# “He only changed his answer because they shouted at him”: children use affective cues to distinguish between genuine and forced consensus

**Emory Richardson (**emory.richardson@yale.edu**) & Frank Keil (**frank.keil@yale.edu**)**
Yale Department of Psychology, 2 Hillhouse Ave. New Haven, CT 06511 USA

**Abstract**

Learning frequently forces us to rely on the good judgment and epistemic vigilance of sources with no more firsthand knowledge of a topic than ourselves, but who may have more second or third-hand knowledge. Yet, being forced to rely on their judgment doesn’t prevent us from evaluating their judgment: one might trust information because it was passed on to you by someone whose epistemic vigilance you trust, but reject it from someone whom you believe lacks good judgment. We present two experiments suggest that by integrating affective cues like anger and surprise along with perceptual access and consensus, children infer what others believe and what the correct answer to a question is. We discuss implications for consensus-based social learning strategies.

**Keywords:** social learning; emotion; development; group processes

## Introduction

Across the hall, you can see a group of five policy makers making a decision; though you can’t hear them, they’re visible through a window. One has just disagreed with the rest of the group, who are clearly angry with what the dissenter says. The four shout the dissenter down, and he concedes. Did this person change their mind, or simply give in? By contrast, suppose the four had shown surprise at the dissenter’s opinion, and the dissenter conceded after discussing the matter with the group. Did this person genuinely change their mind? While the use of intimidation seems to make belief change suspect in the angry group, there is less reason to doubt that the dissenter was genuinely convinced in the surprised group. But would your doubt or belief in the dissenter’s conversion change your own judgment about which answer is more likely to be accurate? While the angry group’s intimidation may be concerning, there was a 4-to-1 consensus even without the dissenter’s forced concession. You may decide that the affective signals are not enough to cause you to trust the dissenter *instead of* such a strong consensus, at least without more information. Is the same true of children? Here, we ask how children weigh affective cues against other criteria for evaluating collective judgment, and compare their judgments to adult weightings.

The ability to recognize and evaluate judgments from multiple sources at once is a crucial skill for agents capable of learning from others. One common social learning strategy is majority rule: across cultures (Boehm, 1996; van Leeuwen et al., 2018), age groups (Haun & Tomasello, 2011; Mannes, 2009), and species (Haun, van Leeuwen & Edelson, 2013; Claidière & Whiten, 2012), learners are more likely to trust majority opinion than minority opinion in a variety of contexts — and stronger majorities elicit greater conformity (Morgan, Laland, & Harris, 2015; Morgan, Rendell, Ehn, Hoppitt & Laland, 2012). Trust in majority rule may often be adaptive: majority rule outperforms other decision rules in simulation studies (Hastie & Kameda, 2005), and is provably more reliable than alternative decision rules under certain conditions (Condorcet, 1785, List & Goodin, 2001). One of these conditions is that individuals’ judgments are statistically independent of each other. When judgments are *not* independent, majority rule is less robustly accurate because it is then regulated by the accuracy of the most influential individuals — and while an accurate influencer may improve collective judgment, an *in*accurate influencer will distort it (Hahn, von Sydow, & Merdes, 2018; Becker, Almaatouq, & Horvat, 2020).

One way to avoid the pernicious effects of social influence on consensus accuracy is to discount the judgment of anyone who does not have firsthand, independent knowledge of the topic. For instance, even 4-year-olds will discount a majority judgment if *every* individual member of the majority is less informed than the dissenter (Einav, 2014; Kim & Spelke, 2020; Hu et al., 2015). However, older children seem to selectively ignore the independence criterion in other cases. For instance, only by age 6 do children trust a single informant relaying the testimony of three eyewitnesses over three informants relaying the testimony of a single eyewitness (Aboody et al., 2021). Similarly, when contrasting two informants who conspicuously look at a third informant’s answer to a general trivia question with three informants who independently give an alternative answer, only by age 8-9 do children favor the independent consensus, with 6-year-olds at chance and 5-year-olds actually preferring the *non*-independent consensus (Einav, 2018).

However, the protracted developmental trajectory of discounting these “false” consensuses does not necessarily indicate a flaw in humans’ capacity for social learning. Indeed, a strategy that limits learners strictly to firsthand judgments would be of little use in the real world. For instance, a single scientific paper may require expertise from an entire laboratory; and even a comparatively simple task like reconstructing a perceptual event may require testimony from multiple partially-informed eyewitnesses. By accepting the conclusions of groups in which no single individual can account for every aspect of their collective work, we endorse — at least tacitly — the process by which individual knowledge was aggregated into a single collective judgment. Moreover, as a population’s reliance on

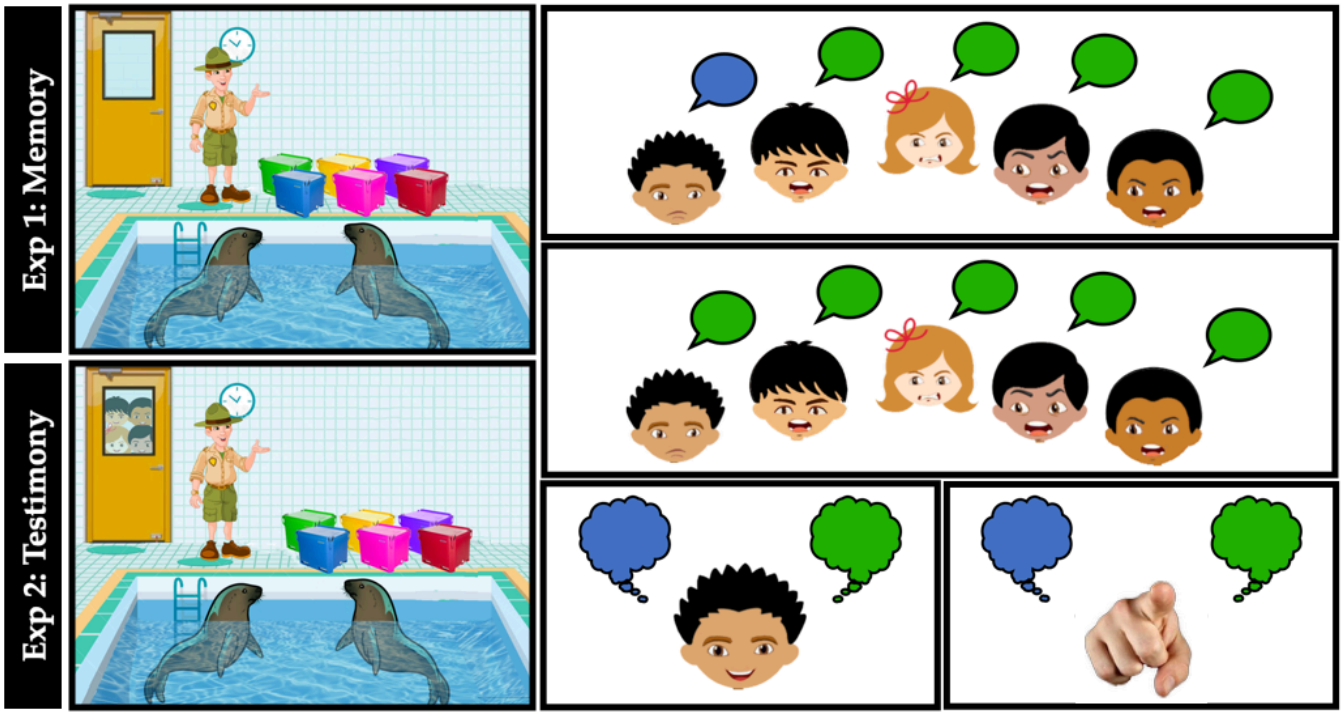


**Figure 1.** Procedure slides from Anger trial in Experiments 1 and 2. The teacher asked the students to look in the boxes to remember which contained an animals' favorite food. In each trial, the dissenter changed his answer after the other 4 either shouted at (Anger) or talked with (Surprise) him. Participants were then asked what the dissenter *really* thought (Belief), and asked which answer they thought was correct (Accuracy). In Exp. 1, all 5 students had been taught the correct answer "last month"; in Exp. 2, the 4-person consensus (but not the dissenter) had seen the answer by chance when they arrived early to class.

social learning increases, the most influential individuals may be many degrees removed from firsthand knowledge. Teachers, textbook authors, and journalists have considerable influence over collective opinion not because of their independent knowledge of the topic, but because of their power to amplify information to a large audience. Yet, while learners need to be selective in their choice of sources, a teacher's knowledge of calculus is not less reliable because they learned it from a textbook instead of independently deriving it from first principles. In other words, while reliance on secondhand evaluations of socially transmitted evidence is ubiquitous in the real world, it is not necessarily a flaw in our use of consensus. Rather, it may reflect learners' assumptions about our informants' reliability as epistemic "filters" of information: we expect our informants to vet each others' judgments rather than accepting them at face value. However, recognizing when these filters are likely to fail us may require challenging theory-of-mind reasoning, and it will certainly require learners to evaluate social influences in addition to simply "counting votes". Learning from consensus will require learners to consider how robust their informants' judgments are to social pressures, and how receptive they are to reliable information.

To the extent that other agents can be expected to rationally update their beliefs, their belief changes may be as informative as the degree of consensus itself. Indeed, in multi-armed bandit games (i.e., players choose between slot machines whose payoffs have different expected values) in which the high-payoff option changes after a certain number of rounds, adults copy the most rapidly "trending" decisions of other players to discover the new location of the high-payoff option, independently of their tendency to copy majority behaviors — consistent with buying patterns in real-world stock markets (Toelch, Bruce, Meeus, & Reader, 2010; Barber, Odean, & Zhu, 2009). This is a challenging theory of mind task: in addition to recognizing both consensus and changing trends, learners may need to consider whether the other agents' decisions are rational, given the evidence available to those agents. Yet, even 4-to-7 year-old children predict that agents will update their beliefs rationally. For instance, if an agent is shown new evidence about the proportion of prizes in a box, children expect the agent's belief-updating to integrate their prior beliefs, the new (asocial) evidence, and whether the person taking prizes from the box to present as evidence was doing so randomly or selectively (Magid, Yan, Siegel, Tenenbaum, & Schulz, 2017). However, unlike these experiments in which participants had access to the same asocial evidence that the agents themselves had seen, real world judgments often require us to infer whether or not to trust an informants' judgment even when we have neither the access

nor the competence to evaluate the critical evidence ourselves. Nevertheless, learners may still be able to recognize when informants' judgments are unlikely to reflect a reliable evaluation of the evidence available to them.

Recent work suggests that affective signals like anger and surprise could help learners evaluate social influences on consensus by suggesting whether changes in opinion are *genuine* or *forced* (Richardson & Keil, 2021). An agent who *genuinely* changes their belief not only changes the degree of consensus, but may provide evidence in favor of their new belief and against their prior belief; an agent who is *forced* to conform changes the "consensus", but their conversion is uninformative. Thus, the initial consensus may be more informative than the post-conversion consensus. Adults reason in precisely this manner (Richardson & Keil, 2021). For instance, while participants trusted a 6-to-4 majority even after the 6 shouted at the 4 to become a 9-to-1 majority, they were no *more* confident in the final 9-to-1 consensus than they had been in the initial 6-to-4 consensus; however, they did increase their confidence proportionally if the consensus grew to 9-vs-1 after the 6 simply expressed surprise at the 4 and talked with them. In general, when consensus changed because one faction gained or lost endorsers after expressing surprising at the opposing view, participants' own confidence shifted to reflect the final degree of consensus, suggesting that participants inferred that the converts had genuinely changed their belief during the meeting, and took this belief change to be informative. However, while participants' confidence shifted to reflect the *final* degree of consensus when the angry faction *lost* endorsers, it tracked the *initial* consensus when the angry faction *gained* endorsers by shouting at them, suggesting that participants inferred that the converts had been forced to conform but had not genuinely changed their beliefs.

Here, we use a simplified version of the same paradigm to examine the components of this reasoning in development, asking children to infer both what a convert "really" believes, and which answer they themselves think is correct. Past work suggests that preschoolers have a "benevolence bias": between the ages of 3 and 5, children trust a nice informant over a mean informant even if the mean informant is described as having greater expertise or past accuracy (Johnston, Mills, & Landrum, 2015; Landrum, Mills, & Johnston, 2013). Surprisingly, while these studies suggest that children do become less likely to reject a mean expert between the ages of 3 and 5, one study suggests that benevolence may even override visual access as a cue to perceptual knowledge; children as old as 5 believe a nice informant who tells them the contents of a box without looking inside more than a mean informant who did look (Lane, Wellman, & Gelman, 2012). Importantly, it's unlikely that these children were simply avoiding negative affect; other work suggests that if two informants give conflicting names for unfamiliar objects, but two bystanders smile and nod at one informant's label and frown and shake their heads at the other's, children trust the informant whose answer the bystanders approved of — treating the negative affect and gesture as an evaluation of the informants' accuracy, not as "meanness" (Fusaro & Harris, 2008).

Young children's benevolence bias suggests an early-developing mechanism for adults' asymmetric evaluations of angry and surprised consensuses in Richardson & Keil (2021), as well as a qualitative shift in those evaluations across development. Attributing a "benevolence bias" to other agents could lead both children and adults to infer that an informant who changes their answer after a group shouts at them has not genuinely changed their belief. However, the benevolence bias may lead children and adults to different conclusions about the accuracy of the group consensus. Children may reject even a strong majority if it uses anger to force conformity; yet, while adults may lose some confidence in an angry group, anger is less epistemically significant for adults than cues like majority consensus (Richardson & Keil, 2021). In contrast, if the group expresses surprise at a dissenter, children and adults alike may be more likely to infer that the agent's belief change was genuine and that the consensus judgment is correct than if the group expresses anger.

We test these predictions in children ages 6-9 and adults. Though the benevolence bias is strongest among 3-5 year olds, children's ability to weigh consensus against other relevant cues and their ability to diagnose others' conformist tendencies emerges later (Einav, 2018; Aboody et al., 2021; Cordonier, Nettles, & Rochat, 2018), and reasoning about social influences on third-party beliefs may involved advanced theory-of-mind capacities thought to emerge only in middle childhood.

## General Method

Participants were presented with a slideshow-story in which a protagonist changes their answer about an animal's favorite food about the contents of different colored boxes after talking with four other informants, who express either Anger or Surprise (within-subjects, order counterbalanced) at the protagonist's original answer. After hearing the story, participants made two forced-choice judgments: which answer the *protagonist* thinks is correct (Belief), and which answer *the participant themselves* think is correct (Accuracy). Names, faces, answer colors, and animals were changed for the Anger and Surprise trials.

Importantly, because the participants themselves could not see the contents of the colored boxes, they had to rely on other cues to evaluate the characters' beliefs. The story itself was written to enable minimal script changes across experiments when altering the kind of evidence available to the characters. In each experiment an initial 4-vs-1 consensus became a 5-vs-0 consensus when the protagonist changed their answer after the other students [ *shouted at / talked with* ] him. However, in Experiment 1, the protagonist's informational access was equal to their classmates (they each looked into the boxes and tried to remember what they had been taught the month before), while in Experiment 2, access was unequal (the consensus informants had seen the answer right before class, while the protagonist had not seen the answer). Pre-registrations are

available at the first author's OSF repository (https://osf.io/zsgfu/?view_only=b6b5a14e4940404389504db554f818b0).

## Experiment 1: Equal Access (Memory)

In Experiment 1, all informants were on an equal footing in terms of informational access: each person could consult their own memory of what the zookeeper had taught them, as well as other people's memories after meeting together.

**Participants**. We recruited 40 adults through MTurk, as well as 80 children on Zoom in two age groups (40 age 6-7, M=6.96, SD=.53; 40 age 8-9, M=8.98, SD=.56; 54% girls). An additional 4 children were excluded according to pre-registered exclusion criteria (2 repeated parent interference, 2 experimenter error); these were replaced with new participants.

**Procedure**. First, participants were introduced to a zookeeper who kept the animals' favorite food in different colored boxes so that he knew which kind to give to each animal. Next, participants were introduced to a protagonist and their four classmates. They were told that last month, the zookeeper had taught the students about how he kept the animals' favorite food in different colored boxes, and that today, the teacher was testing the students to see what they had learned. In the Anger trial, the teacher said "*first look in each of the boxes, and let's see what you remember from last month: which box has the sea lion's favorite kind of fish? Think for yourself, and then talk about it with the other people at your table.*" Following this, participants saw that 4 of the students all gave the same answer, but 1 student gave a different answer. Next, participants were told that the consensus group either shouted angrily at the dissenter or were surprised and talked with him, after which he changed his answer (see Fig. 1): "*Jack's classmates all said that the fish in the green box was the sea lions' favorite kind. But Jack said the blue box had the sea lions' favorite kind of fish. But the other students were very* ***[ angry / surprised ]*** *that Jack answered blue. And after the other people all* ***[ shouted at / talked with ]*** *Jack, Jack changed his answer to green too.*" Finally, participants were asked which answer the protagonist really believed, and which answer they themselves thought was correct: "*what if <u>we</u> ask Jack what he <u>really thinks</u> the answer is? Does he really think the blue box has the sea lions' favorite fish, or does he really think the green box does? Okay! And why do you say so? Okay! And my other question is, what do <u>you</u> think the right answer is? Do you think the sea lions' favorite fish is really in the blue box, or is it really in green box?*". Names, faces, animals/answer colors and the description of the emotion were changed for the Anger and Surprise trials, but the script was otherwise identical. The order of the Surprise and Anger trials and the color of the answers was counterbalanced.

**Results.** We fit separate binomial logistic regressions for the Belief questions and Accuracy questions with random intercepts for each participant and Emotion and age in years (centered on the mean of the child sample in order to use the model's intercept to compare the child sample to chance) as fixed effects. This model suggested that while children inferred that the protagonist had *not* genuinely changed his belief after being shouted at ($\beta_{Int}$ = -2.67, *SE* = 0.38, *z* = -7.068, *p* < .001), with no effect of age ($\beta_{Ct\_AgeYears}$ = -0.28, *SE* = 0.41, *z* = 0.487, *p* =.487), they were also significantly more likely to infer that he had genuinely changed his answer when the group expressed surprise and talked with him than when they shouted at him ($\beta_{EmotionSurprise}$ = 2.61, *SE* = 0.42, *z* = 6.163, *p* < .001). However, there was also significant Emotion*Ct_AgeYears interaction, produced by an increasing willingness across development to believe that the protagonist had genuinely changed his belief in the Surprise trial ($\beta_{Emotion*Ct\_AgeYears}$ = 0.88, *SE* = 0.45, *z* = 1.97, *p* = .049). Following our preregistered analysis plan, we computed 95% confidence interval for each AgeGroup and Emotion to compare to chance. In the Anger trial, only 13% of the younger children, 3% of the older children, and 5% of adults (*Younger*: 5 of 40, 95CI: 4.2–26.8; *Older*: 1 of 40, 95CI: 0.0–13.2; *Adults*: 2 of 40, 95CI: 0.0–16.9) inferred that the protagonist had genuinely changed his belief. However, while in the Surprise trial 68% of adults inferred the belief change was genuine (27 of 40, 95CI: 50.9–81.4), only 33% of the younger children (13 of 40, 95CI: 18.6–49.1) did so; older children were split, with 45% believing the change was genuine (18 of 40, 95CI: 29.3–61.5).

Despite this general skepticism among younger children that the protagonist had genuinely changed his belief for either emotion, evidence for either a "benevolence bias" or trust in consensus among children was weaker than expected. Though the endorsements of the consensus judgment as accurate increased across development ($\beta_{Ct\_AgeYears}$ = 1.18, *SE* = 0.29, *z* = .63, *p* =.532), and participants were overall more likely to endorse the consensus on the Surprise trial than the Anger trial ($\beta_{EmotionSurprise}$ = 1.21, *SE* = 0.42, *z* = 2.85, *p* = .004), only Older children endorsed the protagonist more than the consensus on Anger trials (12 of 40, 95CI: 16.6–46.5). While adults did endorse the consensus on both Anger trials (34 of 40, 95CI: 70.2–94.3) and Surprise trials (39 of 40, 95CI: 86.8–99.9), the youngest children did not reject consensus on the Anger trial (16 of 40, 95CI: 24.9–56.7), contrary to prediction. Meanwhile, neither Older nor Younger children's endorsements differed from chance for the Surprise trial (Younger: 19 of 40 endorse consensus, 95CI: 31.5–63.9; Older: 23 of 40 endorse consensus, 95CI: 40.9–73.0).

## Experiment 2: Unequal Access (Testimony)

In Experiment 2, we manipulated the informants' informational access. Rather than telling participants that the informants had *all* learned from the zookeeper the month before and could try to remember what they had learned (as in Experiment 1), the four-informant consensus was said to have seen the answer before class (unbeknownst to the teacher and protagonist, who had not seen the answer), giving the informants unequal informational access.

Though trust in consensus is highly contingent even in adults, children trust informants with direct perceptual access over those without access from early in development.

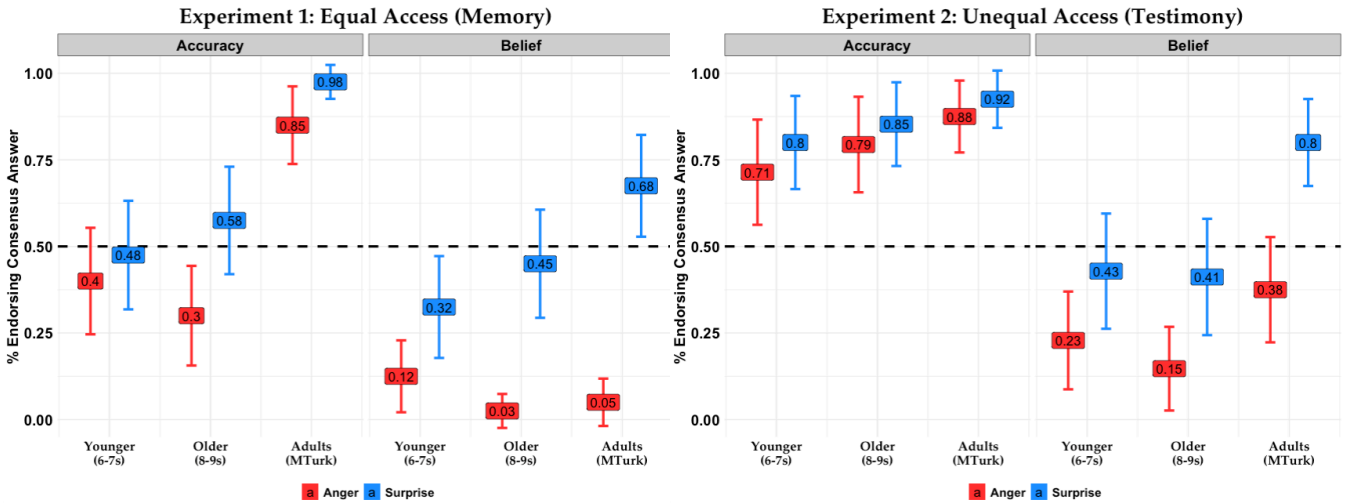


**Figure 2.** Means and 95% CIs for Anger & Surprise trials, binning to compare each age group to chance. Developmental model treated age as a continuous variable.

Thus, in Experiment 2, we expected participants of all ages to endorse the consensus judgment over the protagonist, regardless of affective cues. However, evaluations of third-party beliefs may be more subject to affective cues. In our procedure, participants learned from the experimenter him-/herself (who they had no reason to mistrust) that the consensus informants had seen the answer directly and that the protagonist had not seen the answer. In contrast, the protagonist of the story could only rely on the testimony of their group. Thus, while the protagonist would have little reason to trust a group that mistreated, they would have little reason *not* to trust a group who did not mistreat them. We predicted that even though participants would disregard affective cues in their own Accuracy endorsements, they would infer that the protagonist had *not* genuinely changed their Belief in the Anger trial. In contrast, we expected adults to infer that the belief change was genuine in the Surprise trial, but given children's apparent suspicions in Experiment 1 that the protagonist was "only" conforming, we did not expect children to infer that the belief change was genuine in the Surprise trial at rates greater than chance, only greater than in the Anger trial.

**Participants**. We planned to recruit 40 adults through MTurk, as well as 80 children on Zoom in two age groups. However, because data collection was slowed by COVID, our initial submission reports an analysis of only 69 of the 80 children in the child sample (35 ages 6-7 and 34 ages 8-9).

**Procedure**. The procedure was similar to Experiment 1, with the following changes: first, the slideshow story was expanded to show the protagonist's classmates arriving to the zoo early and peeking in through the window as the zookeeper fed the animals: "*so Jack's classmates saw which color box had the sea lions' favorite fish, and which color boxes had the other animals' favorite fish.*" Second, instead of asking the students to remember what the zookeeper had taught them last month, the teacher announced that before the zookeeper explained what each animal's favorite fish was, she wanted them to look in the boxes and guess. Third, when the characters gave their answers, participants were reminded that the protagonist's classmates had seen the zookeeper feeding the animals from the boxes and the protagonist themselves had not. Thus, in Experiment 2, the protagonist.

**Results.** We fit separate binomial logistic regressions for the Belief questions and Accuracy questions with random intercepts for each participant and Emotion and age in years (centered on the mean of the child sample) as fixed effects. This model suggested that children inferred that the protagonist had *not* genuinely changed his belief after being shouted at ($\beta_{Int}$ = -1.73, *SE* = 0.44, *z* = -3.98, *p* < .001), with no effect of age ($\beta_{Ct\_AgeYears}$ = -0.20, *SE* = 0.29, *z* = -0.69, *p* = .491); however, they were significantly more likely (with data from n=69 children out of the planned n=80) to infer that he had genuinely changed his answer when the group expressed surprise and talked with him than when they shouted at him ($\beta_{EmotionSurprise}$ = 1.34, *SE* = 0.47, *z* = 2.87, *p* = .004), with no interaction ($\beta_{Emotion*Ct\_AgeYears}$ = -0.07, *SE* = 0.36, *z* = 0.29, *p* = .85). Following our preregistered analysis plan, we computed 95% confidence intervals for each AgeGroup and Emotion to compare to chance. In the Anger trial, 23% of the younger children and 15% of the older children inferred that the protagonist had genuinely changed his belief (*Younger*: 8 of 35, 95CI: 10.4–40.1; *Older*: 5 of 34, 95CI: 5.0–31.1). However, 38% of Adults inferred that the belief change was genuine (*Adults*: 15 of 40, 95CI: 22.7–54.2). In the Surprise trial, though 80% of adults inferred the belief change was genuine (32 of 40, 95CI: 64.4–90.9), only 43% of the younger children (*Younger*: 15 of 35, 95CI: 26.3–60.6) and 41% of the older children (*Older*: 14 of 34, 95CI: 24.6–59.3) believed the change was genuine.

Despite this general skepticism among younger children that the protagonist had genuinely changed his belief for either emotion — and in contrast to Experiment 1 — children overall were equally like to endorse the consensus

answer regardless of affect ($\beta_{Int}$ = 1.65, *SE* = 0.55, *z* = 3.01, *p* =.003, with no change across development or interaction (all other factors not significant; with data from n=69 children out of the planned n=80). Following our preregistered analysis plan, we also computed 95% confidence intervals for each AgeGroup and Emotion to compare to chance. As predicted, all ages endorsed the consensus on both the Anger trial (*Adults*: 88%, 95CI: 73.2–95.8; *Older*: 79%, 95CI: 62.1—91.3; *Younger*: 71% 95CI: 53.7—85.4) and Surprise trial (*Adults*: 93%, 95CI: 79.6–98.4; *Older*: 85%, 95CI: 68.9—95.0; *Younger*: 80% 95CI: 63.1—91.6).

## Discussion

Human societies are built on outsourced knowledge: from the cutting edge of technological development and policy decisions to educational curriculums and coordinating carpool schedules, learners must rely on the good judgment of groups of individuals who are themselves subject to influences from first- and secondhand sources, the pressure of consensus, and social motivations such as in-group status and prestige. Because these influences can both clarify and corrupt collective judgments, the process by which individual judgments are aggregated into a collective judgment may be as informative as the judgments themselves.

The influence of affective cues on children's own endorsements as well as their inferences about the protagonists' beliefs suggest that children do reason about how individual answers are aggregated into a collective response. The difference in Accuracy endorsements between Experiments 1 and 2 suggests that while both children and adults will disregard negative affect if a consensus has firsthand knowledge and a dissenter does not, children and adults weigh affective and consensus cues differently if the two sides both have firsthand knowledge: while adults give more weight to consensus, children may give more weight to affect. Still, both older children and adults were less likely to endorse an angry consensus than a surprised consensus in Experiment 1, consistent with past work.

The reasons for children's uncertainty about whose answer in Experiment 1 was more accurate are less clear. Though we predicted that they would reject the angry consensus in favor of the protagonist's original answer, their judgments were at chance on the anger trial as well as the surprise trial. Given that in past work, children endorse a single informant relaying the testimony of three eyewitnesses over three informants relaying the testimony of a single eyewitness by age 6 (Aboody, Yousif, Sheskin, & Keil, 2021), it seems unlikely that the 6- and 7-year-olds in our experiments failed to recognize the consensus. Moreover, given their strong inferences about the protagonist's beliefs in both experiments and their endorsement of the angry consensus in Experiment 2 (pending complete data), it seems unlikely that they simply failed to understand the question more generally. Rather, children's general skepticism about the genuineness of the protagonist's belief change — with the youngest children doubting the conversion even on the Surprise trial — seems to suggest that younger children may have attempted to evaluate the consensus from the protagonist's point of view. Indeed, children who inferred that the protagonist *had* genuinely converted on the Surprise trial in Experiment 1 were more likely to also endorse the consensus answer themselves, while those who inferred that the protagonist had *not* genuinely converted on the Surprise trial were equally likely to endorse either the consensus or the protagonists' original answer.

If children were relying on the protagonist's judgment in the Surprise trial in Experiment 1, then children should be more likely to trust the consensus if the protagonist converts after the consensus has the opportunity to demonstrate their accuracy to the protagonist. For instance, if instead of having to adjudicate conflicting memories, the students were told the number of food items in the box and the number of items each animal would eat, the potential for an unambiguous demonstration of the correct answer might lead children to infer that Max would only change his answer if he was satisfied with the demonstration. Future work will address this possibility. Indeed, Experiment 2 may show some evidence of this kind of "vicarious persuasion" in that adults and children appear to be more willing to believe that the protagonist genuinely changed his mind even after being shouted at: adults, for example, may have inferred that the angry group could have simply told the protagonist that they had seen the answer. Adults are more willing to believe statements that a source claims to know firsthand; inferring similar behavior from a third-party could produce inferences similar to those observed for the Belief measures in Experiment 2.

## Acknowledgments

This research was funded by NSF grant DRL 1561143 to Frank Keil.

## References

Aboody, R., Yousif, S. R., Sheskin, M., & Keil, F. C. (2022). Says who? Children consider informants' sources when deciding whom to believe. *Journal of Experimental Psychology: General*.

Barber, B. M., Odean, T., & Zhu, N. (2009). Systematic noise. *Journal of Financial Markets*, 23.

Becker, J., Almaatouq, A., & Horvat, E. A. (2020). Network Structures of Collective Intelligence- The Contingent Benefits of Group Discussion. *PsyArxiv*.

Boehm, C. (1996). Emergency Decisions, Cultural-Selection Mechanics, and Group Selection [and Comments and Reply]. *Current Anthropology*, *37*(5), 763–793.

Claidière, N., & Whiten, A. (2012). Integrating the study of conformity and culture in humans and nonhuman animals. *Psychological Bulletin*, *138*(1), 126–145.

Cordonier, L., Nettles, T., & Rochat, P. (2018). Strong and strategic conformity understanding by 3- and 5-year-old children. *British Journal of Developmental Psychology*, *36*(3), 438–451.

Einav, S. (2014). Does the Majority Always Know Best? Young Children's Flexible Trust in Majority Opinion. *PLoS ONE*, *9*(8), e104585.

Einav, S. (2018). Thinking for themselves? The effect of informant independence on children's endorsement of testimony from a consensus. *Social Development*, *27*(1), 73–86.

Fusaro, M., & Harris, P. L. (2008). Children assess informant reliability using bystanders' non-verbal cues. *Developmental Science*, *11*(5), 771–777.

Hahn, U., von Sydow, M., & Merdes, C. (2019). How Communication Can Make Voters Choose Less Well. *Topics in Cognitive Science*, *11*(1), 194–206.

Hastie, R., & Kameda, T. (2005). The Robust Beauty of Majority Rules in Group Decisions. *Psychological Review*, *112*(2), 494–508.

Haun, D. B. M., & Tomasello, M. (2011). Conformity to Peer Pressure in Preschool Children: Peer Pressure in Preschool Children. *Child Development*, *82*(6), 1759–1767.

Haun, D. B. M., van Leeuwen, E. J. C., & Edelson, M. G. (2013). Majority influence in children and other animals. *Developmental Cognitive Neuroscience*, *3*, 61–71.

Hu, J., Whalen, A., Buchsbaum, D., Griffiths, T., & Xu, F. (2015). Can Children Balance the Size of a Majority with the Quality of their Information? *Proceedings of the Annual Meeting of the Cognitive Science Society*, *35*, 6.

Johnston, A. M., Mills, C. M., & Landrum, A. R. (2015). How do children weigh competence and benevolence when deciding whom to trust? *Cognition*, *144*, 76–90.

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.

Landrum, A. R., Mills, C. M., & Johnston, A. M. (2013). When do children trust the expert? Benevolence information influences children's trust more than expertise. *Developmental Science*, *16*(4), 622–638.

Lane, J. D., Ronfard, S., & El-Sherif, D. (2018). The Influence of First-Hand Testimony and Hearsay on Children's Belief in the Improbable. *Child Development*, *89*(4), 1133–1140.

Lane, J. D., Wellman, H. M., & Gelman, S. A. (2013). Informants' Traits Weigh Heavily in Young Children's Trust in Testimony and in Their Epistemic Inferences. *Child Development*, *84*(4), 1253–1268.

List, C., & Goodin, R. E. (2001). Epistemic Democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy*, *9*(3), 277–306.

Magid, R. W., Yan, P., Siegel, M. H., Tenenbaum, J. B., & Schulz, L. E. (2018). Changing minds: Children's inferences about third party belief revision. *Developmental Science*, *21*(2), e12553.

Mahr, J. B., & Csibra, G. (2021). The effect of source claims on statement believability and speaker accountability. *Memory & Cognition*, *49*(8), 1505–1525.

Mannes, A. E. (2009). Are We Wise About the Wisdom of Crowds? The Use of Group Judgments in Belief Revision. *Management Science*, *55*(8), 1267–1279.

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.

Morgan, T. J. H., Rendell, L. E., Ehn, M., Hoppitt, W., & Laland, K. N. (2012). The evolutionary basis of human social learning. *Proceedings of the Royal Society B: Biological Sciences*, *279*(1729), 653–662.

Richardson, E., & Keil, F. (2021). You can't trust an angry group- asymmetric evaluations of angry and surprised rhetoric affect confidence in trending opinions. *Proceedings of the Cognitive Science Society*.

Toelch, U., Bruce, M. J., Meeus, M. T. H., & Reader, S. M. (2010). Humans copy rapidly increasing choices in a multiarmed bandit problem. *Evolution and Human Behavior*, *31*(5), 326–333.

van Leeuwen, E. J. C., Cohen, E., Collier-Baker, E., Rapold, C. J., Schäfer, M., Schütte, S., & Haun, D. B. M. (2018). The development of human social learning across seven societies. *Nature Communications*, *9*(1), 2076.