Title

Investigating Sensitivity to Shared Information and Personal Experience in Children's Use of Majority Information

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

Authors
Otsubo, Kay
Whalen, Andrew
Buchsbaum, Daphna

Publication Date
2017

Peer reviewed

Investigating Sensitivity to Shared Information and Personal Experience in Children's Use of Majority Information

Kay Otsubo (kay.otsubo@mail.utoronto.ca)
Department of Psychology, University of Toronto
Toronto, ON M5S 3G3 CAN

Andrew Whalen (awhalen@roslin.ed.ac.uk)
The Roslin Institute, University of Edinburgh
Midlothian, SCT EH25 9RG UK

Daphna Buchsbaum (buchsbaum@psych.utoronto.ca)
Department of Psychology, University of Toronto
Toronto, ON M5S 3G3 CAN

Abstract

When learning from others, rather than simply following the majority's opinion, we need to accurately evaluate the quality of the information both the majority and the minority provide, and integrate that information with our own personal experience. This is especially true when the majority's opinion is based on lower quality information, because they shared the same evidence rather than collecting evidence independently. Previous work demonstrated that adults are sensitive to the quality of the majority's information, consistent with the predictions of a Bayesian rational model (Whalen, Griffiths, & Buchsbaum, in press). In two behavioural experiments, we investigated how preschoolers combine testimony from a majority that conflicts with a minority or with the child's own personal evidence. Unlike adults, children over-relied on the majority when given only testimony. However, when also given their own conflicting evidence, children relied significantly less on the majority and over-relied on their own evidence. These findings help explain why children may follow the majority at times, but in others trust their own judgements.

Keywords: Selective Trust; Conformity Bias; Children; Statistical Dependency; Bayesian Modeling; Social Learning

Introduction

Learning from others is a valuable strategy to use when encountering uncertainty. Human's use of social information is thought to underlie our ability to live in almost every known environment, and to underpin the evolution of human culture (Boyd & Richerson, 1985). The information we receive from others is also integrated with the information we gather ourselves through personal experience. Rather than blindly following what others say, it is important we evaluate the information we receive by understanding how others may have formed their opinions, especially when there is disagreement amongst individuals or when their testimony conflicts with our own personal observations.

Evaluating informant testimony is particularly important when the testimony is based on shared information. For example, imagine you are reading four restaurant reviews. In the first two reviews, the reviewers independently visited the restaurant and ordered different dishes, and they both recommended the restaurant. In the third and fourth reviews, the reviewers went to the restaurant together and shared a single dish, and then both did not recommend the restaurant. Which set of reviews should you trust?

From a pure numbers standpoint, the number of positive and negative reviews is equal. However, the first and second reviews may provide additional information about what an average experience (your experience) at the restaurant would be like. The third and fourth reviews provide less information because the reviewers' shared experience makes their responses statistically dependent on each other. For instance, if the two reviewers shared an unusually salty dish they are both likely to write negative reviews and as such, given the third review, the fourth review provides no new information. Thus, being aware of this statistical dependence will help social learners avoid the mistake of placing trust in a group based on just their number of opinions.

This ability to assess the quality of information being provided is especially important for young children, who are learning much about the world from the testimony of others (e.g., Mills, 2013). Whether young children can use statistical dependency to evaluate testimony quality is an interesting question because much of our social learning occurs during early childhood, when the ability to understand mental states is still developing (Wellman, Cross, & Watson, 2001). The ability to assess quality of information may require a complex form of "theory of mind" that goes beyond simply copying the majority. To accurately assess the quality of information, children must consider not only the testimony each person gave, but also the unseen information leading to that testimony, and how that information was gathered.

In addition, one's own personal experience can also conflict with what others say, and must be integrated with the testimony received. Imagine that you have a negative experience at that restaurant, and are debating whether to go there again. If there are enough other positive reviews, you may be willing to disregard your own judgement, and give the restaurant another chance. Can children evaluate the quality of their own information relative to a conflicting majority in a similar manner?

In two experiments, we examined how 4- and 5-year-old children evaluated their own private information and the information they received from informants who either shared a piece of evidence or collected evidence independently. Specifically, we investigated whether children can distinguish the quality of information provided by multiple informants and exhibit a sensitivity to shared data (as suggested by Hu et al., 2015) or if they merely conform to the majority (Corriveau, Fusaro, & Harris, 2009). We then compared children's performance to that of adults' on a similar task (Whalen et al., in press), and the predictions of a Bayesian rational model to understand the extent to which children may conform to the majority despite the amount of information the majority provided.

In Experiment 1, we investigated whether children were sensitive to evidence being shared by a majority group when a single dissenter with independent evidence was present. We found that children were biased towards following a majority opinion and were not sensitive to statistical dependencies between informants. In Experiment 2, we highlighted the source of informants' knowledge by providing children with their own private evidence that conflicted with the majority. We found that given conflicting personal evidence, children no longer followed the majority and instead sided with their own evidence regardless of the quality of the majority's information. Compared to both adults and to the rational model, children were not sensitive to dependency, trusting statistically dependent informants more in Experiment 1, and placing more weight on their own evidence in Experiment 2.

Background

Previous work by Whalen et al. (in press) demonstrated that adults are sensitive to statistical dependency between informants. Participants correctly rated that an option was more likely when it was endorsed by a majority group with independent evidence than a group with shared evidence (see Figure 1(b) for results). Adults also integrated their own evidence with testimony, appropriately demonstrating no bias towards their own evidence when it conflicted with the majority endorsement. In particular, they endorsed the majority opinion when the group had a higher quality of information than provided by their own personal evidence.

These findings were consistent with a Bayesian model of social learning which captures how an idealized learner might learn from multiple informants with shared information. The model illustrates that conforming to the majority is rational when the majority has a greater quality of information because (like our independent restaurant reviewers) each member contributed additional independent information. Thus, although in some cases adults disregard their own evidence and favour the majority, this may be a product of rationally integrating the two sources of information and assessing their quality, and not a bias towards the majority.

In this paper, we investigated whether 4- and 5-year-old children could assess the information quality provided by a majority when it conflicted with the information of a minority or with the child's own personal evidence. At the age of 4, children already start to implement strategies in choosing who to listen to by selectively trusting informants, for instance by preferring those who are knowledgeable or accurate (Koenig & Harris, 2005), or experts in the field (e.g., Kushnir, Vredenburgh, & Schneider, 2013). However, the current literature is unclear on whether children value conformity or information quality during social learning.

Previous studies argued that children value a consensus even when it conflicts with the child's own perception. For instance, children sometimes followed the majority even when they understood and identified the endorsement of the majority to be incorrect (Corriveau & Harris, 2010; Haun & Tomasello, 2011). Using the Asch (1956) paradigm, children were observed to conform especially when answering in public in front of their peers (Haun & Tomasello, 2011). At the age of four, children are already capable of recognizing a consensus and conforming to them even in ambiguous tasks such as labelling a novel object (Corriveau et al., 2009). These findings then suggest that children may have a bias to conform to a majority, even when the conflicting information comes from their own perception.

On the other hand, some studies have argued that children do exhibit the ability to evaluate the quality of information they receive from multiple informants (Hu et al., 2015). Hu and colleagues (2015) found that, when given testimony from two groups, children preferred the group with the highest quality of knowledge – favouring the group that received direct knowledge via visual perception over those who received indirect knowledge via hearsay. However, when group sizes were not equal, children preferred the group with the most members, even if the members of the larger group had only received hearsay.

Additional work has shown that children avoid a conformity bias if the majority group is proven to be unsuccessful in reaching an apparent goal (Wilks et al., 2014), provide implausible functions for a novel object (Schillaci & Kelemen, 2014), or have lower expertise than the minority (Burdett et al., 2016). These findings then suggest that, at least in some cases, children have a preference for informants with a greater quality of knowledge, rather than having a preference for the majority per se.

Preschoolers also demonstrate the ability to integrate a single informant's testimony with their own observations. When the two sources conflict, preschoolers acknowledge the confidence and statistical data provided by an informant to assess causal relationships of novel toys (Bridgers et al., 2015), and acknowledge an informant's awareness for appearance-reality when considering their own perception as misleading (Lane et al., 2014). These findings suggest that children can integrate both sources of information, which contrasts with previous theoretic models that emphasized reliance on only social learning (e.g., Rendell, Fogarty, & Laland, 2010; but see Perreault, Moya, & Boyd, 2012).

Therefore, whether children can appropriately integrate the quality of informants' knowledge given a majority and conflicting information from either a minority or from personal observation is an open question. However, it is not always obvious how these different sources of information should normatively be integrated.

Bayesian Model of Learning from Independent and Dependent Informants

To further understand how an individual can combine the information they receive from testimony and personal evidence, we followed the Bayesian model developed by Whalen et al. (in press) which captures how an idealized learner would integrate information provided by groups with different sources of data – shared or independent – with personal evidence. In this model, a learner collects personal evidence about the state of the world, e, and receives testimony from n informants, $t_1, \ldots, t_n$ who collect their own evidence about the state of the world, d. Learners evaluate a potential hypothesis, h, using Bayes' rule,

$$p(h|e, t_1, \ldots, t_n) \propto p(t_1, \ldots, t_n|h)p(e|h)p(h), \quad (1)$$

where $p(h|e, t_1, \ldots, t_n)$ is the posterior probability of h, the probability that a hypothesis about the state of the world is true given the personal evidence and testimony, while $p(h)$ is the prior probability of h, the probability the hypothesis is true before any evidence is given. Finally, $p(e|h)$ is the probability of getting that evidence given the hypothesis, and $p(t_1, \ldots, t_n|h)$ is the probability of getting that testimony.

When multiple informants provide independent testimony, the probability of a series of testimony is equivalent to the product of the probability of each individual testimony:

$$p(t_1, \ldots, t_n|h) = \prod_{i=1}^{n} p(t_i|h). \quad (2)$$

The testimony of each informant is based on their private data d_i, so $p(t_i|h)$ is obtained by marginalizing over d_i:

$$p(t_i|h) = \sum_{d_i} p(d_i|h)p(t_i|d_i), \quad (3)$$

where $p(t_i|d_i)$ is the probability that the informant produces testimony t_i after observing d_i. On the other hand, when multiple informants base their testimony on shared private data, denoted as d', the probability of a series of testimony is obtained by marginalizing over the shared private data:

$$p(t_1, \ldots, t_n|h) = \sum_{d'} p(d'|h) \prod_{i} p(t_i|d'). \quad (4)$$

In both cases, we assume that informants give testimony in support of a hypothesis proportional to the product of the informant's evidence given the hypothesis and the prior probability, $p(t_i = h_i|d_i) \propto p(d_i|h_i)p(h_i)$ (for more information, see Whalen et al., in press). This Bayesian model illustrates that, in many cases, conforming to the majority is rational when the majority collects independent evidence, increasing their quality of information (see Figure 1(a) for example predictions based on our experiment task). In addition, this model accurately predicted the performance of adults in our experimental task suggesting that adults integrate both sources of information rationally (see Figure 1(b) for adult performance).

Following the approach of Whalen et al. (in press), we ran two behavioural experiments that examined how children evaluated the information they were provided by a majority group with shared or independent evidence, along with either a dissenting informant or conflicting private evidence.

Experiment 1: Dissenting Informant

In Experiment 1, children were shown a video about two jars with differing proportions of red and yellow balls and were asked to guess which jar was being sampled from, given the testimony of three friends who received a ball from the chosen jar. The first two informants endorsed the same jar and made up the majority group while the third informant dissented and endorsed the opposing jar. Children were randomly assigned to either the *Shared* condition, where the majority shared one ball, or the *Independent* condition, where majority members each received their own ball.

Our model predicts that a rational learner would choose the jar endorsed by the majority only when each member collected independent evidence, but be at chance when the majority shared one piece of evidence, Figure 1(a). On the other hand, if children have a conformity bias, we should expect them to pick the majority's jar in both conditions.

Methods

Participants A total of 29 preschoolers (female = 17, male = 12; mean age = 4 years 11 months; range = 49 – 71 months) were recruited either through local museums or in lab. They were randomly assigned to one of two conditions: The *Independent* condition (n = 14) or the *Shared* condition (n = 15). An additional 6 children were excluded due to atypical development (1), provided ambiguous answers (1), did not provide an answer (2), or experimenter error (2).

Procedure In this experiment, children were shown a video on a laptop, where an experimenter introduced two jars comprised of coloured balls – one with mostly yellow balls, a few red balls, and one green ball and one with mostly red balls, a few yellow balls, and one white ball, and introduced her three adult friends. The experimenter explained that she would pour just one of the two jars into her bag and give each of her friends a ball from the bag. Each of her friends would then tell her which of the two jars they thought she picked.

Once the experimenter filled the bag with one of the two jars, she used a cup to randomly scoop a ball from the bag to hand to each of her friends. After looking inside the cup, the informants provided testimony as to which jar they thought the bag was filled from, either the jar with mostly red or mostly yellow balls in it. The first two informants always endorsed the same jar and made up the majority group, while the last informant always chose the opposite jar. The jar

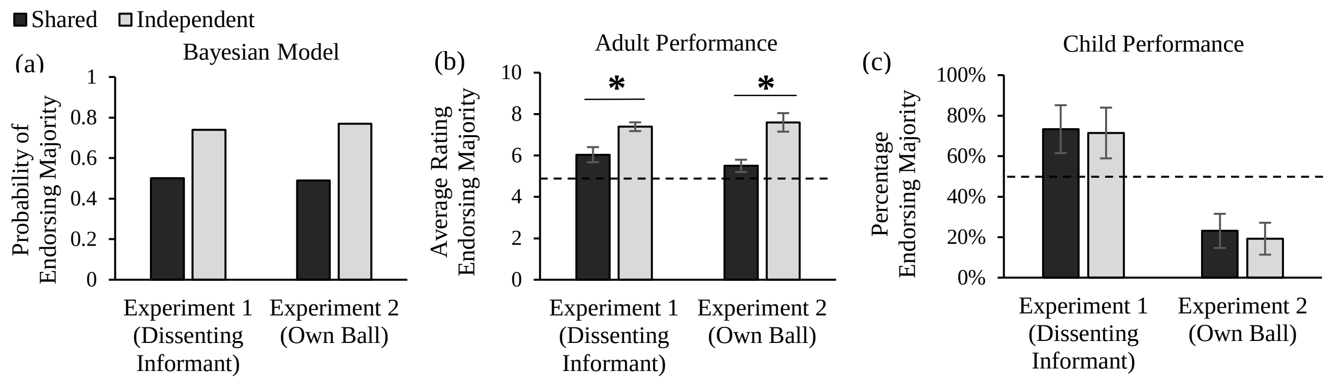

Figure 1: Probability of endorsing the majority opinion in both experiments as predicted by the Bayesian model (a), the average rating of endorsing the majority in both experiments as performed by adults (b) from Whalen et al. (in press), and the percentage endorsing the majority in both experiments as performed by children (c) in the current study.

endorsed by the majority and the actor playing the minority informant were counterbalanced.

In the *Independent* condition, informants were in the room one at a time and were each given their own randomly sampled ball to view. Each informant stated, for instance, "I looked at the ball and I think that the bag has mostly red balls in it." In the *Shared* condition, all three informants were present in the room and the first two informants shared a single randomly sampled ball. After providing a testimony, the first informant was asked to pass the same cup to the second informant who then agreed with the first, e.g. "I looked at the ball and thought about what my friend said. I agree with Jessie. I think that the bag has mostly red balls in it." While the third informant received a different random ball and disagreed with the rest, e.g. "I looked at the ball and thought about what my friends said. I disagree with Jessie and Sarah. I think that the bag has mostly yellow balls in it."

Once the video was completed, the on-site experimenter, a different person than the one in the video, reminded the child which jar each informant endorsed and if they saw the same or a different ball as the previous informant, and that all the balls came from just one jar. Finally, she asked a forced-choice question of which jar the child thought the bag was filled from, either the jar with mostly red or yellow balls in it. The order in which the jars were stated was randomized.

Results

Each child was given a score of 0 or 1, with 1 as agreeing with the majority and 0 as disagreeing. Results are shown in Figure 1(c). Overall, children chose the jar endorsed by the majority significantly more often than chance, regardless of how the members of the majority collected their information (binomial test, 21 out of 29 endorsed the majority, $p = 0.024$).

We analyzed the differences between the conditions using a Fisher's exact test. The difference between the *Independent* and *Shared* condition was not significant (10 out of 14 endorsed the majority in the *Independent* condition, 11 out of 15 in the *Shared* condition, $p = 1$). Children chose the majority's jar equally often when the majority had higher quality independent information and lower quality shared information. Finally, compared to adults and to our model, children appeared to place more weight on the statistically dependent testimony in the *Shared* condition.

Discussion

Unlike our model predictions and adults' performance, children were not able to appropriately evaluate the quality of information in an informant's testimony. When two informants received the same ball and gave the same testimony, children over-weighed the majority's shared information relative to the dissenter's independent information. These results support previous findings by Corriveau et al. (2009) who found that children conform to a majority when faced with an ambiguous decision.

Based on the results of Experiment 1, we wanted to identify ways to help children avoid relying on the majority and instead, evaluate which group has the greater quality of information. To do this, we highlighted the independent nature of the minority information by having the child receive private evidence that conflicted with the majority testimony which mimicked many real-world scenarios where our own private experience conflicts with testimony. If the child then has to integrate social learning with personal observation, this may help identify the source of knowledge each individual has and overcome a conformity bias.

Experiment 2: Own Ball

In Experiment 1, children were making a decision based on testimony alone. However, in most real-world cases, we take in the information that others provide us and evaluate it with our own information. Therefore, in Experiment 2, children no longer saw a minority group, and instead were given their own ball from the bag that conflicted with the testimony. For example, if the informants all endorsed the jar with mostly red balls, the child received a yellow ball from the bag.

As predicted by the Bayesian model, children should choose the jar endorsed by the majority when the members independently collected data and have more information than provided by the child's own single piece of evidence. If, however, the members of the majority shared a single piece of evidence, children should endorse the majority at chance, as the child's own evidence would be as reliable as the majority's. Similar to Experiment 1, if children present a conformity bias, they will follow the majority regardless of the quality of information provided by the group.

Methods

Participants A total of 52 preschoolers (female = 24, male = 28; mean age = 4 years 11 months; range = 48 – 71 months) were recruited through local museums and daycares, or in lab. They were randomly assigned to one of two conditions: The *Independent* condition (n = 26) or the *Shared* condition (n = 26). An additional 14 children were excluded due to experimenter error (10), previous participation in Experiment 1 (1), inattentiveness (1) and ambiguous answers (2).

Procedure Experiment 2 had the same jars and actors, and similar sampling procedures. However, all three informants endorsed the same jar and made up the majority group for both conditions. As in Experiment 1, children were randomly assigned to either the *Independent* or *Shared* condition.

In the *Independent* condition, each informant received their own distinct randomly sampled ball and provided their testimony in the room one at a time. In the *Shared* condition, all three informants were present in the room and shared a single randomly sampled ball and provided testimony agreeing with the previous informants.

After the video ended, the on-site experimenter reminded the child which informant endorsed which jar, if they looked at the same or different ball as the previous informants, and that all the balls came from just one jar. In this experiment, all the informants endorsed only one jar. Next, the on-site experimenter brought out an identical bag and stated that it was the same bag from the video containing the same balls. Similar to the experimenter in the video, she used a plastic cup to give the child their own ball from the bag. The on-site experimenter pretended to scoop up a ball at random, but in fact the child always received a ball that was a different colour from the majority testimony. After the child looked inside the cup, the on-site experimenter asked which jar they thought all the balls came from, as in Experiment 1.

Results

Results are shown in Figure 1(c). Overall, children chose the jar endorsed by the informants below chance regardless of how the informants collected their information (binomial test, 11 out of 52 endorsed the majority, $p < 0.001$). Similar to Experiment 1, we found that children did not choose the informant's jar more in the *Independent* condition compared to the *Shared* condition (5 out of 26 endorsed the majority in the *Independent* condition, 6 out of 26 in the *Shared* condition, $p = 1$, Fisher's exact test). In both conditions, children weighed their own evidence more, compared to both adults and the predictions of the Bayesian model.

Discussion

We found that in Experiment 2, children relied heavily on their own evidence and chose the jar consistent with their own ball regardless of whether the majority collected independent or shared evidence. As in Experiment 1, we found no significant difference between the *Independent* and *Shared* conditions, suggesting that children were not sensitive to the statistical dependency. These results support previous findings suggesting that children may rely on their own evidence that they personally collected over the evidence collected by others (Kushnir & Gopnik, 2005; Kushnir, Wellman, & Gelman, 2009).

General Discussion

We investigated how children weighed the value of information they received from multiple individuals and their own personal evidence. We compared the performance of 4- and 5-year-old children to the performance of adults on a similar task (Whalen et al. in press) and to the predictions of a Bayesian rational model. Experiment 1 showed that children were not sensitive to the shared information of a group and were instead following the majority. Experiment 2 demonstrated that children would no longer use the strategy of conforming to the majority if they themselves collected conflicting evidence and instead, relied on their own evidence, regardless of the quality of the majority's information. Therefore, compared to adults, children applied a different strategy in the integration of information.

Children's apparent conformity bias is consistent with previous findings that argued that children prefer to rely on the majority (e.g., Corriveau et al., 2009; Haun, Rekers, & Tomasello, 2012; Haun & Tomasello, 2011). Children may exhibit this reliance on the majority because it is often a useful and reliable social learning strategy (Haun et al., 2014). After all, the majority made their choices for a reason.

However, in the presence of children's own conflicting evidence, a conformity bias was no longer present. Although previous work has suggested that children may present a conformity bias even in the face of conflicting direct perception (Corriveau & Harris, 2010; Haun & Tomasello, 2011), it is important to note that in those studies, a majority of children still favoured their own evidence over the testimony of the majority.

This bias to rely on personal evidence over the evidence collected by others has previously been observed in the causal domain as a self-agency bias, especially when the evidence seemed to be ambiguous or probabilistic (Kushnir et al., 2009). Kushnir and Gopnik (2005) discovered that children would weigh their own causal interventions more heavily than the causal interventions of others. They suggested that children had this bias because they viewed their own actions to be more controlled and reliable and less likely to be confounded than those of other individuals. However, this bias to one's own evidence has not yet been observed in a non-causal domain like that of our current study. Similarly, children in our study might have considered their direct perception of their own ball to be more reliable than the information they received from the informants' testimony.

One possible explanation for the presence of both a bias towards conformity and towards personal evidence in this task is that preschoolers learning how to integrate information have yet to develop more complex aspects of theory of mind (e.g., Gweon et al., 2012). At this age, children might have had difficulty reasoning about how the informants generated their testimony based on the evidence that they likely received. In other words, 4- and 5-year-old children know that people can have beliefs, but may have difficulty in knowing how these people came to believe something. As a consequence, children might rely on the number of endorsements given rather than on their quality, leading to the appearance of a conformity bias. On the other hand, children were likely confident in what they themselves saw which appeared as a bias towards their own evidence.

Future work should investigate whether children have difficulty inferring the evidence informants likely received based on their testimony, by testing how children respond when they can observe this evidence directly, for instance by presenting it in clear cups. We expect that children would then compare the amount of evidence between the majority and minority group rather than the number of endorsements. If children can identify the statistical dependency when the evidence is visible, they should no longer demonstrate a conformity or a personal evidence bias.

In addition, in ongoing follow-up studies, the child's evidence is presented on-screen within the video rather than performed live to equate saliency. If children still present a bias towards their own evidence, the salience from a live performance as a reason for this bias can be ruled out.

Taken together, our findings suggest that children implement a different social learning mechanism than adults and our Bayesian model. When integrating testimony alone, children over-weighed the quality of information provided by the majority. On the other hand, when the child was given their own conflicting evidence, children under-weighed the quality of information provided by the majority and relied on their own perception. Thus, unlike adults, children require further development in their social learning and perhaps their reasoning of mental states to avoid biases and become sensitive to statistical dependency.

Acknowledgments

We are grateful for all the families for their participation, and to the Computational Cognitive Development lab for their help with data collection, especially Nafisa Bhuiyan and Elina Kadriu, and Tom Griffiths for helpful discussions. This work was supported by the Natural Science and Engineering Research Council of Canada and the Canada Foundation of Innovation.

References

Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, *70*(9), 1–70.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.

Bridgers, S., Buchsbaum, D., Seiver, E., Griffiths, T. L., & Gopnik, A. (2015). Children 's causal inferences from conflicting testimony and observations. *Developmental Psychology*, *52*(1), 9–18.

Burdett, E. R. 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).

Corriveau, K. H., Fusaro, M., & Harris, P. L. (2009). Going with the flow: Preschoolers prefer nondissenters as informants. *Psychological Science*, *20*(3), 372–377.

Corriveau, K. H., & Harris, P. L. (2010). Preschoolers (sometimes) defer to the majority in making simple perceptual judgments. *Developmental Psychology*, *46*(2), 437–445.

Gweon, H., Dodell-Feder, D., Bedny, M., & Saxe, R. (2012). Theory of mind performance in children correlates with functional specialization of a brain region for thinking about thoughts. *Child Development*, *83*(6), 1853–1868.

Haun, D. B. M., Rekers, Y., & Tomasello, M. (2012). Majority-biased transmission in chimpanzees and human children, but not orangutans. *Current Biology*, *22*, 727–731.

Haun, D. B. M., Rekers, Y., & Tomasello, M. (2014). Children conform to the behavior of peers; other great apes stick with what they know. *Psychological Science*, *25*(12), 2160–2167.

Haun, D. B. M., & Tomasello, M. (2011). Conformity to peer pressure in preschool children. *Child Development*, *82*(6), 1759–1767.

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 37th Annual Conference of the Cognitive Science Society*.

Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*(6), 1261–1277.

Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science*, *16*(9), 678–683.

Kushnir, T., Vredenburgh, C., & Schneider, L. A. (2013). "Who can help me fix this toy?" The distinction between causal knowledge and word knowledge guides preschoolers' selective requests for information. *Developmental Psychology*, *49*(3), 446–453.

Kushnir, T., Wellman, H. M., & Gelman, S. A. (2009). A self-agency bias in preschoolers' causal inferences. *Developmental Psychology*, *45*(2), 597–603.

Lane, J. D., Harris, P. L., Gelman, S. A., & Wellman, H. M. (2014). More Than Meets the Eye: Young Children's Trust in Claims That Defy Their Perceptions. *Developmental Psychology*, *50*(3), 865–871.

Mills, C. M. (2013). Knowing When to Doubt: Developing a Critical Stance When Learning From Others. *Developmental Psychology*, *49*(3), 404.

Perreault, C., Moya, C., & Boyd, R. (2012). A Bayesian approach to the evolution of social learning. *Evolution and Human Behavior*, *33*(5), 449–459.

Rendell, L., Fogarty, L., & Laland, K. N. (2010). Rogers' paradox recast and resolved: Population structure and the evolution of social learning strategies. *Evolution*, *64*(2), 534–548.

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.

Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: the truth about false belief. *Child Development*, *72*(3), 655–684.

Whalen, A., Griffiths, T. L., & Buchsbaum, D. (in press). Sensitivity to shared information in social learning. *Cognitive Science*.

Wilks, M., Collier-Baker, E., & Nielsen, M. (2014). Preschool children favor copying a successful individual over an unsuccessful group. *Developmental Science*, *18*(6), 1014–1024.