When does the majority rule? Preschoolers' trust in majority informants varies by task domain

Jane C. Hu (jane.hu@berkeley.edu)
Daphna Buchsbaum (daphnab@berkeley.edu)
Thomas L. Griffiths (tom_griffiths@berkeley.edu)
Fei Xu (fei_xu@berkeley.edu)
Department of Psychology, University of California, Berkeley, Tolman Hall
Berkeley, CA 94720 USA

Abstract

In order to learn about the world, young children rely on information provided by social partners. Past research has shown children consider a variety of factors when learning from others, including consensus. Corriveau, Fusaro, and Harris (2009) found that in an object labeling task, children trust responses that receive majority support, and they concluded that children prefer members of a majority as social informants. However, it is possible that children prefer majority members only in domains that rely strongly on socially constructed norms, such as object labeling, where non-social information is unavailable. We formalized this prediction using a rational model of learning from testimony across tasks, and compared our model's predictions to children's responses in object labeling and causal learning tasks. We find that in a causal learning task, a domain that relies less on socially constructed norms, children rely more on their personal observations than informant testimony.

Keywords: social learning; Bayesian modeling; social cognition; consensus; testimony; epistemic trust

Introduction

We humans are inherently social creatures, and throughout our daily interactions, we openly share our thoughts and opinions with one another. The ubiquity of our social sharing and learning is rare among animal species (Tomasello et al., 2005; Warneken & Tomasello, 2009), and has been cited as an explanation for the robustness of human culture (Boyd, Richerson, & Henrich, 2011). Listening to others who share their knowledge can save precious time and effort, as learning through experience can be difficult and time-consuming. In listening to others' testimony, we can instantly and effortlessly learn how to prepare a dish, where to hunt, or who to hire to fix the kitchen sink.

Learning from others is especially important for young children, who have a relatively small pool of life experiences to draw on in new situations. However, one potential drawback to social learning is the possibility of receiving incorrect or misleading information. Therefore, it would be advantageous for children to employ mechanisms to evaluate sources' reliability. Previous work has found that children use informants' past accuracy as an indicator of trustworthiness (Sabbagh & Baldwin, 2001, Birch, Vauthier, & Bloom, 2008; Koenig, Clement & Harris, 2004; Pasquini et al., 2004; Corriveau & Harris, 2009) and selectively imitate others (Gergely, Bekkering, & Kiraly, 2002; Brugger et al., 2007; Buchsbaum, Gopnik, Griffiths, &

Shafto, 2011; Schulz, Hooppell, & Jenkins, 2008). On the other hand, other studies suggest children's social learning is sometimes surprisingly unselective and irrational (Lyons, Young, & Keil, 2007; McGuigan & Whiten, 2009).

We can learn not only from reliable individuals, but also from "crowd sourcing" information from a group of people. Adults often turn to others for advice, assuming that opinions held by many must be valid by virtue of their popularity. This intuition echoes the law of large numbers in probability theory: the more individual testimonies, the more likely the collective conclusion of those testimonies is accurate. Corriveau, Fusaro, and Harris (2009) found that three- and four-year-old children view consensus as an indication of reliability; they were more likely to endorse novel object labels that received majority support, and to choose a member of the majority group as an informant about other object labels. The authors concluded that preschoolers prefer information endorsed by the majority, and prefer members of a majority as informants.

The *extent* to which children prefer members of a majority as informants is still unclear. One possibility is that children prefer majority members as informants in all situations where multiple testimonies are available. In this view, children would indiscriminately weigh information from others as the most valuable source of information, perhaps prizing it above their own observations. Social psychologists have discovered that consensus opinions can override adults' existing opinions (Asch, 1956; Cialdini & Goldstein, 2004), which can result in internalization of the consensus opinion (Kelman, 1958; Nolan et al., 2008).

However, if children are rational learners, they should not always prefer majority testimonies. Domain demands should affect the weight children place on others' testimony. When learning about domains that are heavily socially constructed (e.g. object labels or tool use conventions), testimony from others should be highly valuable because the relevant knowledge is transmitted through others, and children cannot learn this type of information on their own. By virtue of the social conventions that dictate object labeling, typically only one label is regarded as correct (Markman, 1989). Alternatively, learning about domains that are not socially constructed (e.g. causal relationships, or naïve physics), children should consider not only informants' endorsements but also their own knowledge, because this type of knowledge can be gained through personal observation and is not typically bound by social

convention. This leaves open the possibility that an effect can have multiple causes.

In this paper, we explore how children's endorsement of majority testimony varies as a function of domain type, and compare these empirical results to rational behavior as predicted by a Bayesian model of learning from testimony. Specifically, we compare children's endorsement of majority testimony in an object labeling task versus a causal learning task. We predict that when given two options – one endorsed by a three-person majority, and one endorsed by a single minority informant – children should be more likely to endorse the majority's testimony when learning socially constructed facts (object labeling) than when learning non-socially constructed facts (causal learning).

Modeling Testimony Across Task Domains

In order to rationally learn from others' testimony, children must consider several types of information: the testimonies themselves, their own observations, and social and pragmatic cues that can affect the interpretation of others' testimony. The specific cues and information available to children vary depending on domain, leading them to rely more heavily on pure testimony in some domains. Learning object labels is a task that is especially dependent on social conventions. Speakers of a language must implicitly agree that certain words refer to specific objects, concepts, or ideas (Clark, 1988; 1990) and use them accordingly. In contrast, causal knowledge can be gained through non-social cues, like personal experience, that also provide reliable information.

A Bayesian ideal observer model is a natural way to formalize our assumptions about the types of evidence available in these different domains, and about the pre-existing biases and pragmatic assumptions that learners may bring to linguistic versus causal inferences. Buchsbaum et al. (2012) developed a model of how a rational learner should make causal inferences from both informant testimony and direct observations of causal outcomes. In this model, the learner receives testimony from one or more informants about the causal efficacy of one or more actions, and may also observe the causal outcome of these actions. The learner's goal is to choose a causally effective action. Here, we adapt this model to compare rational inferences from testimony in object labeling and causal tasks.

Model Details

Our model for causal inference from testimony is very similar to the model presented in Buchsbaum et al. (2012). In this model, learners receive testimony $r_{c,i}$ from informant i about whether they think candidate cause c is effective. Learners can also directly observe the effects $e_{c,j}$ of those causes (with N_c being the number of observations of the effect of cause c) Each cause c has a true underlying causal strength w_c, where $p(w_c = \rho) = \gamma$ and $p(w_c = 1 - \rho) = 1 - \gamma$, where ρ is a relatively high causal strength value, and γ is the probability of a cause having high causal strength. The probability of an effect e following c is w_c. Each informant i

has knowledge about the strength of cause c, $k_{c,i}$.[1] We assume that $k_{c,i} \in \{0, 1\}$, corresponding to two possible states of knowledge of a cause: knowledgeable and naïve. If $k_{c,i} = 1$ (informant i knows about the causal strength of c), then $p(r_{c,i} = w_c \mid k_{c,i} = 1, w_c) = 1 - \varepsilon$ — an informant with knowledge of cause c will give correct testimony about the causal strength of c with probability $1 - \varepsilon$, where ε is a small probability of giving incorrect testimony. In this work, we use $\varepsilon = 0.01$. On the other hand, if $k_{c,i} = 0$ then $p(r_{c,i} = w_c \mid k_{c,i} = 0, w_c) = p(r_{c,i} = w_c \mid k_{c,i} = 0) = 0.5$ – the informant will guess uniformly at random between the two possible actions. The probability of informant i being knowledgeable about a particular cause is $p(k_{c,i} = 1) = \tau$.
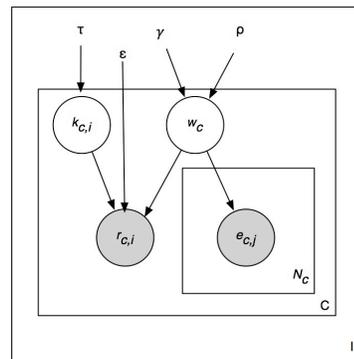

Figure 1: Dependencies of the variables in our Bayesian ideal observer model.

Finally, we assume $p(\text{choose } c) \propto p(\text{effect} \mid c, \text{obs})$ — children choose causes in proportion to how likely they think they are to produce the effect, given the evidence. We can use this model to compute the probability that the learner should choose to perform a particular action to bring about the effect, using the dependencies defined in our graphical model shown in Figure 1 (for further details, see Buchsbaum et al., 2012). We can use the same model to infer novel object's labels from testimony. In this case, we have unnamed objects instead of causes, and instead of a causal strength, each object has a probability of corresponding to the novel label. However, unlike the causal case, there are no independent observations to incorporate into the model; you cannot "see" whether a label truly names an object.

Model Predictions

We can examine a simple contrast between object labeling versus causal learning, where in each task, we assume two possibilities per task: i.e. two objects that could be the referent of a novel label, versus two actions to perform on a toy to elicit music. Corriveau, Fusaro, and Harris (2009) showed children three majority informants making one prediction, and a minority informant making an alternate prediction, so in a similar object labeling task, the majority

---

[1] The previous model also represented informants' expressed confidence, overall knowledgeability, and tendency towards overconfidence. These are not included in the current work.
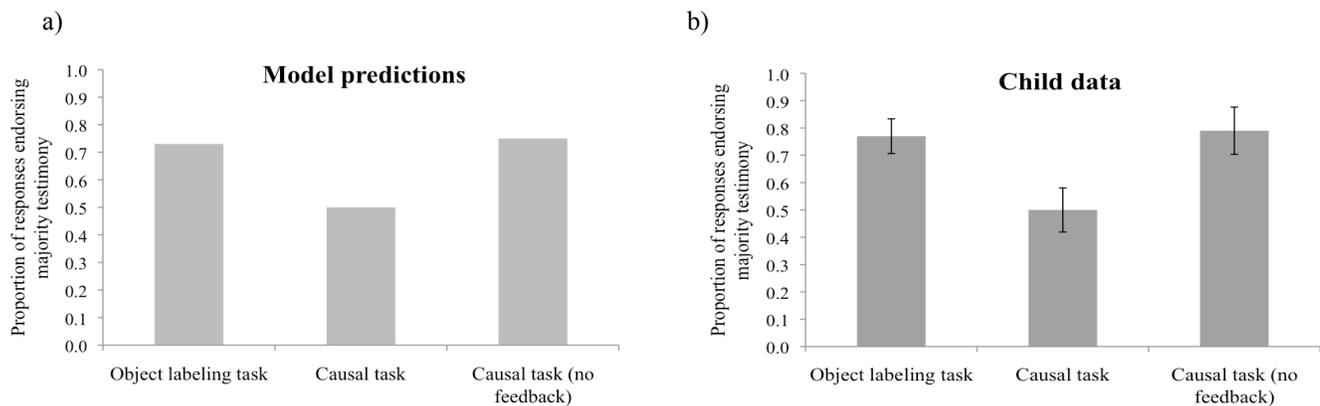
Figure 2: Proportion of responses endorsing majority testimony from (a) model predictions and (b) child data.

may each label Object 1 *modi* once, while the minority informant labels Object 2 *modi* three times. In a comparable causal task, the majority may all activate a toy using Action 1, while the minority informant activates the same toy three times using Action 2. Here, statements drawing attention to the demonstrated action are treated as testimony that the action is causally effective. As in the graphical model, the effect of an action is independent of the actor. We can use the model defined above to formalize some of the differences between these two tasks, then examine the model's predictions for whether rational learners should endorse majority testimony.

In labeling objects, we know that there exists a pragmatic mutual exclusivity assumption (Markman, 1989). If an informant labels Object 1 as the *modi,* this strongly implies that they believe that Object 2 is *not a modi*. In contrast, using one causal action does not necessarily imply that other actions are ineffective. We can capture this difference by having an informant's testimony that Object 1 is the *modi* implicitly include testimony that Object 2 is not a *modi*. In contrast, testimony in the causal case about the efficacy of Action 1 is left neutral with respect to the efficacy of Action 2. Instead, we treat testimony about Action 2 as unobserved for this informant (as are any demonstrations of Action 2 they might have performed).

Our remaining modeling assumptions are similar for both the causal and object labeling tasks. From previous work, we know that children assume that causes are relatively rare – most effects can only be brought about in one or two ways (Buchsbaum et., al, 2011; Bonawitz & Lombrozo, 2012). Similarly, children generally assume that an object has only one basic-level label (Markman, 1989), so if it is a *modi,* it is probably not also a *toma* or a *blicket*. We can represent both of these prior biases by using a small value for γ, making multiple causes and multiple labels relatively unlikely. We also know that children are biased to assume that causes are deterministic or near-deterministic (Schulz & Sommerville, 2006), and similarly that if an object is a *modi*, it is probably a *modi* every time, rather than occasionally something else. We can represent both of these assumptions using a high value for ρ. Finally, we know that children are a priori biased to assume adults are generally knowledgeable and helpful (Taylor, Cartwright, and Bowden, 1991), which can be represented by using a high value of τ.

We can now look at model predictions for the simple object labeling and causal inference tasks described above. We present predictions using the example parameter values $\gamma = 0.05$ $\rho = 0.9$ and $\tau = 0.8$ in Figure 2a. However, the qualitative differences in model predictions described below are robust to a wide range of parameter values, and in particular hold for any combination of values consistent with our assumptions. Given object-label testimony from a majority of three informants and one minority informant, the model predicts that learners should strongly favor the majority label. This is true not only if we explicitly force the model to consider only hypotheses where exactly one object is a *modi* (representing a hard mutual exclusivity constraint), but also if we remove this constraint, but continue to hold the softer pragmatic assumption that an informant who calls one object a *modi* is also saying that the other is not a *modi*.

In contrast, in the case where three informants activate a toy one way, and the minority informant activates the toy in another, the model predicts that after observing both actions bringing about the effect equally often, learners should be equally likely to choose either action themselves, despite the conflicting testimony. Finally, we examine a case where informants make causal predictions, but do not demonstrate the actions, paralleling the lack of non-testimony evidence in object labeling. In this case, we do not assume that predicting that one action is effective entails that the other action is not. Here, the model again predicts that the learner should endorse the majority's action choice, but *only* if they believe causes are rare. If they believe that causes are very common, they should continue to be roughly evenly split. Given our assumption that children are biased to believe causes are rare, we predict that they will again endorse the majority's demonstration in this case.

Experiment 1: Comparing tasks

In this study, we present preschoolers with four informants' conflicting testimony about objects. In the object labeling condition, informants identify the referent of a novel label, and in the causal learning condition, they demonstrate a novel action on the object that results in a song.

Methods

Participants Participants were 64 preschoolers, 29 male and 27 female (mean age = 4 years 2 months; range = 36 – 65 months). Participants were recruited in the San Francisco Bay Area by mail and phone calls or from local preschools

and museums. An additional five children were tested, but were excluded due to fussiness (4) or experimenter error (1).

Materials In the object labeling condition, stimuli were four novel objects. In the causal condition, stimuli were two plush toys, each of which contained a wireless, battery-powered doorbell chime box. The boxes played short melodies when activated by a handheld remote to create the illusion that children's actions were causally efficacious. Pre-recorded video clips of informants' testimonies were shown to children on a 13" laptop screen.

Procedure Participants were randomly assigned to either the object labeling condition or causal condition. Each participant participated in two test trials of their condition.

In each condition, the experimenter introduced novel objects to participants and explained that they were unknowledgeable about their labels or causes. Participants then watched four video clips of four informants evaluating the objects.

Each clip began with a female informant sitting at a table with the novel objects. She visually inspected them, then picked up one of the toys and called it by the novel label (e.g. *modi*), or acted on the toy, resulting in the toy playing a short song. In three of the four video clips, the "majority" informants each endorsed one object as a *modi* or performed one action to elicit music, and the minority informant informant endorsed the other object the *modi*, or performed an alternate action to elicit music. The minority informant always repeated the novel label or alternate action three times so that each participant heard the label used to refer to each object an equal number of times.

After participants watched the video clips, the experimenter presented the child with the objects from the video clips and asked children to identify the referent of the novel label, or to make the toy play music. Participants' first gestural or vocal response was recorded. Participants in the causal condition were invited to activate toys three times.

Half of the video clips were mirror images of original recordings to control for the location of objects (object labeling condition) and handedness of informants when manipulating toys (causal condition). The trial presented first and identity of the minority informant were also counterbalanced.

Results and Discussion

Participants were assigned a score (0, 1, or 2) based on the number of trials in which they endorsed the majority informants' testimony (0-2) first responses in the two trials they participated in (see Table 1).

The distribution of scores in the object labeling condition was significantly different from those in the causal condition, $\chi^2(1, N = 64) = 6.72$, $p < .03$. The proportion of endorsements of majority testimony over the minority informant's was significantly greater in the object labeling trials (49/64) than in causal trials (32/64), $\chi^2(1, N = 128) = 8.61$, $p < .003$ (see Figure 2b). These results closely match

our model's predictions. There were no significant differences in responses based on gender or age (younger vs. older than mean age).

Table 1: Participant scores by condition.

Condition	Score		
	0	*1*	*2*
Exp 1: Object labeling task	4	7	21
Exp 1: Causal task	13	6	13
Exp 2: Causal task (no feedback)	2	3	12

In the causal condition, participants were invited to activate each toy three times. Not all participants made three attempts, but collectively, participants made a total of 179 attempts to activate toys in the causal trials. Ninety-four attempts (53% of total attempts) were actions performed by majority informants, 84 attempts (47% of total attempts) were actions performed by the minority informant, and one attempt was a novel action performed by none of the informants. All participants in this condition attempted at least one action performed by the minority informant.

As predicted by our model, these data show children were more likely to endorse majority testimony when learning socially constructed facts (object labels) than non-socially constructed facts (cause-and-effect relationships).

Though there was no formal coding scheme for children's spontaneous comments during the study sessions, anecdotal evidence suggests that children's intuitions matched our model assumptions about mutual exclusivity. In the causal condition, children's comments suggested they accepted both the majority testimony and the minority informant's testimony ("Both [actions] make it go!"). Furthermore, all children in this condition attempted an action performed by a minority informant at least once in the study, suggesting that children were open to multiple possibilities when learning about cause and effect. In the object labeling condition, however, several children expressed the belief that there was only one correct answer ("That one isn't the *modi*!" about the minority-endorsed object).

A possible alternative explanation of these results is that children in the causal condition did not use information gained through their observations or informants' testimony; rather, they were simply confused by the task and randomly imitated informants' responses. To rule out this possibility, we designed another causal condition in which we expected children to endorse the majority testimony.

Experiment 2: The effect of feedback

In the causal condition of Experiment 1, children indiscriminately imitated the majority and minority informant actions, presumably because they were able to rely on their own observations, which suggested both demonstrated actions were equally effective at activating the toy. This second experiment examines how children behave when they do not have their own observations to rely on, but instead only have information from informants. We predict

that when children lack personal observations indicating the efficacy of informants' testimony, they will be more likely to endorse the majority's testimony.

Participants

Participants were 17 preschoolers, 7 male and 10 female (mean age = 4 years 4 months; range = 40 – 62 months). Participants were recruited in the San Francisco Bay Area by mail and phone calls or from preschools. An additional three children were tested, but excluded due to fussiness.

Materials and Procedure

The materials and general procedure of Experiment 2 were identical to those used in Experiment 1. The crucial difference between the two experiments was the content of the video clips participants watched. While in Experiment 1, children watched informants in the video clips perform actions that resulted in the toy playing music, in Experiment 2, informants in the video clips only mimed the actions they endorsed, and no music played as a result of miming the actions. In other words, children who viewed the Experiment 2 video clips received no information about the efficacy of the informants' testimony.

The script of the videos also differed from Experiment 1. Unlike the informants in Experiment 1, who did not verbally describe the action they performed, informants in the Experiment 2 video clips explicitly described their endorsed action and its hypothetical causal effect before miming the action, in order to provide context to children about why the action was being mimed: "It plays music if you pull the pink one!"

Results

As in Experiment 1, Experiment 2 participants were given a score of 0, 1, or 2 based on their first responses (see Table 1). The number of first responses endorsing majority informants' testimony was significantly higher in Experiment 2 than in the Experiment 1 causal condition, $\chi^2(1, N = 94) = 6.84, p < .008$, and was correctly predicted by our model (see Figure 2b).

Participants in Experiment 2 were also invited to attempt to activate each causal toy three times. Participants collectively made 99 attempts to activate the toys. Of those attempts, 64 (65% of total attempts) were actions performed by majority informants, and 35 (35% of total attempts) were actions performed by the minority informant. Unlike in Experiment 1, where all causal condition participants attempted at least one action performed by the minority informant, four of the 14 participants in Experiment 2 attempted only actions endorsed by the majority.

General Discussion

In this set of studies, we found that children do not indiscriminately endorse majority opinions; rather, their endorsement of majority opinions varies by task domain type and availability of alternate sources of knowledge. In

Experiment 1, children were significantly more likely to endorse majority testimony when learning about socially constructed facts (object labels) than non-socially constructed facts (causal relationships). Experiment 2 found that in the absence of information about the efficacy of informants' actions, children endorse majority testimony. Children's responses were predicted by a Bayesian model, suggesting that children make rational inferences from informants' testimony, and, when available, weigh other sources of information (e.g. personal observations) more heavily than testimony.

Though these results suggest children consider different sources of information in a non-socially-constructed domain, it is unknown whether they would do so in a socially constructed domain. In the causal conditions, the amount of feedback (i.e., hearing the toy play music) children received about actions' effects was easily quantifiable; however, it is less straightforward what would demonstrate positive or negative feedback about informants' endorsements in a object labeling condition. Future studies could explore how to convey feedback in an object labeling condition – perhaps showing successful or unsuccessful communication achieved through using the label – and the effect it would have on children's inferences.

Follow-up studies could also examine the effects of the informants' language in Experiment 2. In designing Experiment 2, we tried to make informants' video demonstrations as natural as possible while maintaining a similar script to Experiment 1, but creating parallel conditions proved difficult. Recall that the informants in Experiment 1 video clips performed their endorsed action without naming their actions. The informants in Experiment 2 narrated their actions and those actions' hypothetical effect ("It plays music if you pull the pink one!") to explain why informants were miming actions and to present possible actions for children to attempt at test. The hypothetical language used could imply to children that informants had prior experience with or knowledge about the toy. This prior knowledge, combined with informants' explicit demonstration, could be interpreted as evidence that informants were acting pedagogically, or upholding a social norm. Future work could examine how much of Experiment 2's effect was driven by children's lack of access to personal observations, versus pedagogical effects or social norm adherence.

Another difference between the conditions is the number of objects used in object labeling and causal conditions. In each of the causal conditions, participants saw the informants perform one of two actions on a single object, but the object labeling condition, participants saw informants call one of two objects by a novel label. A follow-up study to the object labeling condition could feature video clips with informants calling one object by two names, so that procedures of the object labeling and causal conditions would be more parallel.

In Experiment 1, we found that children's endorsement of majority testimony varies by task domain, but the flexibility

with which children incorporate conflicting information has yet to be determined. Accepting majority testimony as universally informative could potentially mislead a learner; individual members of the majority opinion could be mistaken, or the majority opinion as a whole could be flawed (see Esser, 1998 for a review on groupthink). Future work can identify the cues used to identify a reliable or unreliable majority. Children could discount informant testimony for rational reasons – for instance, if an informant is unreliable or unknowledgeable – or for less rational reasons – for instance, bias against out-group members (Kinzler & Spelke, 2011).

Additionally, the nature of the beliefs underlying children's endorsements has yet to be examined. Children may endorse majority testimony in the moment in order to conform to societal norms, but do not truly believe this testimony correct. In social psychology, this is called compliance. It is also possible that children internalize this new social knowledge and believe it to be true. Future studies could explore whether children are merely complying with social norms in similar object labeling tasks, or whether they internalize the majority's testimony. Children could be asked to teach others the names of objects, or to recall objects' novel labels in sessions hours or days later.

Overall, the similarity between our model's predictions and empirical data suggest that young children are discerning when considering others' testimony; the extent to which they prefer majority testimony is dependent on task domain type. This work also adds to the growing body of literature that suggests children consider information from multiple sources to make rational inferences.

References

Asch, S. E. (1956). Studies of independence and conformity: a minority of one against a unanimous majority. *Psychological Monograms, 70*(9).

Birch, S. A. J., Vauthier, S. A., & Bloom, P. (2008). Three-and-four-year-olds spontaneously use others' past performance to guide their learning. *Cognition, 107*(3), 1018-1034.

Bonawitz, E.B, & Lombrozo, T. (2012) Occam's Rattle: Children's use of simplicity and probability to constrain inference. *Developmental Psychology, 48*, 1156-1164.

Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academies of Sciences, 108*(2), 10918-10925.

Brugger, A., Lariviere, L. A., Mumme, D. L., & Bushnell, E. W. (2007). Doing the right thing: Infants' selection of actions to imitate from observed event sequences. *Child Development, 78*(3), 806-824.

Buchsbaum, D., Gopnik, A., Griffiths, T. L., & Shafto, P. (2011). Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition, 120*(3), 331-340.

Buchsbaum, D., Bridgers, S., Whalen, A., Seiver, E., Griffiths, T. L., & Gopnik, A. (2012). Do I know that you know what you know?

Modeling testimony in causal inference. Proceedings of the 34th Annual Conference of the Cognitive Science Society.

Cialdini, R. B., & Goldstein, N. J. (2004). Social Influence: Compliance and Conformity**. *Annual Review of Psychology, 55*, 591-621.

Clark, E. V. (1988). On the logic of contrast. *Journal of Child Language, 15*, 317–335.

Clark, E. V. (1990). On the pragmatics of contrast. *Journal of Child Language, 17*, 417–431.

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. (2009). Choosing your informant: weighing familiarity and recent accuracy. *Developmental Science, 12*(3), 426-437.

Esser, J. K. (1998). Alive and well after 25 years: a review of groupthink research. *Organizational Behavior and Human Decision Processes, 73*(2–3), 116–141.

Gergely, G., Bekkering, H. & Kiraly, I. (2002). Rational imitation in preverbal infants. *Nature, 415*, 755.

Kelman, H. C. (1958). Compliance, Identification, and Internalization: Three Processes of Attitude Change. *The Journal of Conflict Resolution, 2*(1), 51-60.

Kinzler, K.D., & Spelke, E.S. (2011). Do infants show social preferences for people differing in race? *Cognition*, 119, 1-9.

Koenig, M., Clement, F., & Harris, P. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science, 15*(10), 694-698.

Lyons, D. E., Young, A. G., & Keil, F. C. (2007). The hidden structure of overimitation. *Proceedings of the National Academy of Sciences, 104*(50), 19751.

Markman, E. M. (1989). Categorization and Naming in Children: Problems of Induction. MIT Press.

McGuigan, N., & Whiten, A. (2009). Emulation and "overemulation" in the social learning of causally opaque versus causally transparent tool use by 23- and 30-month-olds. *Journal of Experimental Child Psychology, 104,* 367–381.

Nolan, J. M., Schultz, P. W., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2008). Normative social influence is underdetected. *Personality and Social Psychology Bulletin, 34*(7), 913-923.

Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology, 43*(5), 1215-1226.

Sabbagh, M. A., & Baldwin, D. A., (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development, 72*(4), 1054-1070.

Schulz, L. E., Hooppell, C., & Jenkins, A. C. (2008). Judicious imitation: Children differentially imitate deterministically and probabilistically effective actions. *Child Development, 79*(2), 395–410.

Schulz, L.E., & Sommerville, J. (2006). God does not play dice: Causal determinism and preschoolers' causal inferences. *Child Development, 77*(2), 427-442.

Taylor, M., Cartwright, B. S., and Bowden, T. (1991). Perspective taking and theory of mind: Do children predict interpretive diversity as a function of differences in observers' knowledge? *Child Development, 62*(6), 1334-1351.

Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28.

Warneken, F., & Tomasello, M. (2009). Helping and cooperation at 14 months of age. *Infancy*, 11(3), 271-294.