Can Children and Adults Balance Majority Size with Information Quality in Learning from Preferences?

Rebekah A. Gelpí[1], A. Whalen[2], Thomas L. Griffiths[3,4], Fei Xu[5], and Daphna Buchsbaum[6]

[1] Department of Psychology, University of Toronto

[2] Roslin Institute

[3] Department of Psychology, Princeton University

[4] Department of Computer Science, Princeton University

[5] Department of Psychology, University of California, Berkeley

[6] Department of Cognitive, Linguistic, and Psychological Sciences, Brown University

Author Note

Rebekah A. Gelpí https://orcid.org/0000-0001-6899-0520

A. Whalen https://orcid.org/0000-0002-6922-0947

Thomas L. Griffiths https://orcid.org/0000-0002-5138-7255

Fei Xu https://orcid.org/0000-0002-8220-7528

Daphna Buchsbaum https://orcid.org/0000-0002-8716-7756

We have no known conflicts of interest to declare.

Correspondence concerning this article should be addressed to Rebekah A. Gelpí, Department of Psychology, University of Toronto, 100 St. George Street, Toronto, ON, Canada M5S 3G3.

Email: rebekah.gelpi@mail.utoronto.ca

Word count: 13,697 (excluding figures and references)

This study was not preregistered. Anonymized data and analysis code are available at the following link: https://osf.io/ekzbm/?view_only=67fc7cdefc284240a17204e895d97880

Abstract

We investigate how 3- to 5-year-old US and Canadian children ($N = 189$) and US adults ($N = 241$) balance the number of endorsements for a given option with the quality of the informants' source of information when deciding which of two boxes contains the better option. When choosing between two different boxes endorsed by groups of equal sizes, both children (Experiments 1–3) and adults (Experiment 6) tend to choose boxes endorsed by informants with visual access to the boxes over informants with hearsay. However, children's choices were biased towards the larger group when the size of the group conflicted with the quality of the source of the groups' information (Experiments 4–5), while adults more often chose the option endorsed by the group with the higher quality information (Experiment 6). Children were more likely to conform to a majority opinion when compared to both adults and to a normative computational model that endorses a group proportional to the number of independent, direct observations made by that group's informants. These findings suggest that, while adults balance the size of a majority with the quality of the informants' information source, preschoolers can evaluate when groups differ in the source of their information, but may assume that the presence of a majority endorsing an option is inherently informative over and above the information source group members' testimony relied on.

Keywords: social learning; testimony; consensus; conformity bias

Public Significance Statement

This study suggests that young children's intuitions about what kinds of information to trust is similar to adults' in some ways; children considering that people with direct access to a piece of information should be relied upon more than people whose information comes from hearsay.

However, our study finds that children consider a larger number of people endorsing one option over another inherently informative, while adults balance the number of people and their access to information appropriately. This finding offers us insight into children's emerging understanding about how to evaluate the quality of a piece of information based on its source.

Can children and adults balance majority size with information quality in learning from preferences?

Imagine you want to try a new restaurant and ask some friends for suggestions of what to order. Four friends suggest that the pizza is better, while another friend suggests that the pasta is better. All else being equal, you would probably order the pizza. It often makes sense to follow a majority, especially if we have little or incomplete information, because we assume that others are broadly rational, and have good reasons for their behaviors and preferences, and they may have based their decisions on information or evidence we do not have access to (e.g., Morgan et al., 2012). A body of theoretical work has suggested that conforming to a majority is one of several contextually successful social learning strategies that people engage in (e.g., Henrich & Boyd, 1998; Hoppitt & Laland, 2013; Kendal et al., 2018; Rendell et al., 2011; Whalen, Griffiths & Buchsbaum, 2018).

For children, who have comparatively little expertise and fewer life experiences, learning from others' actions can be especially beneficial, offering the opportunity to acquire large amounts of information without having to engage in time-consuming, costly, and possibly even dangerous trial-and-error. This capacity for social learning is a cornerstone of human society, and it has been proposed to be a driving force in our cultural evolution and ultimate success as a species (Boyd & Richerson, 1985; Boyd et al., 2011; Csibra & Gergely, 2009; Tomasello, 1999).

However, depending on how the people we are learning from came to their own decisions, there are cases where following a majority can also lead us astray (Bikhchandani et al., 1992; Anderson & Holt, 1997). People can be ignorant, make mistakes, or even intentionally mislead others, and those learning from them may receive information from multiple people

whose testimony conflicts. If people are not discerning in evaluating majority information, they may accept inaccurate information, and conform to an incorrect majority. Further, people must keep track of other cues to the reliability or informational quality of others' testimony beyond the size of the group that endorses an option, such as the degree to which individuals within a group are sharing a source of evidence. If the majority of a group endorses an option (e.g., that a restaurant's pizza is better than its pasta), but this endorsement results from a single, shared primary source of evidence (all hearing from the same friend who once had a bad pizza), their endorsements may be less informative than if their endorsements result from independent converging evidence (e.g., each individual tried the pizza separately and separately preferred it).

Several recent studies have sought to understand the contexts in which adults do or do not exhibit a bias towards numerical majorities[1], above and beyond the information they provide, in situations where groups of people disagree or prefer different options. In some cases, adults seem to show an "illusion of consensus", wherein a consensus that exhibits statistical dependency (i.e., all relying on a single source) is considered to be as reliable as a "true consensus" of multiple independent sources (Alister et al., 2022; Desai et al., 2022; Yousif et al., 2019). However, when the source of the information that informants are basing their testimony on is made transparent, adults appropriately adjust their degree of endorsement of majority, rating majorities with a greater number of converging sources of data as more credible than those with fewer

[1] Judging when humans' reliance on a majority endorsing an option is appropriate or inappropriate can be difficult, in part because there are differing definitions of what constitutes a majority bias or "conformity" (see Whiten, 2019). For example, copying a behavior or belief in proportion to how often it appears in a group can result in individuals exhibiting a numerical tendency to endorse a majority—that is, endorsing a majority behavior or belief more often than a minority one—but this is not the same as conformity or a majority *bias*, which we define as a situation in which an individual endorses a majority to a greater degree than normatively predicted, for instance if a numerical majority nonetheless provides lower quality information.

independent sources of data (Alister et al., 2022; Desai et al., 2022; Mercier & Miton, 2019; Whalen et al., 2018).

Understanding when young children develop the ability to monitor the independence of an informant's sources of information—and when they may be susceptible to a similar "illusion of consensus" as adults—is particularly important given children's reliance on learning from others in early life. Here we examine whether, and when, children and adults are sensitive to the source and quality of informants' testimony, and how they use this to assess the quality of not only individual informants, but also of groups of informants who differ in the source and quality of their testimony.

Cues to Information Quality in Children's Selective Trust

To effectively learn about the world, children must develop a sense of selective trust, believing those whom they consider accurate and reliable sources of testimony. A large body of literature about children's trust in testimony has found that children selectively trust informants, and are sensitive to a wide variety of cues to informant reliability, including past accuracy and perceived expertise (for reviews see e.g., Harris et al., 2018; Landrum et al., 2015; Mills, 2013; Robinson & Einav, 2014; Sobel & Kushnir, 2013).

One valuable cue to informant quality that children use is perceptual access. For example, if a child knows that a potential informant has seen inside a box, then that person's statements about the contents of the box are more useful than someone who has not looked inside. By age three, young children understand that visual experience provides informants with knowledge (e.g. O'Neill et al., 1992; Pillow, 1989; Sodian & Wimmer, 1987); consequently, they prefer to

get their information from people who have seen something directly (e.g., Butler et al., 2018, 2020; Povinelli & deBlois, 1992; Robinson et al., 2008; but see Palmquist & Jaswal, 2012).

However, in many situations, children may not have information about the past accuracy or knowledge states of a potential informant. In situations like this, children may instead rely on other cues to information quality, such as evaluating what the majority of people believe (Corriveau et al., 2009), and endorsing or imitating the majority's choice. For example, 3- and 4-year-olds endorse novel object labels given by a majority over those given by a dissenter (Corriveau et al., 2009; Pham & Buchsbaum, 2020), and 2-year olds are more likely to imitate a majority's actions over those of an equally successful minority (Haun et al., 2012). Children endorse majorities more consistently in conventional domains such as language tasks, compared to domains where asocial learning is also possible, such as causal learning (Pham & Buchsbaum, 2020). Children may also endorse a majority's judgment when their own perceptual evidence is uncertain (Bernard et al., 2015; Morgan et al., 2015). The finding that children conform to a majority's choice across multiple contexts has led to the suggestion that children may have a consistent bias to conform to the majority, regardless of the quality of the majority's testimony, as this would be an efficient and generally accurate social learning heuristic (e.g., Walker & Andrade, 1996; Haun & Tomasello, 2011).

However, the fact that a numerical majority makes a certain choice or engages in a certain behavior does not always indicate that an option is the best; majorities can be less successful at a task, make implausible claims, or base their choices on fewer primary sources. Nevertheless, the existing evidence about children's ability to make inferences about groups' information quality is mixed. Some studies suggest that as young as 4 years of age, children preferentially attend to quality of information over the size of the group endorsing the claim: for

instance, 4-year-old children will copy a successful dissenter over an unsuccessful majority in an instrumental learning task (Wilks et al., 2015), are less likely to endorse a majority's description of an object's function if that function is implausible (Schillaci & Kelemen, 2014), and will endorse the identity of a drawing given by the artist rather than that given by a conflicting majority (Einav, 2014). Others have found evidence showing that children under age six are swayed by the presence of a majority, even when there are other cues to information quality available: for example, 4-year-olds did not consistently endorse an informant with a past history of success over a conflicting majority with unknown expertise (Burdett et al., 2016; Sampaio et al., 2019). Likewise, Bernard and colleagues (2015) found that 4-year-olds endorsed a previously unreliable majority rather than a previously reliable minority, while 6-year-olds endorsed the previously reliable minority.

Another cue to information quality is the degree of statistical independence of sources: that is, understanding that multiple informants who received their data from a single source do not inherently have more information than a single informant with a single source. Here, young children also appear to display a bias towards conforming beyond what is rational. For example, 4- and 5-year-old children endorsed a majority that shared a single data point as often as a majority with independent data points (Otsubo et al., 2017). Aboody and colleagues (2022) also found a developmental transition in the consideration of information quality: 6-year-old children believed an individual whose claim was supported by multiple independent informants more than multiple individuals whose claims relied on a single informant. However, 4-year-olds did not display a clear tendency to endorse either the majority with a single source or an individual with multiple sources.

Given 4- and 5-year-old children's ability to reason about sources of information, and to selectively trust informants along many dimensions (e.g., Birch et al., 2008; Jaswal & Neely, 2006; Koenig & Harris, 2005), the mixed pattern of results in studies of conformity to a majority among 4- and 5-year-old children may reflect multiple possibilities. In many previous studies, the size of a majority and the quality of the statistical information provided by the informants was not clearly differentiated; therefore, the degree to which endorsement of a majority would reflect conformity—rather than the normative choice given the data presented to children—has not been clear. Ambiguity about the quality of a majority's source of information has also been offered as an explanation for why adults sometimes fall victim to an "illusion of consensus" and other times do not (e.g., Alister et al., 2022; Desai et al., 2022).

By explicitly manipulating the size of the majority and the quality of the information that children receive we can clarify whether children are likewise capable of using cues to the quality of a group's testimony when the nature of the group's sources are clear, or whether children simply exhibit a strong conformity bias (as suggested by e.g., Walker & Andrade, 1996; Haun & Tomasello, 2011) above and beyond what is rational.

Learning About Preferences from Others

Additionally, children's evaluation of information quality may extend beyond trying to determine factual information. Many studies of children's endorsement of testimony rely on their evaluation of facts, such as the location of a hamster (Aboody et al., 2022) or how to open a puzzle box (Wilks et al., 2015). In these cases, while someone might consider the perspective of multiple informants if they themselves are uncertain of the answer, there is an underlying ground truth: the hamster must really be in one location, and the puzzle box has a true solution.

In contrast to factual testimony, there is no ground truth when we hear testimony about another person's preferences: if I prefer broccoli and you prefer goldfish, neither of us is objectively "right". Despite this, many preferences are in fact broadly shared, and so testimony can serve as probabilistic evidence that a person may prefer the same item as the informant. Consistent with this, children can use information provided by others to learn their own preferences, such as food preferences (e.g., Birch, 1999; Ventura & Worobey, 2013) or music preferences (e.g., Hargreaves et al., 2015; Lamont & Crich, 2022). By 3 years of age, young children have developed an understanding that preferences are often broadly shared (Vélez et al., 2018), but can differ between individuals (Lucas et al., 2014). Other people's preferences may be particularly informative when we have little personal information to go on (e.g., whether we will like a movie we haven't seen, or a restaurant we've never been to), as children often are early in life.

However, the cues to information quality that children and adults consider to be important may differ depending on whether an informant's testimony is about facts or preferences. In factual domains, indirect information such as hearsay does not directly provide additional knowledge about what happened. In the case of preferences, however, people might still attend to indirect informants' testimony because they perceive their agreement with another source to be inherently informative in its own right. This may explain why children may conform more strongly to majorities in conventional domains, such as object labelling, than in domains such as causal learning (Pham & Buchsbaum, 2020). Thus, a child seeing multiple individuals preferring one unseen option over another may serve as a more graded form of evidence that an option will be preferable to the child as well. While previous work has found that adults can sometimes balance the size of a majority with the quality of their information in a factual domain

(Whalen et al., 2018), neither adults nor children's ability to do this in the domain of preferences has been explored.

Thus, the case of preferences provides an interesting opportunity to consider how both children and adults evaluate the quality of an information source. We hypothesize that both children and adults may use cues to information quality such as the presence of a majority or of a greater number of primary sources endorsing a claim not only to reason about facts, but also about what they themselves are likely to prefer when informants state their own preferences.

Here, we examine how both children and adults reconcile conflicting endorsements from groups of informants with varying degrees of first-hand knowledge of options to choose from, where the option are unknown to the learner. We will particularly focus on understanding of individuals with direct knowledge versus indirect knowledge (i.e., hearsay). Given that preschool-age children preferentially seek information from those with first-hand knowledge (e.g., Butler et al., 2018, 2020), in Experiments 1–3 we first explore whether children in this age range use this cue when they are evaluating testimony from equally sized groups of informants about their *item preferences*, which we use here to refer to which of the two options presented in our task an informant prefers[2]. This also allows us to determine how strongly children tend to endorse an option endorsed by informants with first-hand information when no majority is present, allowing us to more systematically test in later experiments whether—and to what extent—children's and adults' endorsements on the task constitute a majority bias.

We then outline two competing computational models of learning from testimony which predict how (1) a rational learner who is able to normatively evaluate both information quality and majority size, (2) a conformity-biased learner who treats majority size as a heuristic

[2] In the context of our experiments, we use the word "preference" throughout the text to refer to an informant's item preferences, or a learner's potential item preferences.

indicating quality, and (3) a learner that mixes both the normative and conformist strategies might evaluate evidence in a number of scenarios when information quality and group size conflict. In Experiments 4–5, we test the predictions of these models for children's behavior, by examining whether children's inferences are similar to those of the normative model, or whether—and to what degree—they instead display a bias to conform to a majority, even when that majority provides lower quality information. Finally, in Experiment 6, we compare children's responses as well as the model predictions to the performance of adults on these same tasks. By comparing the model's predictions with children's and adults' responses, we can illuminate the extent to which their choices to follow the majority are a rational result of the majority's additional informativeness, and under what conditions they are not.

Experiment 1: Direct knowledge vs. hearsay

In Experiment 1, participants watched as informants gave opinions about which of two boxes contained the better hidden option. Equal numbers of informants endorsed each box, but one box was endorsed by informants who had looked in the boxes and had direct knowledge of what was inside, whereas the other box was endorsed by only one informant with direct knowledge while the other three received hearsay about which box was better. Choosing the box endorsed by the direct group would suggest that children are monitoring individual informants' knowledge quality and not just the number of endorsements per item.

Methods

Participants. Participants were 22 3- to 5-year-old children (mean age = 49 months; range = 43 – 66 months) recruited from a large US metropolitan area, and were tested in the lab, their preschools or at local museums. The sample size was chosen as it is appropriately powered to detect moderate-to-large effect sizes in a summary score of 2 repeated trials (power ≥ 0.80 for

detecting average correct performance of 70% or greater relative to chance, e.g. Rosner (2015); see also Supplementary Material for a derivation specific to our experiment). A range of ethnicities representing the demographics of the local population was represented (see Supplementary Material). Three additional children were excluded due to experimenter error (2) and inattentiveness (1).

Materials. Materials included two black boxes, each of which contained a toy (a toy vehicle or a stuffed animal) or a snack (Goldfish cracker or Froot Loop™). Informants in each trial were eight 7" tall paper dolls (four male, four female), made available online by illustrator Kyle Hinton, glued to a wood block base. Each trial included a set of novel informants (i.e., informants were different across trials).

Procedure. Children participated in two trials: a snack trial and a toy trial. Trial order was counterbalanced, and new materials (i.e., different boxes and different informants) were used for each trial. In each trial, the experimenter first showed the participant the two boxes and explained that each box contained a [toy/snack], but that she did not know what was inside. Children were not shown the contents of the boxes ahead of time, so that their differing levels of familiarity with the option, or pre-existing preferences would not influence their evaluation of the testimony. Then, the child watched as dolls gave opinions about which box contained the better option (Figure 1). A group of four dolls endorsed one box and a second group of four endorsed the other. In the direct group, all four dolls received direct (visual) knowledge before giving their opinions. One at a time, each doll walked over to each box and looked inside, then stood beside the same box and said, "I think this [toy/snack] is better!".

In the indirect group, only the first doll in the group received direct knowledge of the box's contents. The first doll looked inside both of the boxes, then stood next to the box not

endorsed by the direct group and said, “I think this [toy/snack] is better!” This doll then crossed paths with a second doll, and the experimenter made indiscriminate whispering sounds to convey that the two dolls were conversing. The second doll gave their opinion, saying, “[S]he said this [toy/snack] is better, so I think this [toy/snack] is better,” and passed on their hearsay to a third doll, who stated his or her opinion, and then passed the hearsay on to the fourth doll. Each group included equal numbers of male and female dolls, and group order (direct or indirect first) was counterbalanced. The side of the box endorsed by the direct group was also counterbalanced.

After all dolls gave opinions, the experimenter brought all eight dolls back on stage and placed them in front of the box they endorsed, and reminded children that the dolls were all standing in front of the box they had said was better. With both groups of dolls still visible, the experimenter asked the child to choose the box they wanted to try. Once children selected a box, they were presented with the object inside. They were not shown the contents of the unchosen box. The experimenter cleared all materials from the table, and proceeded to the second trial.

Results and Discussion

Results for Experiment 1 are summarized in Table 1. For each trial, children received a 1 if they chose the box endorsed by the direct informants, and a 0 if they chose the box endorsed by the indirect informants. Children chose the direct box over the indirect box significantly more often than chance, $B = 1.07$, $SE = 0.46$, 95% CI = [0.36, 2.49], $OR = 2.92$, $z = 2.35$, $p = .019$. There was no significant difference in responses between the first and second trial, $B = -0.24$, $SE = 0.36$, 95% CI = [-1.01, 0.44], $z = -0.69$, $p = .490$, or for the two trial types (snack vs toy), Fisher exact test, $OR = 0.39$, $p = .31$.

When choosing between two boxes, each endorsed by four informants, children choose the box endorsed by informants with direct knowledge of the boxes’ contents. This suggests that

children monitor the knowledge quality of individual informants within a group, and not just group size. Additionally, this suggests that they understand that visual access is a more reliable source of information than hearsay, even when learning about non-factual domains like preferences.

Experiment 2: Hearsay vs shared knowledge

In Experiment 1 we manipulated two different cues to the quality of the indirect group's testimony. First, the indirect group was making their response based on hearsay, and second, the indirect group was making their response based on a shared source of knowledge: only the first informant directly observed boxes. Both hearsay and shared information could reduce the perceived quality of a group's testimony, so given the results of Experiment 1 it is not possible to determine if children are sensitive to hearsay, shared information, or both. To examine the role of hearsay in a situation without shared knowledge, in Experiment 2 each indirect informant gives testimony based on hearsay from a different (unseen) individual.

Methods

Participants. Participants were 24 3- to 5-year-old children (mean age = 58 months; range = 46 to 70 months; 14 female, 10 male) recruited from a large Canadian metropolitan area, and were tested in the lab, their preschools and local museums (see Supplementary Material for a replication of Experiment 1 in the same geographic region). A range of ethnicities representing the demographics of the local population was represented (see Supplementary Material). 11 additional children were tested but excluded due to experimenter error ($N = 9$), or inattentiveness (N = 2).

Materials. All materials were the same as in Experiment 1.

Procedure. The procedure of Experiment 2 was identical to Experiment 1, except that in the indirect group, the first informant did not look into either box, informants did not cross each other after producing testimony and did not whisper information to each other (Figure 2). Instead, each informant said "My friend [Jane] said that this [toy/snack] is better, so I think this one is better". The name [Jane] was replaced by a different name (e.g., Tom) for each informant, always of the opposite gender of the informant.

Results and Discussion

Results for Experiment 2 are summarized in Table 1. For each trial, children received a 1 if they chose the box endorsed by the direct informants, and a 0 if they chose the box endorsed by the indirect informants. Children selected the box endorsed by the direct group significantly more often than chance, $B = 0.80$, $SE = 0.32$, 95% CI = [0.21, 1.49], $OR = 2.23$, $z = 2.55$, $p = .011$. There was no significant difference in responses between the first and the second trial, $B = -0.29$, $SE = 0.32$, 95% CI = [-0.94, 0.32], $z = -0.97$, $p = .353$, or for the two trial types, Fisher exact test, $OR = 0.93$, $p = .50$.

As in Experiment 1, we find that children choose the option endorsed by the direct group when given an option of following informants with direct visual access over informants with indirect visual access. The result holds true even when source of information is disentangled from shared knowledge.

Experiment 3: Hearsay from multiple sources vs. one source

Experiment 2 clarified that children are sensitive to direct versus indirect sources of knowledge. In Experiment 3 we examine whether they are sensitive to shared knowledge. As in Experiments 1 and 2, participants in Experiment 3 watched as informants gave opinions about which of two boxes contained the better option. In Experiment 3, the informants differed in the

independence of each informants' source of knowledge. Similar to Aboody et al. (2022), in which informants provided second-hand knowledge about a fact, all informants in Experiment 3 gave testimony based on second-hand knowledge (hearsay), but one box was endorsed by informants who each received hearsay from *different sources* (i.e., independent), whereas the other box was endorsed by informants who each received hearsay from *the same source* (i.e., dependent).

Methods

Participants. Participants were 24 3- to 5-year-old children (mean age = 51 months; range = 40 – 62 months; 14 female, 10 male). Participants were recruited from a large US metropolitan area, and were tested in the lab, their preschools and local museums. A range of ethnicities representing the demographics of the local population was represented (see Supplementary Material). An additional three children were tested, but were excluded due to inattentiveness.

Materials. Like Experiments 1 and 2, materials included two black rectangular boxes, each of which contained a snack or a sticker (results from a preliminary condition of Experiment 1 using stickers showed that a condition using stickers did not differ significantly from the original snack or toy conditions). Two additional paper dolls were used, for a total of ten for each trial.

Procedure. Children participated in two trials: a snack trial and a sticker trial. The procedure of Experiment 3 was identical to Experiment 1 with the following changes. In the testimony phase of the experiment, the child watched as the experimenter introduced four dolls (the *source dolls*), who each looked inside both of the boxes (Figure 3). These four dolls were

then put in a separate area on one side of the demonstration table, where they were still visible to the child.

Then, six *informant dolls* came on stage one at a time. Each encountered a source doll who was "taking a walk" away from the source doll area towards the informant doll. The informant doll whispered with this source doll. Of the six informant dolls, three endorsed one box, and three endorsed the other. These two groups differed in which source doll(s) they whispered with before giving their opinions. In the *independent group*, the three informant dolls received information by each individually whispering with their own, independent source doll. In the *dependent group*, all three informant dolls whispered with the same source doll. Group order and side of box endorsed by independent group (left or right) were counterbalanced.

After each informant doll talked with a source doll, (s)he endorsed a box by saying to the source doll: "Oh, you think this box is better? Well, then, I think this box is better, too." Then, the informant doll remained in front of the box they endorsed, while the source doll returned to the source doll area of the table. Once all six informant dolls had given opinions, the experimenter removed the source dolls from the table. Children were then reminded of which box each group of informant dolls had endorsed and asked to choose a box, as in Experiments 1 and 2. Source dolls in trial 1 were always informant dolls in trial 2, and genders of dolls in independent and dependent groups (2 males, 1 female vs. 2 females, 1 male) were also changed between trials.

Results and Discussion

Results for Experiment 3 are summarized in Table 1. For each trial, children received a 1 if they chose the box endorsed by the independent informants, and a 0 if they chose the box

endorsed by the dependent informants. . Children selected the box endorsed by the independent group significantly more than chance, $B = 0.73$, $SE = 0.36$, 95% CI = [0.10, 1.71], $OR = 2.08$, $z = 2.05$, $p = .041$. There was no significant difference in responses between the first and the second trial, $B = 0.20$, $SE = 0.32$, 95% CI = [-0.42, 0.86], $z = 0.62$, $p = .534$, or for the two trial types, Fisher exact test, $OR = 0.93$, $p = .50$.

When all informants have only indirect knowledge of the box contents, children correctly endorse the group whose knowledge comes from independent testimony. This result suggests that the difference in Experiment 1 is not solely due to children's understanding of hearsay, but also due to their understanding of independence and dependence between informant's testimony. Taken together, Experiments 1–3 suggest that children have a robust sensitivity to the source of informants' knowledge, and can use source and quality of knowledge to accurately evaluate groups of informants.

Modeling the Quality of Informant Testimony

Experiments 1, 2 and 3 find that children are sensitive to both the dependency between informants, and to the source of informants' knowledge—whether their testimony is based on directly observed evidence or on hearsay. In both of these cases, children seem to understand that dependent informants, or indirect informants, provide less information than their independent or direct counterparts.

This setup provides a unique way to examine how children learn from multiple informants, and the types of biases they might have. Numerous studies (e.g., Aboody et al., 2022; Bernard et al., 2015; Einav, 2014; Sampaio et al., 2019; Wilks et al., 2015) have found that children under the age of 6 often, but not always, endorse a majority of informants over a minority. In many cases, agreeing with a majority can actually be rational: if each informant

provides an independent source of information, a majority is supported by a greater amount of evidence than a corresponding minority. This means that it can be hard to assess whether or not children are biased towards majorities *above and beyond* what is rational.

To disentangle the amount of information a majority provides from the number of demonstrators in the majority, we need to examine cases where we know that the majority of informants provide less information than the minority, so that it is irrational to follow the majority based on their information quality. Here, we focus on the case where the indirect group has more informants than the direct group but, because they give their testimony based on hearsay, they nonetheless provide less information than the direct group. In this case, children might normatively determine that they should endorse the choice of the minority with direct information. Alternatively, if children have a conformity bias in these tasks, children may conclude that, even if a larger group of indirect informants provides less total information than a smaller group of direct informants, the mere presence of a majority is informative in its own right.

Therefore, to assess whether children have a conformity bias in these tasks, we need to identify cases where children should normatively endorse a smaller direct group of informants over a larger indirect group, and make predictions for the *extent* of that tendency. By developing several scenarios where a rational learner should endorse groups to greater or lesser degrees, we can evaluate children's behavior in greater detail than just whether or not they endorse a majority, providing a more precise measure of the degree to which children deviate from normative inference.

Next, we present a normative model which analyses how a rational learner should make decisions based on indirect and direct testimony, without a conformity bias. We then compare

the predictions of this model to children's performance, and to the predictions of a conformity biased model, in a series of new experiments (Experiments 4-6) to assess whether children conform to the majority more than is rational. The model we build follows from previous Bayesian models of learning from testimony (e.g., Buchsbaum et al. 2012, Shafto et al., 2012, Whalen et al. 2018) where learners use Bayes' rule to perform inference over multiple hypotheses and select a behavior. Bayes' rule indicates that the probability that a hypothesis, h, is true, given some data, such as informant testimony t, is proportional to the probability of the testimony given the hypothesis times the prior probability of the hypothesis, or

$$p(h|t) \propto p(t|h)p(h). \quad (1)$$

$p(h|t)$ is the posterior probability, $p(t|h)$ is the likelihood, and $p(h)$ is the prior probability of the hypothesis.

In general, hypotheses represent claims about the world, and the data represents observations. In this case, the hypotheses represent beliefs about which item is in which box, and the data are the testimonies given by the informants. Unlike previous models of learning from testimony, here the informants make claims about their preferences rather than factual claims. To capture differing preferences, we assume that a proportion λ of the population prefers one item, while the rest prefer the other. We call the item preferred by the proportion λ the *target* item.

Source Knowledge Model

Under our experimental setup (modeled on Experiments 1–3), the learner evaluates two hypotheses, h_d, that the target item is in the box endorsed by the direct group, and h_i, that the target item is in the box endorsed by the indirect group. The probability of each hypothesis can then be calculated via Bayes' rule. For example, evaluating the hypothesis that the box chosen by the direct group is preferred yields the posterior probability

$$p(h_d|\boldsymbol{t_d}, \boldsymbol{t_i}) \propto p(\boldsymbol{t_d}|h_d)p(\boldsymbol{t_i}|h_d)p(h_d) \quad (2)$$

where $\boldsymbol{t_i} = (t_{i1}, \ldots, t_{in})$ refers to the testimony of the indirect group, and $\boldsymbol{t_d} = (t_{d1}, \ldots, t_{dn})$ refers to the testimony of the direct group. In other words, the posterior probability of the hypothesis that the box chosen by the direct group is preferred rests on both the prior probability of the target item's location—which we assume to be equal for both locations, $p(h_i) = p(h_d)$, and the likelihood of the testimony provided by the two groups if the preferred item really is in the box endorsed by the direct group.

Direct Evidence. The likelihood term, $p(\boldsymbol{t_d}|h_d)p(\boldsymbol{t_i}|h_d)$—the probability of observing a particular set of testimony given the hypothesis that the target item is in the box preferred by a direct group—depends critically on how the learner assumes informants generate their testimony. For simplicity, we assume that direct informants observe the contents of the boxes accurately, and report their preferences accurately. This means that the probability that an informant with direct evidence endorses the box containing the target item is simply $p\left(t_{dj}\middle|h_{t_j}\right) = \lambda$, where h_{tj} refers to the hypothesis that the target item is in the box endorsed by direct informant *j*'s testimony, t_{dj}. The direct informants do not hear any other information, so their testimony is not based on the testimony of others, which means that $p(\boldsymbol{t_d}|h_i)$ is just the product of the likelihood of the individual testimonies,

$$p(\boldsymbol{t_d}|h_d) = \prod_{j=1}^{n} p(t_{dj}|h_d). \quad (3)$$

Indirect Evidence. In the case where informants receive indirect evidence in the form of whispers, their testimony is based solely on the information provided by other informants. Future informants must use that information to first infer which item is in which box, and then endorse a box according to their own preference. However, if the learner is also told each informant's

preference, as in our experiments, then they are already aware of all the information that each indirect informant had to make their decision, so that subsequent informants provide no new information. According to the Source Knowledge model, a learner should therefore disregard all but the first informant in the chain, so that

$$p(\boldsymbol{t_i}|h_d) = p(t_{i1}|h_d), \quad (4)$$

where $p(\boldsymbol{t_i}|h_d)$ is the likelihood of the indirect group's testimony as a whole.

Incorporating Preference

Finally, we assume that the learner, like the informants, also has a preference, preferring the target item with probability λ. To choose a box, learners first infer the probability that each box holds the target item, and then use their preference to determine which box they select. The probability that the learner chooses the box endorsed by the direct informants is just the probability that the box contains the learner's preferred item given the testimony (i.e., we assume that some proportion of learners, 1 – λ, do not prefer the target item, so they will choose the box they believe *not* to contain the target item). Taken together, a learner operating under the assumptions of this model should pick the direct informants' box with probability,

$$\lambda \cdot p(h_d|\boldsymbol{t_d}, \boldsymbol{t_i}) + (1-\lambda) \cdot (1 - p(h_d|\boldsymbol{t_d}, \boldsymbol{t_i})), \quad (5)$$

where $p(h_d|\boldsymbol{t_d}, \boldsymbol{t_i})$, is the posterior probability of the target item being in the box endorsed by the direct informants.

Conformity-Biased Model

Alternatively, if children's choices are biased towards conforming to majorities, then they may consider the mere existence of additional informants as being evidence to support the position of these informants, even if their evidence was gathered indirectly. We model conformity bias as treating indirect evidence identically to direct evidence, with the likelihood of the indirect

group's testimony being calculated identically to the likelihood of the direct group's testimony, i.e., by computing the product of the likelihoods of the individual testimonies (Equation 3).

Mixed Model

Lastly, it is possible that children are uncertain about whether to use a source-knowledge based strategy or a conformity-biased strategy when group sizes are unequal. In such a situation, rather than solely weighing the number of independent sources providing information about a preference, or solely relying on the number of informants endorsing an option, children might implement a mixture of these strategies, weighing both the number of independent sources and the absolute number of informants in their reasoning, either within or across individuals. Models including a mixture of strategies have predicted children's learning across a number of social and causal learning scenarios (e.g. Lieder et al., 2015; Nussenbaum et al., 2020); similarly, children might engage in a mixture of strategies to evaluate the testimony they receive. We model this possibility by introducing a parameter, ω, that represents the proportion of the weight placed on the choices predicted by the Source Knowledge model compared to the Conformity-Biased model. At $\omega = 1$, this model is equivalent to that of the Source Knowledge model, while at $\omega = 0$, it is equivalent to the Conformity-Biased model. For simplicity, and to avoid adding another free model parameter, we use a fixed value of $\omega = 0.5$ to reflect an equal mixture of the two models (i.e., averaging their results) throughout the main text (see Supplementary Material for alternate analysis).

Modeling Direct and Indirect Informants

Since in our experiments the two groups of informants always endorse opposite boxes, and since $p(h_i) = p(h_d)$, it is possible to further simplify the posterior probability into a closed form

$$p(h_d|\boldsymbol{t_d}, \boldsymbol{t_i}) = \frac{\lambda^j(1-\lambda)^k}{\lambda^j(1-\lambda)^k+(1-\lambda)^j\lambda^k} \quad (6)$$

where *j* and *k* are the numbers of informants considered to have independent access to the boxes' contents in each group.

For example, under the assumptions of the Source Knowledge model, the number of direct informants with independent access to the boxes' contents in Experiments 1–3 is equal to the number of direct informants, so $j = 4$ (Experiments 1 and 2) or 3 (Experiment 3), while the number of indirect informants with independent access to the boxes' contents is just the first indirect informant, so $k = 1$ (Experiments 1 and 3). In Experiment 2, indirect informants' knowledge is ambiguous, but as there is no evidence that any of the indirect group has obtained knowledge about the boxes' contents, we set $k = 0$.

However, as mentioned previously, a conformity-biased learner may treat all informants as having information of equivalent quality. Thus, in the Conformity-biased Model, both *j* and *k* equal the number of direct and indirect informants, respectively. Since the size of the direct and indirect groups is equivalent in Experiments 1–3, $j = k = 4$ in Experiments 1 and 2 and $j = k = 3$ in Experiment 3 for the Conformity-biased model.

Model Predictions

We can now use our models to make *a priori* predictions about how a rational learner might make inferences when group size and information quality are at odds, and compare these predictions to children's performance, to see whether children endorse a majority above and

beyond the information they provide (i.e., exhibit a majority bias). Experiment 1 provides a baseline case with equally sized direct and indirect groups, where we can be sure that a majority bias could not be playing a role in children's inferences. We therefore first use this experiment to estimate the value of the preference parameter, and then, given that value, make predictions for cases where group sizes differ. Fitting the preference parameter to children's choices in Experiment 1 yields a value of $\lambda = 0.75$, a relatively high value consistent with our intuition that children believe preferences for items such as food and toys are broadly shared.

Model predictions, along with experimental results are presented in Figure 4. Using the best fitting parameter value of $\lambda = 0.75$ for Experiments 1–3 we confirm that, when group sizes are equal, children do not behave consistently with the Conformity-biased model (log likelihood = -94.41), which predicts that children will perform at chance between the direct and indirect groups. Instead, their behavior more closely matches the predictions of the Source Knowledge model (log likelihood = -87.69), choosing the group with a greater amount of direct sources in Experiments 1 through 3, $\chi^2(1) = 13.43$, $p < 0.001$.

In addition to the four direct and four indirect informants (4 vs. 4) case of Experiments 1 and 2 and the three direct and indirect informants (3 vs. 3) case of Experiment 3, we also examined the cases of three direct vs five indirect informants (3 vs. 5), four direct vs six indirect informants (4 vs. 6), and on direct vs seven indirect informants (1 vs. 7). We chose these ratios in order to vary the relative size of the majority while keeping either the number of direct informants (4 vs. 6) or the overall number of informants (3 vs. 5 and 1 vs. 7) consistent with Experiment 1. We examine the model predictions for each case in more detail, below.

In the case of 4 vs. 6 and 3 vs. 5, we find that the Source Knowledge model continues to predict that individuals will be more likely to endorse the direct informants, though at a slightly

lower rate for the 3 vs. 5 case than in the 4 vs. 4 condition. This drop is primarily due to there being one less direct informant in the direct group. Conversely, the Conformity-biased model predicts that children should favor the indirect majority, because the additional two informants are treated as providing additional information.

The case of 1 vs. 7 deviates substantially from the previous cases. In this case, the learner is presented with one informant with direct knowledge in the direct group, and one informant with direct knowledge in the indirect group (the first indirect informant). The Source Knowledge model predicts that a learner should ignore the remaining indirect informants and be at chance between the two groups, while the Conformity-biased model predicts a stronger tendency to endorse the indirect majority.

The three additional cases outlined above provide a range of predictions to investigate whether children have a bias to conform to the majority's behavior above what is rational when group sizes are unequal. Given children's success in Experiments 1–3, it is possible that preschool-age children might successfully use source knowledge when it is available, and understand that the mere presence of a majority does not provide additional evidence, if members of the majority acquired their endorsements from indirect knowledge. If so, children's behavior should closely reflect the predictions of the *a priori* Source Knowledge model. On the other hand, it is possible that children only use source knowledge when group sizes are equal, and may switch to a conformist strategy when these sizes are unequal; in this case, children's choices could be more similar to the predictions of the Conformity-biased model.

Finally, if children do engage in a mixture of strategies, children's choices when the source knowledge and majority conflict would look different from both possibilities. In this case, children would be predicted to choose at chance between the two groups in the 3 vs. 5 and the 4

vs. 6 conditions. However, in the 1 vs. 7 conditions, children would be predicted to choose the indirect group significantly more often than chance, but do so less strongly than the Conformity-biased model. This results in predictions for children's performance across experiments that differentiate the three possible models (Figure 4).

Experiment 4: Source versus consensus

Experiments 1–3 find that children are sensitive to both the dependency between informants, and to the source of informants' knowledge—whether their testimony is based on hearsay. In both cases, children seem to understand that dependent informants, or indirect informants, provide less information than their independent or direct counterparts. We therefore use both of these cues to informant quality in Experiment 4, to examine how children respond to cases where the indirect group has more informants than the direct group but, because they give their testimony based on hearsay, they provide less information than the direct group.

Experiment 4 examines how children respond when presented with an option endorsed by a majority of indirect informants versus an option endorsed by a minority of direct informants. To directly compare children's performance to the predictions of our model, we examined the cases of three direct vs five indirect informants (3 vs. 5), four direct vs six indirect informants (4 vs. 6), and one direct vs seven indirect informants (1 vs. 7). As we anticipated that the presence of unequal groups would be more challenging for children, we increased the sample size collected per condition to 32. Due to recruitment difficulties, one condition (4 vs. 6) had a smaller sample size; a replication of this condition with a full sample of 32 children was conducted in Experiment 5.

Methods

Participants. Participants in the 3 vs. 5 condition were 31 3- to 5-year-old children (mean age = 55 months; range = 44 to 62 months; 18 female, 13 male) recruited from a large US metropolitan area, and were tested in the lab, their preschools and local museums. Three additional children were tested but excluded due to experimenter error. Participants in the 4 vs. 6 condition were 24 3- to 5-year-old children (mean age = 52 months; range = 42 to 61 months; 16 female, 8 male) recruited from a large US metropolitan area, and were tested in the lab, their preschools and local museums. Three additional children were tested but were excluded due to experimenter error. Participants in the 1 vs 7 condition were 32 3- to 5-year-old children (mean age = 56 months; range = 43 to 70 months; 10 female, 22 male) recruited from a large Canadian metropolitan area, and were tested in the lab, their preschools and local museums. 3 additional children were tested but excluded due to experimenter error.

Materials and Procedure. Materials were the same as in Experiment 1, except for the addition of two dolls in in the 4 vs. 6 condition, and the use of stickers (as in Experiment 2) instead of snacks in in the 1 vs. 7 condition. The procedure for Experiment 4 was identical to Experiment 1, except with the number of informants in the direct and indirect groups varying appropriately.

Results

Results for Experiment 4 are summarized in Table 1. For each trial, children received a 1 if they chose the box endorsed by the direct informants, and a 0 if they chose the box endorsed by the indirect informants.

3 vs. 5 Condition. Children were at chance in choosing between the box endorsed by the direct group and the box endorsed by the indirect majority, $B = 0.23$, $SE = 0.32$, 95% CI = [-0.44, 1.01], $OR = 1.26$, $z = 0.72$, $p = .473$. There was no significant difference in responses between

the first and the second trial, $B = 0.16$, $SE = 0.28$, 95% CI = [-0.39, 0.74], $OR = 1.17$, $z = -0.55$, $p = .580$, or for the two trial types, Fisher exact test, $OR = 0.36$, $p = .07$.

4 vs. 6 Condition. Children were at chance in choosing between the box endorsed by the direct group and the box endorsed by the indirect majority, $B = -0.73$, $SE = 0.76$, 95% CI = [-3.98, 0.74], $OR = 0.48$, z = -0.95, $p = .340$. There was no significant difference in responses between the first and the second trial, $B = 0.69$, $SE = 0.47$, 95% CI = [-0.14, 1.92], $OR = 1.99$, $z = 1.46$, $p = .144$, or for the two trial types, Fisher exact test, $OR = 0.71$, $p = .77$.

1 vs. 7 Condition. Children chose the box endorsed by the direct majority significantly below chance, $B = -0.63$, $SE = 0.30$, 95% CI = [-1.36, -0.09], $OR = 0.53$, $z = -2.12$, $p = .034$. There was no significant difference in responses between the first and the second trial, $B = 0.50$, $SE = 0.28$, 95% CI = [-0.03, 1.11], $OR = 1.65$, $z = 1.76$, $p = .079$, or for the two trial types, Fisher exact test, $OR = 1$, $p = 1$.

Discussion

Given children's sensitivity to informants' knowledge source in Experiments 1–3, we predicted that children might continue to use source knowledge when it is available, choosing the item endorsed by the higher quality direct informants, even when source knowledge and group size are in conflict. Instead, we found that unlike children's responses in Experiment 1, and in contrast to the predictions of the normative Source Knowledge model, children in the 3 vs. 5 and 4 vs. 6 conditions of Experiment 4 were at chance when choosing between the boxes endorsed by the direct and indirect groups. However, children in the 1 vs. 7 condition children preferentially endorsed the majority indirect group over the minority direct group, even though the number of informants with direct visual access in both groups was the same.

Across all three conditions of these tasks, children's degree of endorsement of the direct group was lower than the predictions of the Source Knowledge model (Figure 4), which predicts that an idealized learner should endorse the smaller group with a larger number of primary sources in the 3 vs. 5 and 4 vs. 6 conditions, and choose at chance in the 1 vs. 7 condition, where both groups have an equal number of primary sources. These results suggest that a consensus may have the power to diminish children's tendency to endorse testimony from groups with a larger number of primary sources, but it does not shift children's judgments entirely—they do not simply endorse the majority's choice whenever a numerical majority exists, as predicted by the Conformity-biased model.

However, non-significant results can be hard to interpret. On the one hand, these results could be the result of a sensitivity to knowledge source combined with an over-weighting of majority information (e.g., a conformity bias), leading to children being torn between the option endorsed by the majority and the one endorsed by higher quality informants. Although 4- and 5-year-old children can reliably discriminate numerical quantities with a ratio of 1.5 to 1 (Halberda & Feigenson, 2008; Odic et al., 2013), suggesting they should distinguish the size of the groups even in the most challenging group comparison (4 vs. 6), it is also possible that children may find the additional task of interpreting the relative quantity of information provided by the groups more difficult in this case, and thus choose randomly when presented with groups of informants of unequal size, as has been suggested elsewhere (Morgan et al., 2015).

Experiment 5: Replication of 4 vs. 6 Condition

To ensure that the additional complexity of the unequal group sizes did not make Experiment 4 too hard for children to follow, we replicated the 4 vs. 6 condition of Experiment 4

with the addition of a number of control questions evaluating children's understanding of the relative size of the two groups, their memory for the groups' endorsements, and their understanding of the information passed between members of the indirect group.

Methods

Participants. Participants were 32 3- to 5-year-old children (mean age = 58 months; range = 47 to 70 months) recruited from a large Canadian metropolitan area, and were tested in the lab, and local museums. 10 additional children were tested but excluded due to experimenter error, and 3 children did not complete the experiment.

Materials and Procedure. Materials were the same as in the 4 vs 6 condition of Experiment 4, except for the use of stickers (as in Experiment 2) instead of snacks. The procedure for this experiment was identical to the 4 vs. 6 condition of Experiment 4, up until the end of the second trial. Following the child's second trial choice, they were asked three control questions (1) "Do you remember, which people were whispering?" (2) "When the people were whispering, what were they saying?" (3) "Which group has more people?". The dolls remained in front of the boxes they had endorsed throughout these questions.

Results and Discussion

Children were at chance in choosing between the box endorsed by the direct group and the box endorsed by the indirect majority, B = -0.46, SE = 0.33, 95% CI = [-1.31, 0.12], OR = 0.63, z = -1.37, p = .172. There was no significant difference in responses between the first and the second trial, B = -0.15, SE = 0.28, 95% CI = [-0.76, 0.45], OR = 0.86, z = -0.55, p = .579, or for the two trial types, Fisher exact test, OR = 0.88, p = 1. When asked which informants were whispering, 25 of 31 children correctly chose the indirect group (1 child did not choose a group), $p < .001$, exact binomial test. When asked what the informants were whispering, 21 of 25

children gave an answer indicating that the informants were whispering which box contained the better sticker or toy (e.g., "the toy in this box is better"), while 4 children gave a neutral descriptive answer (e.g., "about the sticker"); an additional 7 children did not provide an answer. Finally, 29 of 32 children correctly identified the indirect group as having more people, $p < .001$, exact binomial test.

As with Experiment 4, children were not significantly more likely to choose either the direct or indirect groups. Most children believed that the indirect informants were whispering to each other which toy they liked better. Further, all but three children indicated that the indirect group was larger, consistent with the finding that by 3 years of age, children can consistently distinguish populations differing by a ratio of 1.5 or greater (Odic et al., 2013), even if they have not yet acquired exact numerosity. Together, these findings suggest that poor task understanding did not likely contribute to the non-significant results observed in Experiment 4.

Age Effects

Given previous findings that 3-year-olds sometimes have more difficulty than 4- and 5-year-olds in evaluating informant accuracy (e.g., see Corriveau et al., 2009; Koenig & Harris, 2005), we also examined whether there was an overall effect of age on children's choices—i.e., whether older children were more likely to choose the box endorsed by the direct informants—when the data from all studies was taken together. We found an effect of experimental condition, $\chi^2(6) = 22.33$, $p = .001$, such that children chose the direct box to differing degrees in different studies, but no main effect of age on the degree to which children chose the direct box when considering all of the experiments, $\chi^2(1) = 2.54$, $p = .11$, and no significant interaction between age and experiment in the degree to which children chose the direct box, $\chi^2(6) = 3.15$, $p = .79$, suggesting that age effects are not driving the differences in performance across experiments.

Model Comparison

Comparing children's performance across Experiments 4 and 5 to the Source Knowledge and Conformity-biased models, children were substantially less likely to choose the minority direct group than the predictions of the Source Knowledge model, but also more likely to do so than the Conformity-biased model predicted. If children are considering both source knowledge and the size of a group when making their decisions, their results may reflect a balancing or weighting of both pieces of evidence.

In fact, a simple equal mixture of these two models captured children's performance across the uneven group size conditions very accurately, and significantly better than either the source knowledge or conformity biased model individually. This outcome suggests that while children may use source knowledge alone when there are no conflicting cues in the form of uneven groups, children may use a mix of these strategies when source knowledge cues and group size are in conflict.

As a result, using the source knowledge model (fit to Experiment 1) to predict children's performance in Experiments 1–3, and the mixture of source knowledge and conformity to predict their performance in Experiments 4 and 5 (log likelihood -250.91) provides a significantly better fit to children's performance than making predictions using just source knowledge (log likelihood -279.04, $\chi^2(1) = 56.27$, $p < 0.001$) or just conformity bias (log likelihood -268.90, $\chi^2(1) = 35.97$, $p < 0.001$).

Alternatively, it is possible that children might be able to use source knowledge when neither group is larger, but become conformists in the presence of a majority. To represent this, we tested an alternative model in which children use source knowledge when group size is equal, but rely on the conformity-biased model alone when group sizes are unequal. We found, once

again, that the combination of source knowledge and a mixture of source knowledge and conformity outperformed a model that relied on source knowledge when groups were equally sized and conformity alone when group sizes were unequal (log likelihood -259.55, $\chi^2(1) = 17.28$, $p < 0.001$).

These findings suggest that at least as a group, children could be employing both conformity-biased and source knowledge-based strategies. This supports the interpretation that, even when group sizes are unequal, children might continue to take source knowledge into account, but that they may also treat the mere presence of a majority as an independent source of evidence for the majority's choice, even when the source of each member of the majority's opinion is already known. We will return to a discussion of why this might be the case in the General Discussion.

Experiment 6: Adults

In Experiments 4 and 5, children appeared to be swayed by the size of the indirect majority, suggesting that they believe the size of the majority may provide additional information or an additional cue to informant quality despite the fact that the minority had equal or better information quality. As discussed in the introduction, adults' inferences about the independence and dependence of sources are compatible with a normative model on some tasks (Whalen et al., 2018), but other recent studies have found that adults are sometimes vulnerable to the effect of a "false consensus" (e.g., Yousif et al., 2019). Nevertheless, adults more heavily weight the independence of a source when it is made clear that informants are relying on the independent data they obtained to make their claims (Alister et al., 2022; Desai et al., 2022), and they may find distinguishing between the source quality of the direct and indirect groups less challenging

than children. Here, we therefore examine adults' choices on a task similar to those conducted with children in Experiments 1 and 4.

Methods

Participants. Participants were 241 adult US residents, recruited through Amazon Mechanical Turk (MTurk) and paid $0.50 for their time. Participants were required to have over a 95% lifetime acceptance rate on MTurk. Participants were randomly assigned to one of four conditions: 60 participants to a four direct vs. four indirect condition, 60 participants to a four direct vs. six indirect condition, 60 participants to a three direct vs. five indirect condition, and 61 participants to a one direct vs. seven indirect condition.

Materials. The experiment was an online survey administered using Qualtrics survey software, with custom animations created using JavaScript. The informants were a set of 10 distinct cartoon clip art characters (5 male, 5 female). There were also two pairs of cartoon boxes that differed only in color: a red and blue pair, which participants were told contained games, and a green and yellow pair, which participants were told contained snacks.

Procedure. The procedure closely matched that used with children in Experiments 1 and 4, with the clip art characters replacing the dolls that children saw. Like children, adults each participated in two trials, a snack trial and a game trial, with the order of trials counterbalanced. Adults saw two boxes on opposite sides of the screen. For the direct group, each member of the group was shown one at a time. A character appeared on the screen, then moved to each box while the cartoon text "*Looks inside box*" flashed above the character's head. Then, the character stood by one box and said, "I think the [game/snack] in the [blue] box is better!" For the indirect group, the first member was shown looking inside the boxes, declaring his or her opinion, and moving to stand next to another indirect group member who appeared on screen.

The cartoon text "*whisper*" appeared above both their heads. The second doll then moved to stand by one box, and gave their opinion, "[S]he said the [game/snack] in the [blue] box was better, so I think the [game/snack] in the [blue] box is better". This process repeated for the remaining characters.

After all characters gave opinions, participants were shown an image with each group of characters placed under the box they endorsed, with a reminder that this was the box each character thought was better. Participants were then asked to "Please select the box with the [game/snack] that you would like to try". Group order and side/color of box endorsed by the direct group were counterbalanced. In game trials, the red box always appeared on the left, and in snack trials the green box always appeared on the left. For each participant, characters' group assignments were randomized.

Results and Discussion

Results are shown in Table 2 and Figure 5. Overall, in the 4 vs 4, 3 vs 5 and 4 vs. 6 conditions, adults chose the box endorsed by the direct group significantly more than chance (all $z \geq 3.71$, all $OR \geq 7.84$, all $p < .001$) . In the 1 vs. 7 condition, adults were at chance for choosing the majority or minority box, $B = 0.42$, $SE = 0.33$, 95% CI = [-0.22, 1.06], $OR = 1.52$, $z = 1.30$, $p = .19$. Across experiments, we find that adults choose the option endorsed by the direct group, even when the indirect informants are the majority. In the 1 vs. 7 condition, where there is one direct informant endorsing each option, adults ignore the additional indirect informants and are at chance between the two options.

In comparing adult and child performance, a 2 (age group: adults or children) x 4 (Experiment: 1, 4-6) ANOVA revealed a main effect of age group; adults' and children's responses differed significantly, $\chi^2(1) = 61.18$, $p < .001$. There was also a significant interaction

of experiment with age group, $\chi^2(3) = 9.35$, $p = .025$. Planned comparisons between age groups suggest that this effect was driven by differences in the uneven group size conditions. Adults were significantly more likely than children to choose the box chosen by the direct group in the 4 vs. 6 condition, $B = 1.95$, $SE = 0.31$, 95% CI = [1.35, 2.55], $OR = 7.02$, $z = 6.33$, $p < .001$, the 3 vs 5 condition, $B = 1.36$, $SE = 0.35$, 95% CI = [0.67, 2.04], $z = 3.87$, $p < .001$, and the 1 vs. 7 condition, $B = 0.84$, $SE = 0.32$, 95% CI = [0.22, 1.47], $z = 2.65$, $p = .008$, but there was no difference between age groups in the 4 vs. 4 condition, $B = 0.63$, $SE = 0.42$, 95% CI = [-0.19, 1.45], $z = 1.50$, $p = .13$.

In contrast to children, we find a very close qualitative and quantitative fit between adult's responses and the source knowledge model (Figure 5; log likelihood -262.18), indicating that adults, unlike children, balance the number of informants and the quality of their knowledge source. In contrast, the conformity-biased model was a comparatively poor fit for adults' responses (log likelihood -443.41, $\chi^2(1) = 362.47$, $p < 0.001$). The best fitting preference value for adults is approximately $\lambda = 0.84$. This value is similar to the value found for children, and suggests that the differences in children and adults' inferences are not due to differing assumptions about the extent to which preferences are shared.

Overall, the Source Knowledge model accurately captures adult, but not child, performance across conditions, while a simple additive mixture of source knowledge and conformity bias accurately captures children's performance in the uneven group size conditions, providing further support for the finding that children are making a different kind of inference than adults, one that takes into account source of knowledge, but also comparatively favors the majority. In addition, the source knowledge model does accurately capture children's judgments in the equal group size conditions, supporting the interpretation that children are using source

knowledge appropriately in those cases, suggesting that the difference between children and adults is not due to an inability to monitor and track multiple informants' information quality.

General Discussion

These studies provide the first empirical evidence that as young as three years old, children can weigh multiple informants' opinions using the quality of their knowledge source to assess which option they themselves should choose. They are also the first to demonstrate that adults can normatively balance the size of a majority with the number of primary sources they provide, and that they can do so in the domain of preferences. In contrast, when a larger number of total informants was contrasted with a smaller number of informants with greater direct knowledge, children's choices across conditions suggested a majority bias, though informed by source knowledge.

We find that with equal numbers of informant endorsements (Experiment 1), children favored a box recommended by informants with direct perceptual access over informants who had received knowledge indirectly (hearsay from other informants). This remained true even if the indirect informants gained their knowledge independently of each other, each getting their hearsay from a different source (Experiment 2). Additionally, when children encountered informants who all received only hearsay (Experiment 3), they favored opinions from informants who received hearsay from several independent sources over informants who received hearsay from the same source.

When the box endorsed by a larger number of total informants and the box endorsed by a larger number of the informants with direct knowledge were pitted against one another, children were either at chance in choosing between the boxes (Experiment 4: 3 vs. 5 and 4 vs. 6 conditions) or selected the box endorsed by the indirect majority (Experiment 4: 1 vs. 7

condition). From the perspective of tracking endorsements based on direct knowledge, additional informants in the indirect group provide limited new information, since their endorsements are statistically dependent on the endorsement made by the initial informant with direct knowledge. The Source Knowledge model predictions indicate that an idealized learner, who believes that the informants only have access to the information presented in the experiments, should choose the box endorsed by a larger number of the informants with direct knowledge, not the majority of total informants. Across conditions, adults consistently endorsed the direct group, and behaved in accordance with the predictions of a normative model sensitive to source knowledge. The fact that children did not could indicate that they treat the presence of a majority as additional independent evidence beyond the evidence provided by its individual members, in line with findings that children consider majority opinions and behaviors an important source of information (e.g., Bernard et al., 2015; Corriveau et al., 2009; Haun et al., 2012; Pham & Buchsbaum, 2020).

However, we also find that children do not simply conform whenever a majority is present and were not well captured by a purely conformity-biased model. Instead, children's inferences are best captured by a simple mixture of the Conformity-biased model and the Source Knowledge model, suggesting while children's inferences were influenced by the size of the majority group, they were also sensitive to the source of the informants' knowledge, as work in children's selective trust in informants has found (e.g., Aboody et al., 2022; Birch et al., 2008; Bridgers et al., 2016; Jaswal & Neely, 2006; Koenig & Harris, 2005; Ronfard & Corriveau, 2016).

This study bridges these areas of research, demonstrating that children consider both the degree of first-hand information and the number of endorsements when determining how they

884 should integrate conflicting social information. These findings may help reconcile previous
885 mixed results as to whether children have a conformity bias, by suggesting that both information
886 quality and majority size contribute to children's inferences. For instance, comparing a majority
887 that is unsuccessful on the current task with a dissenter who succeeds (Wilks et al., 2015), may
888 create a greater quality disparity than comparing a previously unsuccessful minority to a majority
889 with no known history (Burdett et al., 2016; Sampaio et al., 2019), leading children to favor the
890 minority in the former but not the latter case. Similarly, a disparity in expertise on the task at
891 hand (e.g., Wilks et al., 2015; Einav, 2014) may be a stronger cue to differing quality than a
892 history of accuracy versus inaccuracy on earlier tasks (Bernard et al., 2015). Finally, if children
893 perceive both majority size and direct perceptual access as independent cues to quality, as our
894 results suggest, then they will be less likely to conform to a lower quality majority if that
895 majority is also smaller (e.g., Schillaci & Kelemen, 2014, majority of 2 vs. minority of 1) and
896 will display reduced conformity biases when the majority's claims lack epistemic strength (Kim
897 & Spelke, 2020). In all of these cases, young children might consistently overweight information
898 provided by majorities—i.e., they may show a majority *bias*—but, because children are sensitive
899 to other characteristics such as information quality and the extent of the majority, this will only
900 sometimes lead children to display a tendency to endorse the majority at a level greater than
901 50%.

902 By testing children's and adults' endorsements on several tasks that differ systematically
903 in the number of informants in each group and the quality of the groups' information, this set of
904 experiments provides evidence that preschool age children weigh information source and
905 selective trust differently than adults. Since our model accurately captures adult, but not child,
906 performance, it provides further support for the finding that children are making a different kind

of inference than adults, one that comparatively favors the majority. There are several possibilities for why children may place additional value on majority information relative to adults. One possibility is that children's tendency to overweight majority information is the result of their emerging theory of mind development. To understand that the presence of a majority does not provide additional evidence if the sources of each member's beliefs are not independent from each other, children need to understand that informants' beliefs are generated from the evidence they observe. While children as young as three years old display an awareness that the claims of individuals with perceptual access to information are more reliable (e.g., Pillow, 1989; Robinson et al., 2011; Butler et al., 2018), children's perspective-taking abilities are still developing considerably from ages 4 to 8 (Frick et al., 2014). Thus, although we found no significant age effects in our experiments, correlating an explicit measure of theory of mind abilities (e.g., theory of mind scale, Wellman & Liu, 2004; theory of mind sub-test NEPSY-II, Korkman et al., 2007), with children's tendency to conform to a majority with indirect information might prove fruitful in future work.

Another possibility is that younger children are more motivated to affiliate themselves with a majority than older children and adults (Bernard et al. 2015; but see e.g., Morgan et al. 2015 for an opposite finding of an increasing tendency to conform with age), so that, unlike adults, children were independently motivated by source knowledge and a desire to affiliate with the larger group. This affiliation may also reflect a perception that informants whose initial endorsement is relied upon as hearsay by other informants are more prestigious, and thus more important to affiliate with. For example, 3- and 4-year-old children show a prestige bias in their learning, attending to demonstrators who are preferentially imitated by bystanders rather than demonstrators whose behavior was ignored (Chudek et al., 2012). Thus, some children in our

experiment may have considered the agreement by informants in the indirect group to be a signal to the quality of the knowledge of the initial informant.

Preferences, in particular, could be a domain in which children might perceive the presence of a majority as intrinsically meaningful and thus disproportionately attend to the number of endorsements. This would be consistent with other findings that children show a greater propensity to endorse majorities in conventional domains (e.g., what to label an object), relative to domains such as causal learning, where asocial learning is possible (Pham & Buchsbaum, 2020). However, it is important to note that such patterns would only be expected in domains where children perceive preferences to be broadly shared; in domains where one's own preferences are expected to be more idiosyncratic and difficult to predict based on the preferences of others, or domains in which the child anticipates or has experienced having a distinct preference from the majority (e.g., food preferences, Repacholi & Gopnik, 1997), we would not anticipate a similar pattern of results. Investigation of when and why this tendency shifts, such that adults on our task endorsed the groups of informants that had the greater number of primary sources, much like they have been shown to do in factual domains (Aboody et al., 2022), and did not consider the endorsements in our task that were based on hearsay as informative as children, despite the fact that adults can also exhibit similar prestige-based learning biases (e.g., Atkisson et al., 2012), could deepen our understanding of the belief system underlying children's selective trust.

Children must often reason about their likely preferences (e.g., with food) before having significant personal experience with the preferred item in question; thus, while objects are not literally hidden, many of the relevant characteristics that might inform a child's preferences, such as the food's taste, are not available to the child before making a choice to try something. On the

other hand, many of children's preferences are learned in an environment in which children already have existing familiar and favoured (as well as disfavoured) items. In these circumstances, children's reasoning about testimony and the degree to which they adjust their beliefs about their own likely preferences are likely to differ in more ways than simply the majority size and the information quality. For example, children may already have a strong belief that they will not enjoy e.g., broccoli more than goldfish, even if they receive testimony from a majority that supports broccoli. Likewise, children may use testimony to make inferences about the informants themselves; much as they make inferences about the reliability of informants based on accuracy (Pasquini et al., 2007; Corriveau et al., 2009), children may reduce trust or reliance on the testimony of informants who endorse an option that is already known to be dispreferred by the child. Thus, an open question in preference learning is how children integrate their own knowledge and pre-existing preferences, as well as new testimony from informants to evaluate both their potential preferences and the quality and relevance of the information they are receiving from informants.

Further, while we find that children as a group are split about midway between a conformity-biased strategy and an arguably more appropriate source knowledge strategy, this does not tell us which mechanism individual children are using to make their choices. This could either be implemented at a between-child level, with some children consistently using a source knowledge strategy, and others using a conformity-biased strategy, or at a within-child level, where the child chooses which strategy to use on each trial, or where the child takes both source knowledge and majority size into account on every trial. For example, in the 4 vs. 6 condition of Experiment 4, children were significantly more likely to consistently choose either the indirect majority or the direct minority on both trials (see Supplementary Material). This may suggest

that individual children are using different strategies in the most ambiguous situations, a finding consistent with some previous work (Burdett et al., 2016). This may also align with findings that adults sometimes exhibit a conformity bias (e.g., Yousif et al., 2019; Desai et al., 2022; Alister et al., 2022) and other times, as in this study and others (e.g., Whalen et al., 2018), do not. If individual young children and adults use comparable strategies when faced with ambiguous situations, but young children perceive more situations to be ambiguous, this could explain why younger children exhibit a conformity bias on our task relative to adults. Extending these findings with older children would help to clarify the nature of this developmental trend.

Extensions of the type of mixture model we apply can be very useful for understanding individual performance when learners have multiple decision-making strategies to choose from (see e.g., Nussenbaum et al., 2020, for an example of children and adults using a mixture of causal hypothesis testing strategies, and Lieder et al., 2015, for an example of children using a mixture of social learning strategies). Future work could use a similar modeling approach to examine the potential for individual differences in more detail.

The presence of a conformity bias in children in situations where it is not present in adults may have striking implications for the development of human culture. Many cultural traits, including language and social conventions, are learned at an early age. Formal models suggest that a conformity bias may lead to the stability of such traits over time (Boyd & Richerson, 1985; Henrich & Boyd, 1998), and recent work has demonstrated a U-shaped trend in a bias toward the majority across 9 countries, with both younger children and adolescents showing a greater frequency of majority-copying behavior (Sibilsky et al., 2022). If children demonstrate a conformity bias at an early age, it may allow them to quickly learn in-group norms, but may allow neutrally beneficial or even detrimental behaviors to persist in the

population. Given that a behavior learned from a majority in childhood may persist through adulthood, a bias towards conformity in children that stems from incorrectly estimating the quality and amount of information provided by each informant would lead to systematic changes in the adoption and maintenance of cultural traits through a population. Though the results from this study do not directly address the transmission of social norms based on informant reliability, future work can explore this issue. Additionally, while some work suggests that children's endorsement of a majority may be particularly strong in conventional domains, in which there is not necessarily a "ground truth" but rather a social convention, relative to domains such as causality where asocial learning is possible (Pham & Buchsbaum, 2020), research into adults suggests that under at least some circumstances, adults can exhibit similar conformity biases in factual domains (e.g., Desai et al., 2022; Yousif et al., 2022), though at other times their behavior appears to be normative (e.g., Whalen et al., 2018). This makes it particularly striking that adults showed no conformity bias in this study. Thus, future work should examine whether the conformity bias that we demonstrate in this set of studies about children's endorsements based on informants' stated preferences extends to other domains, such as facts, and whether variability in adults' tendency to conform is related to the conventionality of the domain, or perhaps to other factors such as the ease of evaluating the informants' sources of knowledge.

Although a conformity bias may allow mildly detrimental behaviors to persist in a population, it may yield benefits. In some cases (e.g., language), the benefit a behavior derives is based solely on the extent to which other individuals in the population also use that behavior. An early-appearing conformity bias may allow children to quickly adopt seemingly arbitrary behaviors (e.g. social norms and customs) which can confer indirect benefits through social bonding and acceptance (e.g., Clegg & Legare, 2016; Evans et al., 2021; Kenward, Karlsson &

Persson, 2011; Schmidt, Rakoczy, & Tomasello, 2011). Moreover, as young children are learning about a wide variety of demonstrators, overestimating adults' knowledge may still be more beneficial than harmful; adults have a wider knowledge base than children, and can draw on this knowledge to provide more accurate information.

Whether picking which snack to eat or deciding which toy to buy, children and adults rely on information they receive from other people every day. Together these experiments go beyond asking whether or not people have a conformity bias, and explore children's and adults' sensitivity to multiple informants' knowledge source when reconciling conflicting endorsements. We find that preschool-age children demonstrate an emerging ability to consider several types of information—directness of knowledge and consensus—when assessing which testimony to use when determining what they themselves are likely to prefer. Despite this, children also exhibit a conformity bias and endorse a majority's opinion disproportionately, even if their testimony is rooted in less first-hand knowledge. Together, these findings may have implications not only for understanding children's social learning but also for understanding the cultural transmission and maintenance of preferences and behaviors.

References

Alister, M., Perfors, A., & Ransom, K. (2022). Source independence affects argument persuasiveness when the relevance is clear. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (pp. 2767–2773).

Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, *87*(5), 847–862.

Atkisson, C., O'Brien, M. J., & Mesoudi, A. (2012). Adult learners in a novel environment use prestige-biased social learning. *Evolutionary Psychology*, *10*(3), 147470491201000309.

Bernard, S., Harris, P., Terrier, N., & Clément, F. (2015). Children weigh the number of informants and perceptual uncertainty when identifying objects. *Journal of Experimental Child Psychology*, *136*, 70–81. https://doi.org/10.1016/j.jecp.2015.03.009

Bernard, S., Proust, J., & Clément, F. (2015). Four-to Six-Year-Old Children's Sensitivity to Reliability Versus Consensus in the Endorsement of Object Labels. *Child Development*, *86*(4), 1112–1124. https://doi.org/10.1111/cdev.12366

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *The Journal of Political Economy*, *100*(5), 1026.

Birch, L. L. (1999). Development of food preferences. *Annual Review of Nutrition*, *19*(1), 41-62.

Birch, S. A., Vauthier, S. A., & Bloom, P. (2008). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition*, *107*, 1018–1034. https://doi.org/10.1016/j.cognition.2007.12.008

Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. University of Chicago Press.

Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, *108*, 10918–10925. https://doi.org/10.1073/pnas.1100290108

Bridgers, S., Buchsbaum, D., Seiver, E., Griffiths, T. L., & Gopnik, A. (2016). Children's causal inferences from conflicting testimony and observations. *Developmental Psychology*, *52*(1), 9–18. https://doi.org/10.1037/a0039830

Burdett, E. R., Lucas, A. J., Buchsbaum, D., McGuigan, N., Wood, L. A., & Whiten, A. (2016). Do children copy an expert or a majority? Examining selective learning in instrumental and normative contexts. *PloS One*, *11*(10), 0164698. https://doi.org/10.1371/journal.pone.0164698

Butler, L. P., Gibbs, H. M., & Tavassolie, N. S. (2020). Children's developing understanding that even reliable sources need to verify their claims. *Cognitive Development*, *54*, 100871. https://doi.org/10.1016/j.cogdev.2020.100871

Butler, L. P., Schmidt, M. F. H., Tavassolie, N. S., & Gibbs, H. M. (2018). Children's evaluation of verified and unverified claims. *Journal of Experimental Child Psychology*, *176*, 73–83. https://doi.org/10.1016/j.jecp.2018.07.007

Chudek, M., Heller, S., Birch, S., & Henrich, J. (2012). Prestige-biased cultural learning: bystander's differential attention to potential models influences children's learning. *Evolution and Human Behavior*, *33*(1), 46-56. https://doi.org/10.1016/j.evolhumbehav.2011.05.005

Clegg, J. M., & Legare, C. H. (2016). Instrumental and conventional interpretations of behavior are associated with distinct outcomes in early childhood. *Child Development*, *87*(2), 527–542. https://doi.org/10.1111/cdev.12472

Coenen, A., Rehder, B., & Gureckis, T. M. (2015). Strategies to intervene on causal systems are adaptively selected. *Cognitive Psychology*, *79*, 102–133. https://doi.org/10.1016/j.cogpsych.2015.02.004

Corriveau, K. H., Fusaro, M., & Harris, P. L. (2009). Going With the Flow: Preschoolers Prefer Nondissenters as Informants. *Psychological Science*, *20*(3), 372–377. https://doi.org/10.1111/j.1467-9280.2009.02291.x

Corriveau, K. H., Pickard, K., & Harris, P. L. (2011). Preschoolers trust particular informants when learning new names and new morphological forms: Preschoolers trust novel morphology. *British Journal of Developmental Psychology*, *29*(1), 46–63. https://doi.org/10.1348/2044-835X.002009

Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, *13*(4), 148–153. https://doi.org/10.1016/j.tics.2009.01.005

Desai, S. C., Xie, B., & Hayes, B. K. (2022). Getting to the source of the illusion of consensus. *Cognition*, *223*, 105023. https://doi.org/10.1016/j.cognition.2022.105023

Einav, S. (2014). Does the majority always know best? Young children's flexible trust in majority opinion. *PLOS ONE*, *9*(8), 104585. https://doi.org/10.1371/journal/pone.0104585.

Einav, S., & Robinson, E. J. (2011). When Being Right Is Not Enough Four-Year-Olds Distinguish Knowledgeable Informants From Merely Accurate Informants. *Psychological Science*. https://doi.org/10.1177/0956797611416998

Evans, C. L., Burdett, E. R. R., Murray, K., & Carpenter, M. (2021). When does it pay to follow the crowd? Children optimize imitation of causally irrelevant actions performed by a

majority. *Journal of Experimental Child Psychology*, *212*, 105229. https://doi.org/10.1016/j.jecp.2021.105229

Frick, A., Möhring, W., & Newcombe, N. S. (2014). Picturing perspectives: Development of perspective-taking abilities in 4- to 8-year-olds. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.00386

Halberda, J., & Feigenson, L. (2008). Developmental change in the acuity of the" Number Sense": The Approximate Number System in 3-, 4-, 5-, and 6-year-olds and adults. *Developmental Psychology*, *44*(5), 1457. https://doi.org/10.1037/a0012682

Hargreaves, D. J., North, A. C., & Tarrant, M. (2015). How and why do musical preferences change in childhood and adolescence. In G. McPherson (Ed.), *The Child as Musician: A Handbook of Musical Development* (pp. 303-322).

Harris, P. L., Koenig, M. A., Corriveau, K. H., & Jaswal, V. K. (2018). Cognitive Foundations of Learning from Testimony. *Annual Review of Psychology*, *69*(1), 251–273. https://doi.org/10.1146/annurev-psych-122216-011710

Haun, D. B., Rekers, Y., & Tomasello, M. (2012). Majority-biased transmission in chimpanzees and human children, but not orangutans. *Current Biology*, *22*(8), 727–731. https://doi.org/10.1016/j.cub.2012.03.006

Haun, D., & Tomasello, M. (2011). Conformity to peer pressure in preschool children. *Child Development*, *82*(6), 1759–1767. https://doi.org/10.1111/j.1467-8624.2011.01666.x

Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, *19*(4), 215–241. https://doi.org/10.1016/S1090-5138(98)00018-X

Hoppitt, W., & Laland, K. N. (2013). *Social learning: An introduction to mechanisms, methods, and models*. Princeton University Press.

Jaswal, V. K., & Neely, L. A. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science*, *17*(9), 757–759. https://doi.org/10.1111/j.1467-9280.2006.01778.x

Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., & Jones, P. L. (2018). Social Learning Strategies: Bridge-Building between Fields. *Trends in Cognitive Sciences*, *22*(7), 651–665. https://doi.org/10.1016/j.tics.2018.04.003

Kenward, B., Karlsson, M., & Persson, J. (2011). Over-imitation is better explained by norm learning than by distorted causal learning. *Proceedings of the Royal Society of London B: Biological Sciences*, *278*(1709), 1239–1246. https://doi.org/10.1098/rspb.2010.1399

Kim, S., & Spelke, E. S. (2020). Learning from multiple informants: Children's response to epistemic bases for consensus judgments. *Journal of Experimental Child Psychology*, *192*, 104759. https://doi.org/10.1016/j.jecp.2019.104759

Koenig, M. A., & Harris, P. L. (2005). Preschoolers Mistrust Ignorant and Inaccurate Speakers. *Child Development*, *76*(6), 1261–1277. https://doi.org/10.1111/j.1467-8624.2005.00849.x

Koenig, M. A., & Sabbagh, M. A. (2013). Selective social learning: New perspectives on learning from others. *Developmental Psychology*, *49*(3), 399–403. https://doi.org/10.1037/a0031619

Korkman, M., Kirk, U., & Kemp, S. (2007). *NEPSY-II: clinical and interpretive manual*. NCS Pearson.

Lamont, A., & Crich, J. (2022). Where do our music preferences come from? Family influences on music across childhood, adolescence and early adulthood. *Journal of Popular Music Education*, *6*(1), 25-43. https://doi.org/10.1386/jpme_00073_1

Landrum, A. R., Eaves, B. S., & Shafto, P. (2015). Learning to trust and trusting to learn: A theoretical framework. *Trends in Cognitive Sciences*, *19*(3), 109–111. https://doi.org/10.1016/j.tics.2014.12.007

Lieder, F., Sim, Z., Hu, J., Griffiths, T. L., & Xu, F. (2015). Children and adults differ in their strategies for social learning. In D. Noelle, R. Dale, A. Warlaumont, J. Yoshimi, T. Matlock, C. Jennings, & P. Maglio (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 1367–1373). Cognitive Science Society.

Lucas, C. G., Griffiths, T. L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T., Markson, L. & Hu, J. (2014). The child as econometrician: A rational model of preference understanding in children. *PLoS ONE*, *9*(3), e92160. https://doi.org/10.1371/journal.pone.0092160

Mercier, H., & Miton, H. (2019). Utilizing simple cues to informational dependency. *Evolution and Human Behavior*, *40*(3), 301-314. https://doi.org/10.1016/j.evolhumbehav.2019.01.001

Mills, C. M. (2013). Knowing When to Doubt: Developing a Critical Stance When Learning From Others. *Developmental Psychology*, *49*(3), 10.1037/a0029500. https://doi.org/10.1037/a0029500

Morgan, T. J. H., Laland, K. N., & Harris, P. L. (2015). The development of adaptive conformity in young children: Effects of uncertainty and consensus. *Developmental Science*, *18*(4), 511–524. https://doi.org/10.1111/desc.12231

Morgan, T., Rendell, L., Ehn, M., Hoppitt, W., & Laland, K. (2012). The evolutionary basis of human social learning. *Proceedings of the Royal Society of London B: Biological Sciences*, *279*, 653–662. https://doi.org/10.1098/rspb.2011.1172

Nussenbaum, K., Cohen, A. O., Davis, Z. J., Halpern, D. J., Gureckis, T. M., & Hartley, C. A. (2020). Causal Information-Seeking Strategies Change Across Childhood and Adolescence. *Cognitive Science*, *44*(9), e12888. https://doi.org/10.1111/cogs.12888

Odic, D., Libertus, M. E., Feigenson, L., & Halberda, J. (2013). Developmental change in the acuity of approximate number and area representations. *Developmental Psychology*, *49*(6), 1103–1112. https://doi.org/10.1037/a0029472

O'Neill, D. K., Astington, J. W., & Flavell, J. H. (1992). Young Children's Understanding of the Role That Sensory Experiences Play in Knowledge Acquisition. *Child Development*, *63*(2), 474–490. https://doi.org/10.2307/1131493

Palmquist, C. M., & Jaswal, V. K. (2012). Preschoolers Expect Pointers (Even Ignorant Ones) to Be Knowledgeable. *Psychological Science*, *23*(3), 230–231. https://doi.org/10.1177/0956797611427043

Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, 43(5), 1216. https://doi.org/10.1037/0012-1649.43.5.1216

Pillow, B. H. (1989). Early understanding of perception as a source of knowledge. *Journal of Experimental Child Psychology*, *47*(1), 116–129. https://doi.org/10.1016/0022-0965(89)90066-0

Povinelli, D. J., & deBlois, S. (1992). Young children's (Homo sapiens) understanding of knowledge formation in themselves and others. *Journal of Comparative Psychology*, *106(3)*, 228. https://doi.org/10.1037/0735-7036.106.3.228

Robinson, E. J., & Einav, S. (2014). *Trust and skepticism: Children's selective learning from testimony*. Psychology Press.

Robinson, E. J., Haigh, S. N., & Nurmsoo, E. (2008). Children's working understanding of knowledge sources: Confidence in knowledge gained from testimony. *Cognitive Development*, *23*(1), 105–118. https://doi.org/10.1016/j.cogdev.2007.05.001

Ronfard, S., & Corriveau, K. H. (2016). Teaching and preschoolers' ability to infer knowl- edge from mistakes. *Journal of Experimental Child Psychology, 150*, 87–98. https: //doi.org/10.1016/j.jecp.2016.05.006

Rosner, B. (2015). *Fundamentals of Biostatistics*. Nelson Education.

Sampaio, L. R., Harris, P. L., & Barros, M. L. (2019). Children's selective trust: When a group majority is confronted with past accuracy. *British Journal of Developmental Psychology*, *37*(4), 571–584. https://doi.org/10.1111/bjdp.12297

Schillaci, R. S., & Kelemen, D. (2014). Children's conformity when acquiring novel conventions: The case of artifacts. *Journal of Cognition and Development*, *15*(4), 569–583. https://doi.org/10.1080/15248372.2013.784973

Schmidt, M. F., Rakoczy, H., & Tomasello, M. (2011). Young children attribute normativity to novel actions without pedagogy or normative language. *Developmental Science*, *14*(3), 530–539. https://doi.org/10.1111/j.1467-7687.2010.01000.x

Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*, *15*(3), 436–447. https://doi.org/10.1111/j.1467-7687.2012.01135.x

Sibilsky, A., Colleran, H., McElreath, R., & Haun, D. B. M. (2022). Expanding the understanding of majority-bias in children's social learning. *Scientific Reports*, *12*(1), 6723. https://doi.org/10.1038/s41598-022-10576-3

Sobel, D. M., & Kushnir, T. (2013). Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, *120*(4), 779–797. https://doi.org/10.1037/a0034191

Sodian, B., & Wimmer, H. (1987). Children's Understanding of Inference as a Source of Knowledge. *Child Development*, *58*(2), 424–433. https://doi.org/10.2307/1130519

Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Harvard University Press.

Vélez, N., Wu, Y., & Gweon, H. (2018). Consistent but not diagnostic: Preschoolers' intuitions about shared preferences within social groups. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (pp. 2621–2626).

Ventura, A. K., & Worobey, J. (2013). Early influences on the development of food preferences. *Current Biology*, *23*(9), R401-R408.

Walker, M. B., & Andrade, M. G. (1996). Conformity in the Asch task as a function of age. *The Journal of Social Psychology*, *136*(3), 367–372. https://doi.org/10.1080/00224545.1996.9714014

Wellman, H. M., & Liu, D. (2004). Scaling of theory-of-mind tasks. *Child Development*, *75*(2), 523–541. https://doi.org/10.1111/j.1467-8624.2004.00691.x

Whalen, A., Griffiths, T. L., & Buchsbaum, D. (2018). Sensitivity to Shared Information in Social Learning. *Cognitive Science*, *42*(1), 168–187. https://doi.org/10.1111/cogs.12485

Whiten, A. (2019). Chapter Two - Conformity and over-imitation: An integrative review of variant forms of hyper-reliance on social learning. In M. Naguib, L. Barrett, S. D. Healy, J. Podos, L. W. Simmons, & M. Zuk (Eds.), *Advances in the Study of Behavior* (Vol. 51, pp. 31–75). Academic Press. https://doi.org/10.1016/bs.asb.2018.12.003

Wilks, M., Collier-Baker, E., & Nielsen, M. (2015). Preschool children favor copying a successful individual over an unsuccessful group. *Developmental Science*, *18*(6), 1014–1024. https://doi.org/10.1111/desc.12274

Yousif, S. R., Aboody, R., & Keil, F. C. (2019). The Illusion of Consensus: A Failure to Distinguish Between True and False Consensus. *Psychological Science*, *30*(8), 1195–1204. https://doi.org/10.1177/0956797619856844

Table 1

Summary of children's performance in Experiments 1–5.

Number of children choosing the direct group's box	0	1	2
Experiment 1 (4 vs. 4)	2	8	12
Experiment 2 (all independent)	1	13	10
Experiment 3 (all indirect)	3	10	11
Experiment 4 (3 vs. 5)	8	12	11
Experiment 4 (4 vs. 6)	11	6	7
Experiment 4 (1 vs. 7)	13	15	4
Experiment 5 (4 vs. 6)	13	12	7

Table 2

*Children's and Adults' choices in Experiments 1, 4, and 6 compared. * indicates a significant result, p < .05, ** indicates p < .01, *** indicates p < .001, via generalized linear mixed model (GLMM).*

Experiment (Children/Adults)	Children's average score for choosing direct group, out of 2 (standard error)	Adults' average score for choosing direct group, out of 2 (standard error)
Experiment 1/6 (4 vs. 4)	1.45** (0.14)	1.67*** (0.07)
Experiment 4/6 (3 vs. 5)	1.10 (0.14)	1.65*** (0.07)
Experiment 4/6 (4 vs. 6)	0.83 (0.18)	1.65*** (0.07)
Experiment 4/6 (1 vs. 7)	0.72* (0.12)	1.13 (0.10)

Figure 1

Experiment 1 Design

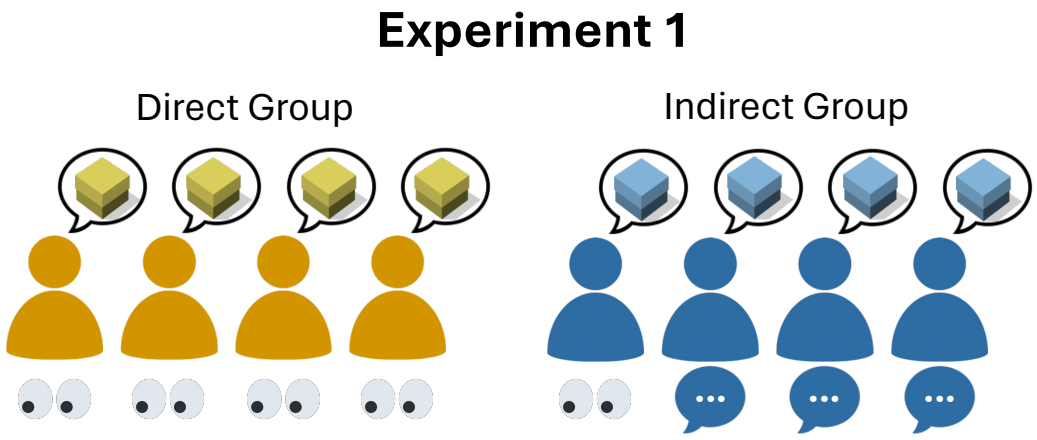

Note. Informant cues for Experiment 1. Children watched as two different groups of informants gathered data directly (eyes) or indirectly (speech bubbles), before endorsing one of the two boxes. Members of the direct group (yellow figures) each independently observed the contents of the boxes before endorsing one of the two boxes (yellow box). In the indirect group (blue figures), one informant directly observed the boxes, and then endorsed the other of the two boxes (blue box). Subsequently, informants in this group would whisper information to the next informant in the chain (speech bubbles), who would also endorse the other of the two boxes.

Figure 2

Experiment 2 Design

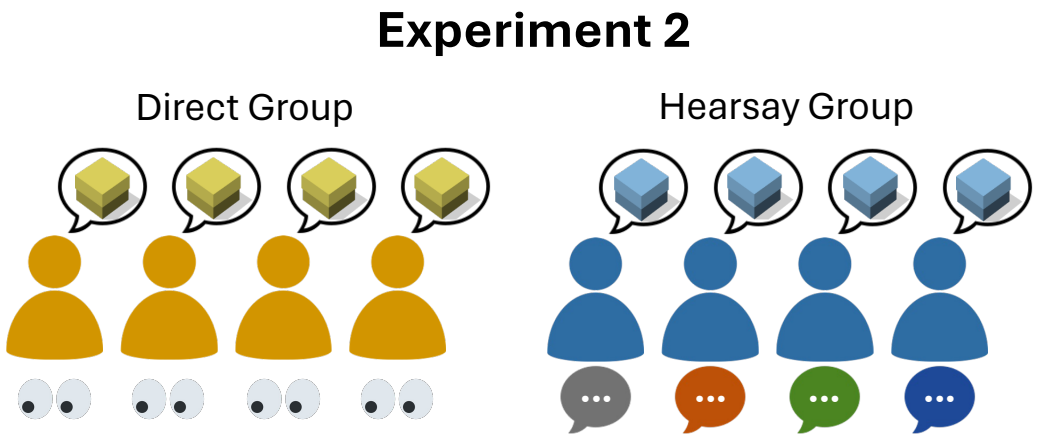

Note. Informant cues for Experiment 2. Children watched as two different groups of informants gathered data directly (eyes) or indirectly (speech bubbles), before endorsing one of the two boxes. Members of the direct group (yellow figures) each independently observed the contents of the boxes before endorsing one of the two boxes (yellow box). In the indirect group (blue figures), each informant reported their source as a different friend (speech bubbles in various colors), and then endorsed the other of the two boxes (blue box).

Figure 3

Experiment 3 Design

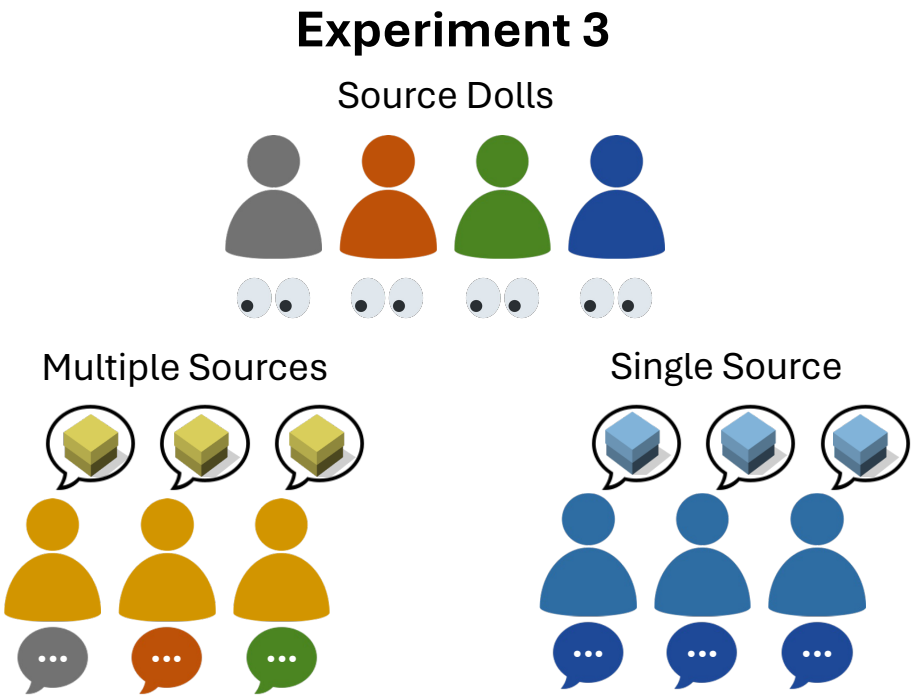

Note. Informant cues for Experiment 3. Children watched as four dolls (figures in various colors) observed the contents of the boxes, then whispered to the informant dolls (yellow and blue figures). Members of the multiple-sources group (yellow) each heard a different source doll whispering (speech bubbles in various colors), and then each endorsed one of the two boxes (yellow box), while members of the single-source group (blue) received information from the same source doll (blue speech bubbles), and then endorsed the other of the two boxes (blue box).

Figure 4

Model Predictions and Children's Choices for Experiments 1-4

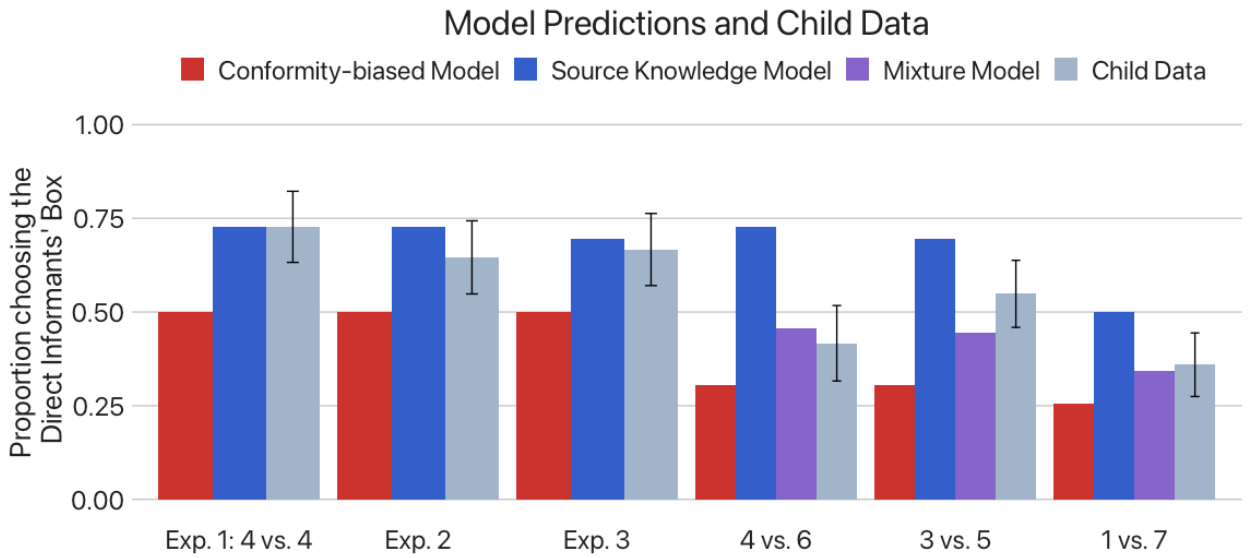

Note. The preference parameter was fit to child performance in Experiment 1.

Figure 5

Model Predictions and Adults' Choices in Experiment 6

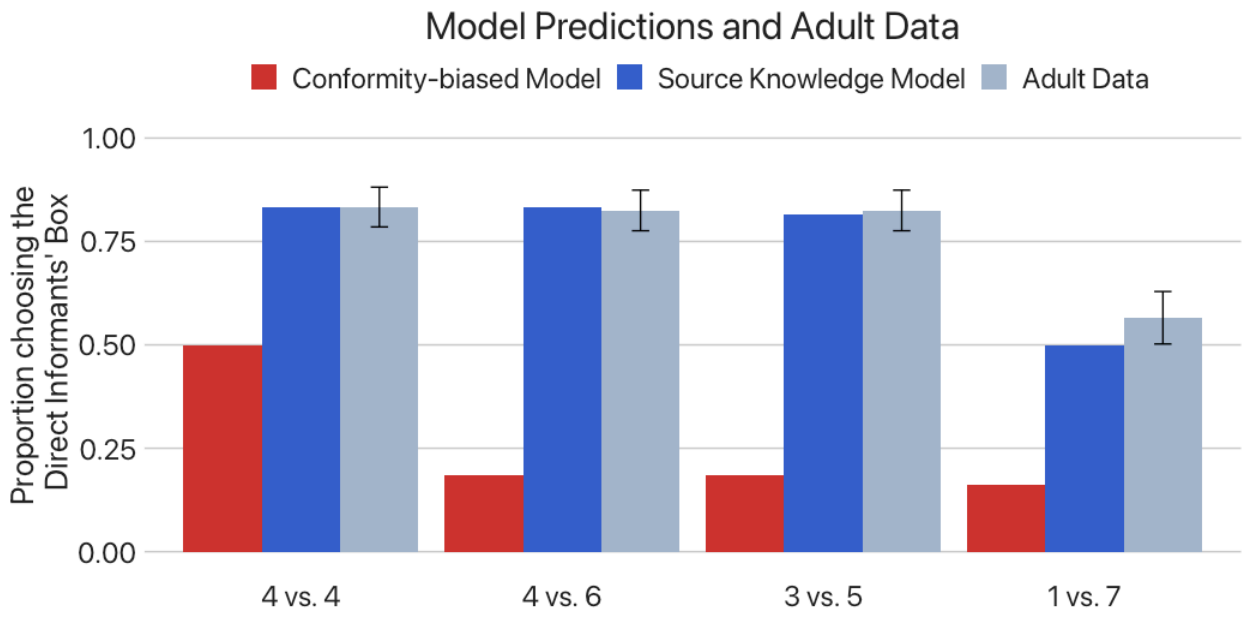

Note. The preference parameter was fit to adult performance in the 4 vs. 4 condition.