribbon, a non-sampled group member with a ribbon, and a non-sampled group member without a ribbon. For each individual, children were asked the following, "This [Hibble/Glerk] is wearing a [clothing item]. Is this against the rule?" and made a forced-choice judgment of "This is against the rule" or "This is not against the rule." Next, children made the same judgment about a visitor (named 'Zorg') who comes to the island and wears the clothing item.

Following this, children were asked to articulate their understanding of the rule in an open response item ("What is the rule?"). If a child was reluctant to give a response, the experimenter prompted, "Do you have any guesses about the rule? Who do you think is allowed to do what?" Lastly, children provided confidence ratings of this understanding of the rule (a 7-point scale, "How sure are you that you know the rule?" with 1 = Not at all, 7 = Very, scale endpoints and mid-point labelled with corresponding emoji expressions). To aid understanding, the experimenter moved the cursor to the end-points and mid-points on the scale and said out loud what each meant (i.e., "Are you not at all sure?", "Are you very sure?", and "Are you somewhere in between?"). Children then repeated the entire procedure in Trial 2, which was identical to Trial 1 except for the individuals were wearing a different clothing item.

3.1.3. Coding

Rule judgments and open responses were coded in an identical manner as Study 1. For ease of reference, we will repeat the coding criteria here. Rule judgments were scored "This is not against the rule" = 0 and "This is against the rule" = 1. Open responses that articulated an interpretable rule (198/240, 82.5% in total) were scored (No = 0, Yes = 1) for whether the response articulated an 'inclusive' (e.g., "No hats allowed"), 'parochial' (e.g., "Glerks cannot wear hats"), 'extreme parochial' (e.g., "Only certain Glerks cannot wear hats"), or 'inverse' rule (e.g., "Glerks cannot wear hats, but Hibbles have to wear hats"). Once again, inter-rater agreement for the open response coding was high (percent agreement: 97% (63/65), Cohen's $\kappa = 0.96$) and remaining disagreements were resolved in discussion between raters.

3.2. Results

3.2.1. Rule judgments

We ran a logistic regression model with rule judgment as the dependent variable and *condition* (dummy coded: 20% = 0, 80% = 1), *individual* (dummy coded, with "Target" as reference group), *participant age*, *item* (dummy coded with "hats" as reference group), and *trial* (1 or 2) as independent variables.² There were significant main effects of *individual* for each comparison to the Target individual, meaning children applied the rule to the other individuals less frequently than they applied the rule to the Target individual ($\beta_{\text{Novel}} = -0.38$, $SE = 0.10$, $p < .001$; $\beta_{\text{NS}} = -0.57$, $SE = 0.10$, $p < .001$). There were not a significant main effect of *condition* or significant two-way interaction effects between *condition* x *individual*. However, there was a significant three-way interaction effects between *condition* x *individual* x *age* for the Novel individual, but not for the Non-sampled individual ($\beta_{\text{Novel} \times 80 \times \text{age}} = 0.20$, $se = 0.10$, $p = .046$; $\beta_{\text{NS} \times 80 \times \text{age}} = 0.18$, $se = 0.10$, $p = .075$).

To examine this interaction further, we split the kids into two age groups (5- & 6-year olds, 7- and 8-year olds) and used the same trial-level scoring system as in Study 1. Table 2 shows the number of

² Attentive readers will note that the models in Study 2 test different statistical hypotheses than the model in Study 1. This is because developmental studies conventionally include age as a covariate in omnibus analyses.

Table 2

Summary statistics for rule judgments across the three individuals (Target, Non-sampled, Visitor) in Experiment 2.

Condition		Age group	
		5- & 6-year-olds	7- & 8-year-olds
20%	<i>Parochial</i> (Rule applies to Target & no other individual)	4/47 (8%)	17/51 (33%)
	<i>More inclusive</i> (Rule applies to Target & 1 other individual)	14/47 (30%)	22/51 (43%)
	<i>Most inclusive</i> (Rule applies to all individuals)	29/47 (62%)	12/51 (24%)
80%	<i>Parochial</i> (Rule applies to Target & no other individual)	7/45 (15%)	8/45 (18%)
	<i>More inclusive</i> (Rule applies to Target & 1 other individual)	12/45 (27%)	12/45 (27%)
	<i>Most inclusive</i> (Rule applies to all individuals)	26/45 (58%)	25/45 (55%)

responses in each category by age group and condition. Using this coding scheme, 22% (52/240) of trials were excluded from analysis for not applying the rule to the target individual. As before, we analyzed these data using an ordinal regression model with trial score as the dependent variable and condition (dummy coded, with “20%” as the reference group), age group (dummy coded, with “5- & 6-year-olds” as the reference group), item (dummy coded, with “hats” as the reference group) and trial (1 or 2) as independent variables. We found a *condition x age group* interaction such that older children were likely to infer a parochial rule in 20% condition ($b_{\text{Condition}=80 \times \text{Age group}=7 \ \& \ 8} = 1.410$, $SE = 0.58$, $p = .01$). In addition, overall we found that older children were more likely to exhibit “Parochial” trials than younger children ($b_{\text{Age group}} = -1.563$, $SE = 0.40$, $p < .001$). This pattern of results indicates that younger children’s judgments did not display population-sensitivity in part because the younger children tended to infer inclusive rules in even the 20% condition.

3.2.2. Open responses

Next, we ran a logistic regression model with Open Response score as the dependent variable and *condition* (coded: 20% = 0, 80% = 1), *rule* (dummy coded, with “inclusive” as reference group), and *age* (z-scored) as independent variables. There were significant main effects of *rule* type for ‘parochial,’ ‘extreme parochial’ and ‘inverse’ rules ($\beta_{\text{Rule=Parochial}} = -0.33$, $SE = 0.09$, $p < .001$; $\beta_{\text{Rule=Extreme parochial}} = -1.32$, $SE = 0.16$, $p < .001$; $\beta_{\text{Rule=Inverse}} = -1.96$, $SE = 0.39$, $p < .001$), meaning, overall children articulated less ‘parochial’ rules (75/198, 38% of responses), ‘extreme parochial’ rules (10/198, 5%) and less ‘inverse’ rules (4/198, 2%) than they articulated ‘inclusive’ rules (109/198, 55% of responses). There was a significant interaction effect of *rule = Parochial x condition = 80* ($\beta_{\text{rule=parochial} \times \text{condition}=80} = -0.26$, $SE = 0.09$, $p = .007$) such that compared to ‘inclusive’ rules (80% = 61/98, 62% of responses, 20% = 48/100, 48%), children articulated less ‘parochial’ rules in the 80% condition (30/98, 31%) than in the 20% condition (45/100, 45%). In addition, there was a significant interaction effect of *rule = Parochial x age* ($\beta_{\text{rule=parochial} \times \text{age}} = 0.51$, $SE = 0.10$, $p < .001$) meaning, as age increased children articulated more ‘parochial’ rules.

In summary: children of all ages who articulated an interpretable rule articulated either a ‘parochial’ rule or an ‘inclusive’ rule. Further, in examining the pattern of open response scores across conditions, we found evidence that children’s responses were sensitive to the population features of their social scope evidence, and this sensitivity increased as age increased.

3.2.3. Confidence ratings

Children’s confidence ratings were high on average ($m = 5.40$, $sd = 1.26$), and there was a near-zero association between confidence rating and age (Pearson’s $r = -0.03$, $p = .546$). There was not a significant difference in confidence ratings across conditions (Welch’s $t(237.7) = 0.63$, $p = .53$; $m_{20\%} = 5.48$, $m_{80\%} = 5.34$). We do not discuss the confidence ratings further.

3.2.4. Comparing children and adults: Rule judgments

To compare adults’ and children’s rule judgments, we ran a logistic regression model with rule judgment score as the dependent variable and *condition* (dummy coded: 20% = 0, 80% = 1), *individual* (dummy coded, with “Target” as the reference group), and participant *age group* (dummy coded: adults = 0, children = 1) as dependent variables. There were significant interaction effects of *Individual = NS x age group* ($\beta = -0.38$, $SE = 0.06$, $p < .001$) and *Individual = Novel x age group* ($\beta = -0.37$, $SE = 0.06$, $p < .001$), meaning, compared to the Target individual, children applied the rule more often to the non-sampled and novel individuals than adults did (see Fig. 4).

3.2.5. Comparing children and adults: Open response

To compare adults’ and children’s open responses, we ran a logistic regression model with open response score as the dependent variable and *condition* (dummy coded: 20% = 0, 80% = 1), *rule* (dummy coded, with “inclusive” as reference group), and participant *age group* (dummy coded: adults = 0, children = 1) as dependent variables. There was a significant interaction effect of *Rule = parochial x age group* ($\beta = -0.28$, $se = 0.08$, $p < .001$), meaning adults articulated parochial rules more often than children (see Fig. 5).

3.3. Discussion

We found age-related change in children’s population-sensitivity to social scope evidence. Younger children’s rule judgments did not show evidence of population-sensitivity. However, those younger children who provided interpretable open responses did so in a population-sensitive manner. Older children showed a similar pattern to adults – they were more likely to infer narrow-scope (parochial) rules in the 20% condition than in the 80% condition.

Our findings suggest that norm learning from statistical evidence may be different from other types of statistical social learning. In particular, in prior work on statistical learning of preferences, there were no age-related differences from children as young as 3.5 years-old (see, e.g., Kushnir et al., 2010; Heck et al., 2021, year; Eason et al., 2018; Diesendruck et al., 2015). Not finding age-related continuity may indicate an important way in which norm acquisition works differently from inferring preferences. Indeed, in studies where children are asked to generalize statistical evidence from one person to another, they readily generalize conventions, labels, and norms to other group members, in contrast with preferences, which they view as idiosyncratic to specific individuals (Buresh & Woodward, 2007; Moll & Tomasello, 2007).

However, it may still be the case that children are more likely to treat some norms (e.g., moral norms) as more universal than others (e.g., conventional norms, religious norms) (see, e.g., Turiel, 2006; Smetana, 2013; Schmidt, Rakoczy, & Tomasello, 2012; Srinivasan, Kaplan, & Dahl, 2019). It may also be the case that children require more information about the groups themselves to view them as distinctive (and their norms as tribal). Indeed, in our studies, we did not give children any information about the distinctiveness of the groups in any way other than their physical appearances and labels.

Given this sort of novel situation, it appears that especially younger children are inclined to infer that novel rules apply inclusively across group boundaries. Furthermore, this pattern of results is consistent with our findings with the adult participants, who also exhibited an inclusive tendency, albeit less so than children. In order to more fully investigate this possibility, we compare our experimental results against the




Fig. 4. Bar chart of average rule judgment scores for 5- and 6-year-olds (left panel; $n = 60$), 7- and 8-year-olds (center panel; $n = 60$), and adults (right panel; $n = 240$). Color corresponds to the candidate subject of the rule, and error bars correspond to 95% CIs.

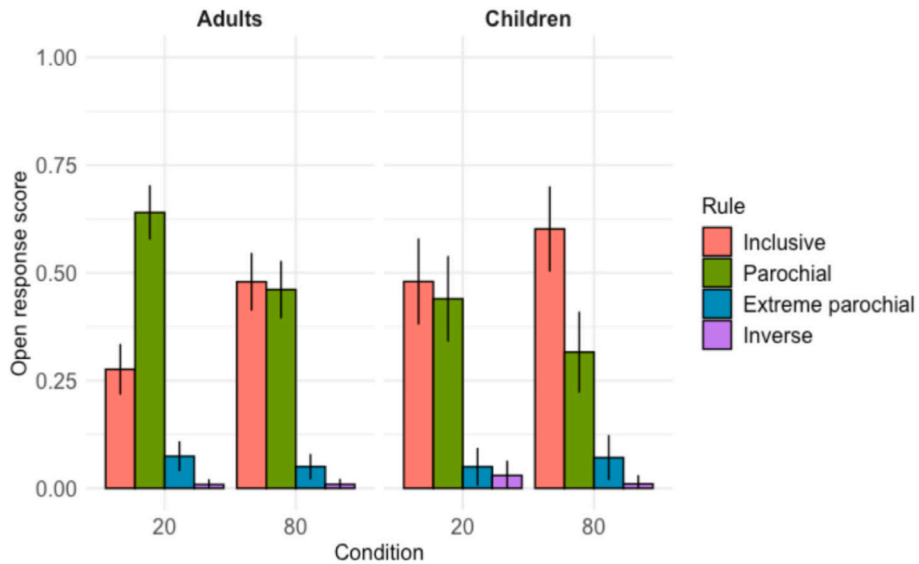


Fig. 5. Bar chart of average open response scores for adults (left panel; $n = 240$), and children (right panel; $n = 120$). Color corresponds to rule type (red = ‘Inclusive’, green = ‘Parochial’, turquoise = ‘Extreme parochial’ and purple = ‘Inverse’). Error bars correspond to 95% CIs. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

predictions from an unbiased Bayesian model to make precise and assess the extent of children and adults’ disposition to infer inclusivity.

4. Computational analysis

Formally, norm acquisition involves inferring a rule R from a certain number of example violations D , where $D = d_1, \dots, d_n$. Per the rational learning account, the learner has access to a hypothesis space H that contains a set of candidate hypotheses for representing the rule R . The learner uses Bayes Rule to compute the posterior probability $p(h|D)$ for each candidate hypothesis:

$$p(h|D) = \frac{p(D|h)p(h)}{\sum_{h' \in H} p(D|h')p(h')}$$

To capture the size principle in formal terms, we specify the likelihood function as follows (per Xu & Tenenbaum, 2007), where n is the number of sample violations observed:

$$p(D|h) = \left[\frac{1}{\text{size}(h)} \right]^n$$

According to the rational learning account, the learner’s hypothesis space H contains only two hypotheses: $h_{\text{Parochial}}$, the rule applies to only members of the group from which violations were observed, and $h_{\text{Inclusive}}$, the rule applies to the entire population. Second, $h_{\text{Parochial}}$ is a subset of $h_{\text{Inclusive}}$: all possible subjects of the parochial rule are also possible subjects of the inclusive rule, but not vice-versa. This allows us to characterize the size of each hypothesis:

$$\begin{aligned} \text{size}(h_{\text{Inclusive}}) &= 1 - \sigma \\ \text{size}(h_{\text{Parochial}}) &= \beta(1 - \sigma) \end{aligned}$$

where σ is a small constant, keeping the respective likelihoods³ from equaling 1, and β is a value that satisfies the conditions $0 < \beta < 1$, keeping the size of $h_{\text{parochial}}$ smaller than $h_{\text{inclusive}}$. Thus, by varying values of β , we can model the learner's credence in $h_{\text{parochial}}$ when n sample violations are observed in a sub-population of approximately β size. For example, if sample violations only apply to members of a 20% sub-group of the total population, then $\beta = 0.20$, whereas if sample violations only apply to members of an 80% sub-group of the total population, then $\beta = 0.80$.

The final step is specifying the learner's prior for each hypothesis. Since there are only two candidate hypotheses, we can define each prior: $p(h_{\text{inclusive}}) = k$ and $p(h_{\text{parochial}}) = r$, where $k + r = 1$ (satisfying the probability axioms). With this machinery in place, we can model the rational learner's posterior credence in each hypothesis, given n violations from a sub-population of β size and flat priors ($k = r$).

Fig. 6a shows the posterior probability for each hypothesis (y-axis) given $n = 4$ violations from a β -sized population (x-axis). As shown, the model predicts an unbiased, ideal learner will display *population-sensitivity* (i.e., $p(h_{\text{inclusive}} | D)$ increases as a function of β). Fig. 6b shows the predicted rate of applying the rule to other members of the target group (from which samples were drawn) and to members of the non-sampled group (i.e., the sub-population of β size) for β parameter values that correspond to our experimental conditions in Study 1. Here, we can also see the model expects that an unbiased learner will be *sample-sensitive* (i.e., for all values of β , the model assigns higher probability that the rule applies to a member of the Target group than a member of non-sampled group). Lastly, since the size principle dictates the smallest hypothesis consistent with the evidence is always preferred, the posterior for $h_{\text{parochial}}$ is always greater than the posterior for $h_{\text{inclusive}}$. Thus, the unbiased model predicts an overall tendency to favor parochial rules.

A key question of interest is how the unbiased model's predictions quantitatively compare to the experimental results (see Fig. 7). Exact binomial tests show that adults more frequently made inclusive judgments than expected from unbiased Bayesian learners (20%: predicted = 0.001, observed = 0.28; 50%: predicted = 0.06, observed = 0.39; 80%: predicted = 0.29, observed = 0.48; 100%: predicted = 0.50, observed = 0.70; all $p < .001$). Likewise, children also made inclusive judgments at a greater rate than unbiased Bayesian learners (20%: predicted = 0.001, observed = 0.52; 80%: predicted = 0.29, observed = 0.55; both $p < .001$). Overall, children made inclusive judgments at a greater rate than adults in the 20% condition ($p < .001$), but not in the 80% condition ($p = .15$). The difference in the 20% condition is mainly attributable to the younger children alone, as the older children's pattern of response does not display a significant difference with the adults' (20%: 7- and 8-year-olds = 0.38, adults = 0.28, $p = .11$; 5- and 6-year-olds = 0.66, adults = 0.28, $p < .001$). The age-related effect here suggests that an inclusive prior emerges early in childhood, and perhaps even weighs more heavily in children's judgments than in adults'.

Our analysis shows that both children and adults conform to a Bayesian inference procedure with a *biased prior toward inclusivity*. Indeed, in one sense, this result confirms a straightforward a priori consequence of the learning problem. Since $h_{\text{parochial}}$ will always be the smallest hypothesis consistent with the data, a learner who follows the size principle will need to have a prior that favors $h_{\text{inclusive}}$ in order to be able to infer $h_{\text{inclusive}}$ when the evidence applies to only one group. In general, as Xu and Tenenbaum (2007) point out, if all logically possible hypotheses are given equal prior probability, the size principle would entail that the best hypothesis is the one that contains only items in the sample. But this would not provide an adequate solution to the learning problem—the best hypothesis would be “a hypothesis that calls for no generalization at all!” (Xu & Tenenbaum, 2007, p. 252). Here, an

³ If $\text{size}(h) = 1$, then the likelihood function would not provide exponentially more likelihood with each additional piece of evidence (because $1^2 = 1^3 = 1^4 \dots$).

inclusive bias is what provides the basis for generalizing from the samples that participants received in our studies.

The exact nature of the bias toward inclusivity is an open empirical question. One possibility is that the bias may be the product of an overhypotheses about social norms, namely that norms apply broadly unless one learns otherwise (c.f., Kemp, Perfors, & Tenenbaum, 2007). Alternatively, the bias could be the product of a post-hoc updating rule whereby $h_{\text{inclusive}}$ receives a ‘boost’ if it is consistent with the available evidence (akin to explanationist updating and the bias for simplicity, see Douven & Schubach, 2015; Lombrozo, 2007). Further, as an empirical matter, the current results cannot tell us whether the bias persists across experimental manipulations and paradigms. More evidence of this sort is needed to determine when and why human learners display an inclusive bias or how the learning bias operates in human psychology.

5. General discussion

The widespread prevalence of parochial norms across history and cultures have led some to suggest parochialism is itself a human universal (Clark et al., 2019; Greene, 2013) in part owing to evolved, group-based biases in social norm acquisition (Chalik & Rhodes, 2020; Chudek & Henrich, 2011; Roberts et al., 2017). In this paper, we investigated whether a rational learning process can also explain this phenomenon. In Study 1, we found that adults can acquire distinctions of social scope in a statistically appropriate manner, and this finding was robust across two forms of measurement (rule judgments and open response). In Study 2, older children displayed the adult-like statistical sensitivity in their rule judgments, and even younger children did so in their open responses. Computational analyses suggests that rule judgments were inclusively biased: compared to an unbiased Bayesian learner, children tended to assume that novel rules apply to everyone in a candidate population. Adults also displayed an inclusive bias, albeit to a lesser extent than children.

Broadly, these findings suggest that rational learning processes can indeed explain the acquisition of parochial norms and highlight an important sense in which children's norm learning can be biased in the *opposite* direction of tribalism. At the least, the finding that children and adults are inclusively biased serves as an existence proof that deep-rooted tribal biases in social learning are not necessary to explain the acquisition of parochial norms. Rather, if children and adults are rational learners, they can acquire a parochial norm when presented with evidence that is consistent with parochialism. However, tribalism can still play a role in norm acquisition, for example, by influencing the sort of evidence that adults seek out, or the evidence to which children are exposed.

Even so, the present demonstration of adults' and older children's statistical sensitivities shows that group-based biases are not necessary for the acquisition of parochial norms. As such, our results stand in contrast to prevailing accounts of norm acquisition, all of which emphasize the role of group-based biases in the acquisition of parochial norms (Chalik & Rhodes, 2020; Chudek & Henrich, 2011; Roberts et al., 2017). In particular, our results cut against the notion that mere group representations automatically elicit biases in norm learning such that one assumes the social scope of norms directly maps onto group behavioral regularities (Roberts et al., 2017, p. 593; Chalik & Rhodes, 2020, p. 80). Rather, we have shown that human learners can relate evidence of social scope with group representations in a statistically appropriate manner, perhaps even with an overall inclusive tendency.

How can these findings be reconciled with past work suggesting a central role for the automatic group bias? One possibility is that certain forms of evidence do indeed elicit biases toward acquiring parochial norms. For example, our paradigm presents learners with population-level evidence, whereas Roberts and colleagues' paradigm places a much greater emphasis on individual-level exemplars. In addition, Roberts and colleagues present participants with generic linguistic evidence (“Hibbles eat purple berries”) whereas our linguistic evidence




Fig. 6. a. (left) Posterior probability for the parochial (red) and inclusive (blue) hypotheses (y-axis), given $n = 4$ violations from a β -sized population (x-axis). b. (right) the predicted mean rate of applying the rule to the other members of the target group and to members of the non-sampled group, for β parameter values that correspond to our experimental conditions in Study 1.

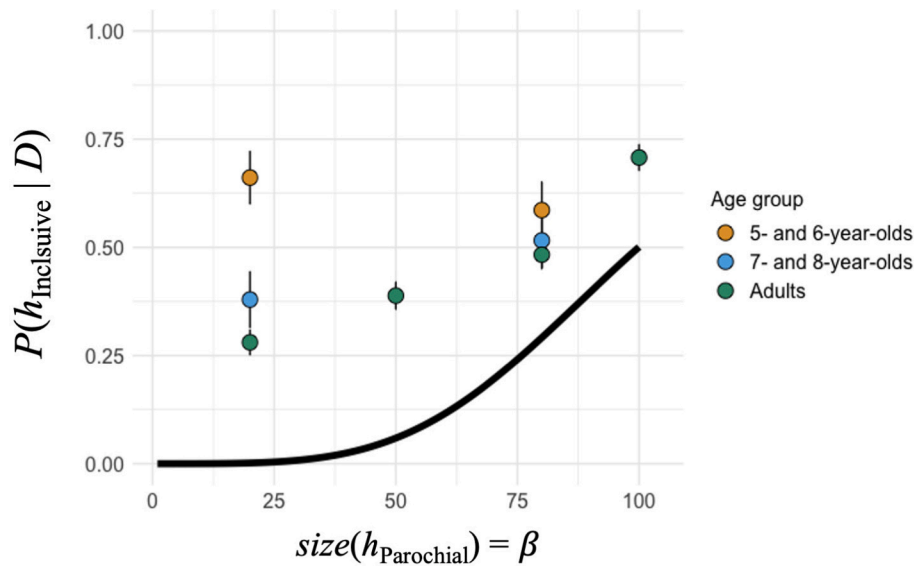


Fig. 7. Black line represents the unbiased Bayesian model predicted posterior for $h_{Inclusive}$ (y-axis), given $n = 4$ violations from a sub-population of β size (x-axis) and $\sigma = 0.001$. Circles indicate mean judgment score for each age group (Adults = blue, 7- and 8-year-olds = green, 5-and-6-year-olds = red), error bars represent standard error. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

picks out explicitly labelled instances of norm violations (“This is against the rule”). It is possible that these differences prompt participants to approach broadly similar learning problems in different ways.

Another possibility is that learners are always making roughly the right sort of statistical inferences and the extant literature has used an inappropriate standard for assessing bias (cf. Pisor & Ross, 2021). So far, the extant literature has assessed bias by comparing participant responses against random chance (a ‘random standard’). However, as revealed from the Bayesian analysis, if naïve learners receive any evidence consistent with a parochial rule, then they *should* judge a parochial norm to be more likely than an inclusive norm. Given this, a more appropriate standard for assessing bias would be an *evidential standard* whereby participant responses are compared against what is most probable, given the evidence. Using an evidential standard, Partington et al. (2021) conducted a meta-analysis of studies that examined children’s norm learning in novel contexts ($N = 1369$ children ages 4- to 13-years-old), and the results suggested that children’s norm learning is indeed appropriately sensitive to the available evidence. At least in the case of children’s learning about novel, third-party norms, the current body of findings are entirely consistent with rational learning processes playing a central role.

To be clear, the present findings do not rule out potential group-based biases in real-world norm acquisition, or even norm acquisition under minimal conditions of group membership. Indeed, our findings are largely silent on matters beyond the sort of novel groups and novel rule learning paradigms characteristic of recent research in cognitive development and moral learning. This is a clear limitation on generalizing from the present findings to the role of rational processes in real-world norm acquisition. Currently, we do not have license to generalize these findings across cultural contexts either; it may be the case that people from different cultural backgrounds would display different patterns of response in these sorts of learning tasks. Finally, the cross-sectional design employed here does not allow for the same depth of insight into developmental trajectories as longitudinal designs could do otherwise. Further tests of the proposed rational learning account that address the above limitations are important steps for future research.

For example, it is possible that our results on simple, subject-based norms (e.g., wearing a ribbon) do not generalize to more complex norms in which agents and patients are both implicated. This is an area that is worth further exploration. Likewise, our studies focused on adults and children’s learning about neutral-valence norms that applied to novel 3rd parties. Of course, not all norms have this sort of content.

Another important aim for future research is investigating whether rational learning can help to explain the acquisition of norms with valenced content. For example, it is plausible that learners have different priors about norms with morally charged content. Previous work (Schmidt et al., 2012) suggests that children are inclusive about moral norms but parochial about conventional norms. Several studies by Chalikh & Rhodes, 2014, 2018 have found that children's willingness to extend a candidate norm across group boundaries depends in part on the harmfulness or seriousness of violating the norm. In addition, Chernyak, Kang, and Kushnir (2019) found that children in the US and Singapore display differences in how they respond to social norm and moral norm violations. The rational learning framework introduced here can also be used as a tool for comparing such cultural differences.

In the meantime, this work contributes to a growing body of research that looks to formalize and test computational models of social learning in children and adults (Cushman & Gershman, 2019; Jara-Ettinger, Gweon, Schulz, & Tenenbaum, 2016), and therefore contributes to the broadly important methodological project of formal theory-building in the psychological sciences (c.f., Borsboom, van der Maas, Dalege, Kievit, & Haig, 2021; Muthukrishna & Henrich, 2019). Along these lines, one advantage of the Bayesian formalism deployed here is that it can be naturally extended to address follow-up questions in future research. For example, individuals often occupy multiple social positions or roles within the same group. It will be important for future studies to assess whether learners make rational statistical inferences from evidence that applies to non-disjoint hypotheses about social categories.

The central claim that we have investigated here—whether mere group representations elicit group-based biases in norm acquisition—is a claim about how the mind is designed. Thus, this claim can be studied in artificial settings where the hypothesized functional components can be (reasonably well) isolated. That being said, real-world evidence about the scope of norms routinely takes the form of sparse, specific instances, which can include explicit normative language. Consider distinctions of scope that are acquired from directed injunctions (“Don’t do that!”) or when a child receives the first piece of evidence about a norm’s scope (“This is *our* rule”). At this moment, it appears that inclusivity remains a live hypothesis in children’s mind. The rational learning account we have developed provides a natural explanation for why this is the case (see Section 4). By contrast, accounts that place tribal or coalitional biases front-and-center will require extra machinery to accommodate this core phenomenon in norm acquisition.

From applying the Bayesian formalism to the case of parochialism, we have shown that group-based biases are not necessary for the acquisition of parochial norms. Instead, we have demonstrated how the acquisition of parochial norms can also result from rational statistical inference over the learner’s evidence. In so far as such rational processes are implicated in real-world norm acquisition, then this would give us reason to believe group-based biases may not be as deep-rooted or immutable as commonly thought. On this picture, the human tendency toward parochialism is not merely due to our group representations, nor is it a hard-wired component of human cognition. Instead, parochialism results from the evidence that learners receive. If this account is correct, then we might explain the existence of troubling parochial norms and moralities as a ‘garbage in, garbage out’ problem. The solution? Provide learners with more evidence that is commensurate with inclusive norm systems and moralities. Of course, that is easier said than done. But in so far as the rational learning account describes key elements of norm acquisition, then efforts in this direction are on the right track.

Data availability

See OSF for all data and code for statistical analyses: https://osf.io/njv3g/?view_only=e89feb548744f0d95803031797c4cf8.

Acknowledgements

The writing of this paper was supported by NSF BCS-1823658, NSF SL- 1955280 and USDA Hatch NYC-321434 awards to T.K. We would like to thank Yeonju Suh, Loie Faulkner, and Elizabeth Latella for assistance with data collection and coding. We would also like to thank participating children and families. Thanks to members of the Early Childhood and Cognition Lab, and audiences at the 2020 meeting of the Cognitive Science Society, the 2021 meeting of the Society for Research in Child Development, and the 2021 meeting of the Society for Philosophy and Psychology for helpful comments and discussion.

Credit statement

TK and SN conceptualized project and TK obtained funding. Design of studies by SP, TK, and SN. Execution of research by SP. SP analyzed data and wrote the first draft of manuscript; all authors contributed to the writing of the final manuscript.

Declarations of interest

Declarations of Interest: none

References

- Anderson, John Robert. (1990). *The Adaptive Character of Thought*. Psychology Press.
- Austerweil, J. L., & Griffiths, T. L. (2013). A nonparametric Bayesian framework for constructing flexible feature representations. *Psychological Review*, 120(4), 817.
- Ayars, A., & Nichols, S. (2017). Moral empiricism and the bias for act-based rules. *Cognition*, 167, 11–24. <https://doi.org/10.1016/j.cognition.2017.01.007>
- Ayars, A., & Nichols, S. (2020). Rational learners and metaethics: Universalism, relativism, and evidence from consensus. *Mind & Language*, 35(1), 67–89.
- Bian, L., Sloane, S., & Baillargeon, R. (2018). Infants expect ingroup support to override fairness when resources are limited. *Proceedings of the National Academy of Sciences*, 115(11), 2705–2710. <https://doi.org/10.1073/pnas.1719445115>
- Birnbaum, D., Deeb, I., Segall, G., Ben-Eliyahu, A., & Diesendruck, G. (2010). The development of social essentialism: The case of Israeli Children’s inferences about Jews and Arabs. *Child Development*, 81(3), 757–777. <https://doi.org/10.1111/j.1467-8624.2010.01432.x>
- Borsboom, D., van der Maas, Dalege, J., Kievit, R. A., & Haig, B. D. (2021). Theory construction methodology: A practical framework for building theories in psychology. *Perspectives on Psychological Science*, 16(4), 756–766.
- Boyd, R., & Richerson, P. J. (2009). Culture and the evolution of human cooperation. *Philosophical Transactions of the Royal Society, B: Biological Sciences*, 364(1533), 3281–3288.
- Buresh, J. S., & Woodward, A. L. (2007). Infants track action goals within and across agents. *Cognition*, 104(2), 287–314.
- Chalikh, L., & Rhodes, M. (2014). Preschoolers use social allegiances to predict behavior. *Journal of Cognition and Development*, 15(1), 136–160.
- Chalikh, L., & Rhodes, M. (2018). Learning about social category-based obligations. *Cognitive Development*, 48, 117–124.
- Chalikh, L., & Rhodes, M. (2020). Chapter three - groups as moral boundaries: A developmental perspective. In J. B. Benson (Ed.), *Vol. 58. Advances in child development and behavior* (pp. 63–93). JAI. <https://doi.org/10.1016/bs.acdb.2020.01.003>
- Chernyak, N., Kang, C., & Kushnir, T. (2019). The cultural roots of free will beliefs: How Singaporean and US children judge and explain possibilities for action in interpersonal contexts. *Developmental Psychology*, 55(4), 866.
- Chudek, M., & Henrich, J. (2011). Culture–gene coevolution, norm-psychology and the emergence of human prosociality. *Trends in Cognitive Sciences*, 15(5), 218–226.
- Clark, C. J., Liu, B. S., Winegard, B. M., & Ditto, P. H. (2019). Tribalism is human nature. *Current Directions in Psychological Science*, 28(6), 587–592.
- Cushman, F., & Gershman, S. (2019). Editors’ introduction: Computational approaches to social cognition. *Topics in Cognitive Science*, 11(2), 281–298. <https://doi.org/10.1111/tops.12424>
- Diehl, M. (1990). The minimal group paradigm: Theoretical explanations and empirical findings. *European Review of Social Psychology*, 1(1), 263–292.
- 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.
- Douven, I., & Schupbach, J. N. (2015). The role of explanatory considerations in updating. *Cognition*, 142, 299–311. <https://doi.org/10.1016/j.cognition.2015.04.017>
- Dunham, Y. (2018). Mere membership. *Trends in Cognitive Sciences*, 22(9), 780–793.
- Dwyer, S., Huebner, B., & Hauser, M. D. (2009). The linguistic analogy: Motivations, results, and speculations. *Topics in Cognitive Science*, 2(3), 486–510.
- Eason, A. E., Doctor, D., Chang, E., Kushnir, T., & Sommerville, J. A. (2018). The choice is yours: Infants’ expectations about an agent’s future behavior based on taking and receiving actions. *Developmental Psychology*, 54(5), 829.

- Edelman, S., & Shabbazi, R. (2011). Survival in a world of probable objects: A fundamental reason for Bayesian enlightenment. *Behavioral and Brain Sciences*, 34(4), 197.
- Efferon, C., Lalive, R., & Fehr, E. (2008). The coevolution of cultural groups and Ingroup favoritism. *Science*, 321(5897), 1844–1849. <https://doi.org/10.1126/science.1155805>
- Finiasz, Z., Sheth, E., Karami, D., Gelman, S., & Kushnir, T. (2022, April). Generic testimony interacts with statistical evidence in adults' and children's category-based induction. In *Poster presented at the 12th biennial meeting of the cognitive development society. Madison, WI.*
- Garvin, L. E., & Woodward, A. L. (2015). Verbal framing of statistical evidence drives children's preference inferences. *Cognition*, 138, 35–48.
- Gelman, S. A., Ware, E. A., & Kleinberg, F. (2010). Effects of generic language on category content and structure. *Cognitive psychology*, 61(3), 273–301.
- Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A Rational Analysis of Rule-Based Concept Learning. *Cognitive Science*, 32(1), 108–154. <https://doi.org/10.1080/03640210701802071>
- Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, 138(6), 1085–1108. <https://doi.org/10.1037/a0028044>
- Graham, J., Waytz, A., Meindl, P., Iyer, R., & Young, L. (2017). Centripetal and centrifugal forces in the moral circle: Competing constraints on moral learning. *Cognition*, 167, 58–65. <https://doi.org/10.1016/j.cognition.2016.12.001>
- Greene, J. D. (2013). *Moral tribes: Emotion, reason, and the gap between us and them.* Penguin.
- Gweon, H., & Schulz, L. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, 332(6037), 1524–1524.
- Heck, I. A., Bas, J., & Kinzler, K. D. (2022). Small groups lead, big groups control: Perceptions of numerical group size, power, and status across development. *Child Development*, 93(1), 194–208.
- Heck, I. A., Kushnir, T., & Kinzler, K. D. (2021). Social sampling: Children track social choices to reason about status hierarchies. *Journal of Experimental Psychology: General*, 150(8), 1673.
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences*, 20(8), 589–604.
- Jin, K., & Baillargeon, R. (2017). Infants possess an abstract expectation of ingroup support. *Proceedings of the National Academy of Sciences*, 114(31), 8199–8204. <https://doi.org/10.1073/pnas.1706286114>
- Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical Bayesian models. *Developmental Science*, 10(3), 307–321. <https://doi.org/10.1111/j.1467-7687.2007.00585.x>
- Kinzler, K. D., Dupoux, E., & Spelke, E. S. (2007). The native language of social cognition. *Proceedings of the National Academy of Sciences*, 104(30), 12577–12580. <https://doi.org/10.1073/pnas.0705345104>
- Kinzler, K. D., Shutts, K., DeJesus, J., & Spelke, E. S. (2009). Accent trumps race in guiding children's social preferences. *Social Cognition*, 27(4), 623–634.
- Kushnir, T., & Gelman, S. A. (2016). Translating testimonial claims into evidence for category-based induction. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th annual conference of the cognitive science society* (pp. 1307–1312). Austin, TX: Cognitive Science Society.
- 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.
- Lakens, D. (2022). Sample size justification. *Collabra: Psychology*, 8(1), 33267.
- Lewis, M. L., & Frank, M. C. (2018). Still suspicious: The suspicious-coincidence effect revisited. *Psychological Science*, 29(12), 2039–2047.
- Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, 55(3), 232–257.
- 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), Article e92160.
- MacDonald, K., Schug, M., Chase, E., & Barth, H. (2013). My people, right or wrong? Minimal group membership disrupts preschoolers' selective trust. *Cognitive Development*, 28(3), 247–259. <https://doi.org/10.1016/j.cogdev.2012.11.001>
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information.*
- Mayo, D. G. (2018). *Statistical inference as severe testing: How to get beyond the statistics wars.* Cambridge University Press.
- Mikhail, J. (2007). Universal moral grammar: Theory, evidence and the future. *Trends in Cognitive Sciences*, 11(4), 143–152.
- Mikhail, J. (2011). *Elements of moral cognition: Rawls' linguistic analogy and the cognitive science of moral and legal judgment.* Cambridge University Press.
- Moll, H., & Tomasello, M. (2007). How 14- and 18-month-olds know what others have experienced. *Developmental Psychology*, 43(2), 309.
- Muthukrishna, M., & Henrich, J. (2019). A problem in theory. *Nature Human Behaviour*, 3(3), 221–229.
- Navarro, D. J., & Perfors, A. F. (2010). Similarity, feature discovery, and the size principle. *Acta Psychologica*, 133(3), 256–268. <https://doi.org/10.1016/j.actpsy.2009.10.008>
- Nichols, S. (2021). *Rational rules: Toward a theory of moral learning.* Oxford University Press.
- Nichols, S., Kumar, S., Lopez, T., Ayars, A., & Chan, H.-Y. (2016). Rational learners and moral rules. *Mind & Language*, 31(5), 530–554. <https://doi.org/10.1111/mila.12119>
- Noyes, A., Dunham, Y., Keil, F. C., & Ritchie, K. (2021). Evidence for multiple sources of inductive potential: Occupations and their relations to social institutions. *Cognitive Psychology*, 130, Article 101422.
- Partington, S., Nichols, S., & Kushnir, T. (2021). Is children's norm learning rational? A meta-analysis. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 43(43).
- Pellizzoni, S., Siegal, M., & Surian, L. (2010). The contact principle and utilitarian moral judgments in young children. *Developmental Science*, 13(2), 265–270.
- Perfors, A., Tenenbaum, J. B., Griffiths, T. L., & Xu, F. (2011). A tutorial introduction to Bayesian models of cognitive development. *Cognition*, 120(3), 302–321. <https://doi.org/10.1016/j.cognition.2010.11.015>
- Pettigrew, R. (2016). *Accuracy and the Laws of credence.* Oxford University Press.
- Pisor, A., & Ross, C. T. (2021). *How generalizable are patterns of parochial altruism in humans?* Unpublished Manuscript.
- Prasada, S., Khemlani, S., Leslie, S.-J., & Glucksberg, S. (2013). Conceptual distinctions amongst generics. *Cognition*, 126(3), 405–422. <https://doi.org/10.1016/j.cognition.2012.11.010>
- Rakoczy, H., Warneken, F., & Tomasello, M. (2008). *The sources of normativity: Young children's awareness of the normative structure of games* (pp. 875–881). Developmental Psychology: In Press.
- Rand, D. G., & Nowak, M. A. (2013). Human cooperation. *Trends in Cognitive Sciences*, 17(8), 413–425.
- Read, K. E. (1955). Morality and the concept of the person among the Gahuku-Gama. *Oceania*, 25(4), 233–282.
- Rhodes, M., & Chalik, L. (2013). Social categories as markers of intrinsic interpersonal obligations. *Psychological Science*, 24(6), 999–1006.
- Riggs, A. E. (2019). Social statistics: Children use statistical reasoning to guide their inferences about the scope of social behavior. *Developmental Psychology*, 55(1), 66.
- Riggs, A. E., & Long, M. (2020). The domain frequency association: A mental shortcut to guide children's generalization of norms and preferences. *Cognitive Development*, 54, Article 100853.
- 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. <https://doi.org/10.1016/j.jecp.2017.03.018>
- Roberts, S. O., Ho, A. K., & Gelman, S. A. (2017). Group presence, category labels, and generic statements influence children to treat descriptive group regularities as prescriptive. *Journal of Experimental Child Psychology*, 158, 19–31. <https://doi.org/10.1016/j.jecp.2016.11.013>
- Roberts, S. O., Ho, A. K., & Gelman, S. A. (2019). The role of group norms in evaluating uncommon and negative behaviors. *Journal of Experimental Psychology: General*, 148(2), 374.
- Roberts, S. O., & Horii, R. I. (2019). Thinking fast and slow: Children's descriptive-to-prescriptive tendency under varying time constraints. *Journal of Cognition and Development*, 20(5), 790–799.
- Schmidt, M. F., Rakoczy, H., & Tomasello, M. (2012). Young children enforce social norms selectively depending on the violator's group affiliation. *Cognition*, 124(3), 325–333.
- Smetana, J. G. (2013). *Moral development: The social domain theory view.*
- Snare, F. E. (1980). The diversity of morals. *Mind*, 89(355), 353–369.
- Srinivasan, M., Kaplan, E., & Dahl, A. (2019). Reasoning about the scope of religious norms: Evidence from Hindu and Muslim children in India. *Child Development*, 90(6), e783–e802.
- Tajfel, H. (1970). Experiments in intergroup discrimination. *Scientific American*, 223(5), 96–103. JSTOR.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, 24(4), 629–640. <https://doi.org/10.1017/S0140525X01000061>
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *science*, 331(6022), 1279–1285.
- Tessler, M. H., Bridgers, S., & Tenenbaum, J. B. (2020). How many observations is one generic worth?. *Proceedings of the annual conference of the cognitive science society.*
- Turiel, E. (2006). *The development of morality.*
- Turnbull, C. (1972). *The Mountain People.* New York: Simon & Schuster.
- Van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science*, 32(6), 939–984.
- Vélez, N., Wu, Y., & Gweon, H. (2018). *Consistent but not diagnostic: Preschoolers' intuitions about shared preferences within social groups.* CogSci.
- Wellman, H. M., Kushnir, T., Xu, F., & Brink, K. A. (2016). Infants use statistical sampling to understand the psychological world. *Infancy*, 21(5), 668–676.
- Wright, J. C., & Bartsch, K. (2008). Portraits of early moral sensibility in two children's everyday conversations. In *Merrill-palmer quarterly* (1982-) (pp. 56–85).
- Xu, F. (2019). Towards a rational constructivist theory of cognitive development. *Psychological Review*, 126(6), 841.
- Xu, F., & Kushnir, T. (2012). *Rational constructivism in cognitive development, (Vol. 43).* Academic Press.
- Xu, F., & Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science*, 22(1), 28–32.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2), 245–272.
- Yamagishi, T., & Mifune, N. (2016). Parochial altruism: Does it explain modern human group psychology? *Current Opinion in Psychology*, 7, 39–43. <https://doi.org/10.1016/j.copsyc.2015.07.015>