Children and adults differ in their strategies for social learning

Falk Lieder[1] (falk.lieder@berkeley.edu)
Helen Wills Neuroscience Institute
University of California at Berkeley

Zi Lin Sim[1] (zi@berkeley.edu)
Department of Psychology
University of California at Berkeley

Jane C. Hu (janehu@uw.edu)
Institute of Learning & Brain Science
University of Washington

Thomas L. Griffiths (tom_griffiths@berkeley.edu)
Department of Psychology
University of California at Berkeley

Fei Xu (fei_xu@berkeley.edu)
Department of Psychology
University of California at Berkeley

[1] These authors contributed equally.

Abstract

Adults and children rely heavily on other people's testimony. However, domains of knowledge where there is no consensus on the truth are likely to result in conflicting testimonies. Previous research has demonstrated that in these cases, learners look towards the majority opinion to make decisions. However, it remains unclear how learners evaluate social information, given that considering either the overall valence, or the number of testimonies, or both may lead to different conclusions. We therefore formalized several social learning strategies and compared them to the performance of adults and children. We find that children use different strategies than adults. This suggests that the development of social learning may involve the acquisition of cognitive strategies.

Keywords: cognitive development; social learning; decision-making; Bayesian models; probabilistic reasoning

Introduction

Many of our beliefs, such as "The Earth revolves around the Sun," are based on other people's testimony rather than direct experience. Learning from others is even more critical for beliefs that are neither true nor false (Gelman, 2009; Harris & Koenig, 2006); this includes cultural beliefs and practices, as well as preferences (e.g., movies, foods). In all of these domains we are likely to encounter multiple conflicting testimonies. So how do we determine whom to believe?

Research on persuasion and consumer choice has demonstrated that "consensus information is an important driver of attitudes and behavior" (Benedicktus, Brady, Darke, & Voorhees, 2010). In other words, we often look towards the majority to guide our beliefs and decisions. Classic studies have shown strong effects of group consensus when adults are asked to make public judgments (Asch, 1956; Sherif, 1936), as well as the influence of word-of-mouth on consumer choice (Katz & Lazarsfeld, 1955). The popularity of rating websites (e.g., yelp.com, tripadvisor.com) shows that people deliberately seek out social recommendations to make choices (Lim & Van Der Heide, 2014).

According to the idea of natural pedagogy children efficiently learn from similar social information (Csibra & Gergely, 2009, 2011). Recent empirical work confirmed that children learn very well from other people's testimony: for example, when an experimenter shows 3 year-olds one way to play a game they readily consider it a general norm (Rakoczy, Warneken, & Tomasello, 2008).

But children do not learn indiscriminately. Just like adults, they are sensitive to the agreement and disagreement among informants, and they learn from the majority. In Corriveau, Fusaro, and Harris (2009), preschoolers watched a group of adults being asked to point to the referent of an unfamiliar word. When the children were later asked about the novel word, they relied on the opinion of the majority. This effect is not limited to word learning; it has been demonstrated for tool use as well (Haun, Rekers, & Tomasello, 2012).

While it is now well-established that adults and children often rely on the majority opinion when learning from others' testimonies, the mechanism underlying how they integrate the conflicting information inherent in multiple testimonies remains unclear. What factors do learners take into account, and how do they do so? Two prominent factors in the evaluation of multiple testimonies are the *overall valence* (e.g. the average number of stars) and the *number* of testimonies. The overall valence provides a quick and easy indicator of what the consensus is. The following example illustrates why the number of testimonies is also important: Suppose you want to choose between Restaurant A, which has 2 ratings with an average rating of 100%, and Restaurant B, which has 300 ratings with an average rating of 90%. If you only considered the overall valence of the testimonies, Restaurant A would seem to be the better option. However, that doesn't take into account the fact that you simply have more information about Restaurant B – and that people consistently consider it to be good. Ultimately, making a decision requires weighing both the overall valence and the number of ratings. How do children and adults reconcile these two sources of information?

Previous studies with adults focused on the effects of overall valence (Benedicktus, 2011; Benedicktus & Andrews, 2006; Benedicktus et al., 2010; Jimnez & Mendoza, 2013; Lim & Van Der Heide, 2014), and in previous studies with children, the overall valence and the number of ratings have always been confounded in the experimental design: the option with the higher overall valence always had the higher number of informants (Corriveau et al., 2009; Corriveau & Harris, 2010; Haun et al., 2012). We thus examine how children and adults reconcile conflicts between these two types of information to learn effectively from others. In addition, we investigate whether children use different learning strategies than adults. A difference between child and adult strategies

would suggest that cognitive development includes discovering cognitive strategies in addition to acquiring knowledge and representations.

Social Learning Strategies

To characterize how children versus adults learn from other people's ratings, we formalize five social learning strategies using probabilistic models and determine which model best explains children's versus adults' choices. All of these strategies provide a way to make a choice based on the same pieces of information: the choice is between two unknown options, one with $n_{1,+}$ positive and $n_{1,-}$ negative ratings and another with $n_{2,+}$ positive and $n_{2,-}$ negative ratings. In brief, the five strategies we consider are to choose the option with more positive ratings (strategy 1, *Count+*), choose the option with the highest proportion of positive ratings (strategy 2, *Ratio*), apply a rational sampling strategy that approximates Bayesian inference (BI) with an optimistic or neutral prior (strategies 3-4; *BI*), or choose at random (strategy 5).

Apart from choosing at random, the first of these strategies is the simplest and the most frugal in its use of cognitive resources: It ignores the negative ratings and prior knowledge about the distribution of ratings. It merely compares the number of positive ratings. Strategy 2, by contrast, also takes into account the number of negative ratings to compute and compare the proportions of positive ratings. Strategies 3 and 4 augment the observed numbers of ratings by prior knowledge. These strategies integrate the fraction of positive ratings (valence) with the number of ratings. This enables strategies 3 and 4 to choose an option with 99 positive and only 1 negative rating over an alternative with 1 positive and 0 negative ratings, whereas strategy 2 would choose the option with only one positive rating.

All models take into account that people's perceptual uncertainty σ_n about how many items there are increases linearly with their numerosity n:

$$\sigma_n = w \cdot n, \qquad (1)$$

where the slope w is known as the Weber fraction (Halberda & Feigenson, 2008). Adults perceive numbers more precisely than children. This difference has been quantified in terms of their Weber fractions (Halberda & Feigenson, 2008): Adults' Weber fraction is $w_{\text{adults}} = 0.108$. By contrast, five-year-olds' Weber fraction is $w_{5y} = 0.229$ and six-year-olds' Weber fraction is $w_{6y} = 0.179$. To model the responses of 5- and 6-year-olds we therefore use the average of their Weber fractions, that is $w_{\text{children}} = 0.21$.

We derived how people's choices would be distributed if they used either of the five strategies. If people used strategy 1 (model m_1), then the probability of choosing the first option ($P(C = 1)$) would be determined by the perceived difference between the two options' numbers of positive ratings:

$$P(C = 1|m_1) = \Phi\left(\beta \cdot \frac{n_{1,+} - n_{2,+}}{w \cdot \sqrt{n_{1,+}^2 + n_{2,+}^2}}\right), \qquad (2)$$

where the denominator is the perceptual uncertainty about the difference in the number of positive ratings, Φ is the cumulative distribution function of the normal distribution, and β is a free parameter that determines choice variability. Our model's prior distribution on this parameter is the standard uniform distribution, that is $P(\beta) = \text{Uniform}([0,1])$.

If people used strategy 2, the choice probability would be determined by the perceived difference between the proportions of positive ratings (i.e. $r_1 - r_2$):

$$P(C = 1|\beta, m_2) = \Phi\left(\beta \cdot \frac{\mathbb{E}[R_1|\mathbf{n}_1] - \mathbb{E}[R_2|\mathbf{n}_2]}{\sqrt{\text{Var}(R_1|\mathbf{n}_1) + \text{Var}(R_2|\mathbf{n}_2)}}\right), \qquad (3)$$

where R_1 and R_2 are the perceived fraction of positive ratings of option 1 and option 2 respectively, $\mathbf{n}_1 = (n_{1,+}, n_{1,-})$, $\mathbf{n}_2 = (n_{2,+}, n_{2,-})$, and the prior on β is again $\text{Uniform}([0,1])$.

Strategies 3 and 4 instantiate the sampling hypothesis (Bonawitz, Denison, Griffiths, & Gopnik, 2014; Denison, Bonawitz, Gopnik, & Griffiths, 2013). According to the sampling hypothesis, the variability of people's judgments reflects approximate Bayesian inference by sampling. From a Bayesian perspective, the goal of preference formation is to infer whether the probability μ_1 that option 1 will generate a positive experience is greater than the probability μ_2 that option 2 will generate a positive experience (i.e. $\mu_1 > \mu_2$). This inference integrates the numbers of ratings $\mathbf{n} = (n_{1,+}, n_{1,-}, n_{2,+}, n_{2,-})$ with prior knowledge about the overall proportion of positive ratings. We model this prior knowledge by a Beta distribution with mean μ and strength τ:

$$\mu_1, \mu_2 | \mu, \tau \sim \text{Beta}(\alpha = \mu \cdot \tau, \beta = (1 - \mu) \cdot \tau). \qquad (4)$$

Under these assumptions, the sampling hypothesis predicts that people's choices C are samples from the posterior distribution on which of the two options is better:

$$P(C = 1|\theta) = P(\mu_1 > \mu_2|\mathbf{n}, \theta) + \frac{1}{2} \cdot P(\mu_1 = \mu_2|\mathbf{n}, \theta), \qquad (5)$$

where the model parameters $\theta = \{\mu, \tau\}$ are the mean μ and strength τ of the prior belief on the frequency of positive ratings. We consider two Bayesian sampling strategies: While strategy 3 assumes that most ratings are positive ($P(\mu|m_3) = \text{Uniform}([0.5, 1])$), strategy 4 does not rely on this assumption ($P(\mu|m_4) = \text{Uniform}([0, 1])$). Both models share a uniform prior on the scale of the strength τ of people's prior knowledge: $p(\log(\tau)) \propto 1$ for $0.1 \leq \tau \leq 10$ and 0 else.

Finally, if people used strategy 5 they would choose both options with equal probability:

$$P(C = 1|m_5) = P(C = 2|m_5) = 0.5. \qquad (6)$$

We used these probabilistic models of the five strategies to design an experiment that can discriminate between them.

Experiment 1: Comparing children and adults

The goal of Experiment 1 was to compare children's versus adults' strategies for learning from other people's ratings.
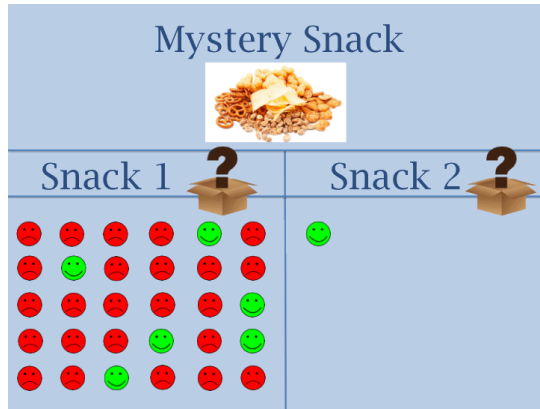
Figure 1: Example trial for distinguishing the *Count+* strategy from Bayesian inference and the ratio strategy.

Table 1: Design and results of Experiment 1.

Trial	Option 1		vs.	Option 2		Choice Freq. Option 2		
	#Pos.	#Neg.		#Pos.	#Neg.	Children	Adults	p
1	6	24		1	0	49%	75%	0.0001
2	5	25		1	0	57%	74%	0.0078
3	9	21		2	0	59%	73%	0.0255
4	8	22		2	0	64%	79%	0.0108
5	1	0		29	1	78%	82%	0.4034
6	1	0		28	1	78%	90%	0.0088
7	0	1		6	24	73%	50%	0.0011
8	1	0		24	6	70%	88%	0.0002
9	7	23		1	0	57%	64%	0.3159
10	8	22		1	0	55%	71%	0.0135
11	2	28		0	1	25%	60%	0.0000
12	2	27		0	1	32%	68%	0.0000
13	30	0		0	30	3%	6%	0.4318
14	10	5		5	10	14%	10%	0.3340
15	0	15		0	2	13%	3%	0.0013
16	15	0		5	0	6%	7%	0.9426
17	15	0		0	30	65%	85%	0.0002
18	30	0		0	15	17%	19%	0.8113

Methods

Participants. Sixty-six English-speaking 4- to 6-year-olds (33 boys and 33 girls; 10 4-year olds, 38 5-year olds, and 18 6-year-olds) and 300 adult participants were tested. Children were recruited from schools and museums surrounding Berkeley, CA. Their average age was 67.5 months (range = 48.9 to 81.6 months). Two additional children were tested but excluded due to experimenter error. Three hundred adult participants were recruited on Amazon Mechanical Turk and paid $0.20 for about three minutes of work. Twelve adults were excluded because of incomplete data or color blindness.

Materials and Procedure. Eighteen pairs of ratings were selected for use in this experiment; see Table 1. The first twelve pairs of ratings were designed to be maximally informative with regards to differentiating between strategies 1–4. The remaining 6 paired ratings were questions designed to ensure that participants understood the task instructions and discriminate strategies 1–4 from random choice (strategy 5). Child participants were tested individually in front of a laptop. An experimenter recorded their responses for each trial. Adult participants were tested through an online survey and their responses were recorded automatically. On each trial, participants viewed an image consisting of two unknown items (termed "mystery items") belonging to one of six domains: books, games, movies, snacks, stickers, and stuffed animals. Each item was presented with its associated binary ratings, presented as smiling and frowning faces (see Figure 1). Participants were told, "Some people played with Item 1, and some people played with Item 2. A green smiley face means the person liked that item, and a red frowny face means the person did not like that item." Upon presenting the ratings for both items, participants were asked, "Which item do you think is better? Item 1 or Item 2?" and their responses were recorded. The order of the trials was randomized, and the domain used on each trial and the placement of the smiling/frowning faces were counterbalanced.

Results and Discussion

The frequencies with which children versus adults chose either option are summarized in Table 1. To estimate the proportion of children and adults using each of the five strategies defined above, we performed random-effects Bayesian model selection at the group level separately for children versus adults (Stephan, Penny, Daunizeau, Moran, & Friston, 2009). This analysis performs hierarchical Bayesian inference on the strategy used by each participant and the relative frequency of each strategy within the population. Figure 2 summarizes our results. We found that a significantly larger proportion of adults than children chose according to Bayesian inference (67.1% vs. 39.9%, $p < 0.0001$). We can be 99.99% confident that, among adults, Bayesian inference with a neutral prior (strategy 4) is the most common of the five strategies. By contrast, the proportion of children using this strategy is not significantly larger than the proportion of children using strategy 1 ($p = 0.19$). Conversely, more children than adults used strategy 1 (31.0% vs. 13.1%, $p < 0.001$) or the ratio strategy (4.84% vs. 0.36%, $p < 0.03$).

We re-analyzed the data under the assumption that each person might use multiple strategies, and the results led to the same conclusion. The difference between children and adults was most pronounced in trial types 11 and 12 (see Table 1). In trial 11, about 75% of our child participants preferred the option with 2 positive and 28 negative ratings to the option 0 positive and 1 negative rating, whereas Bayesian inference and more than 70% of adults prefer the second option; and similarly for trial 12.

While the experiment discriminated the strategies considered so far, people might use other strategies such as choosing the option with the highest difference between the numbers of positive ratings and negative ratings. In the present experiment this strategy would have lead to exactly the same decisions as performing Bayesian inference with the neutral prior for all but two choices. To resolve this ambiguity, we designed a second experiment.
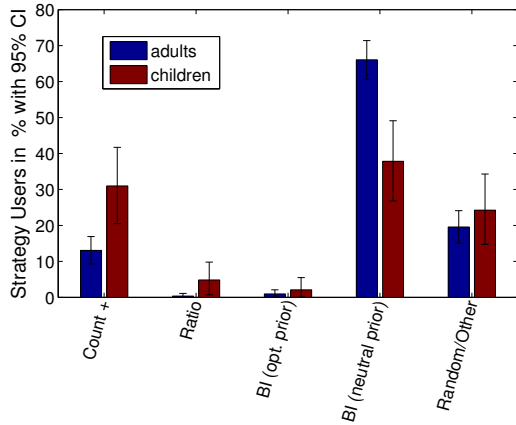
Figure 2: Proportion of children versus adults forming preferences according to Bayesian inference or heuristic strategies.

Experiment 2: Testing more models with adults

The goal of Experiment 2 was to test if adults use a rational sampling strategy or a simple heuristic that Experiment 1 could not distinguish from Bayesian inference.

Methods

Modeling alternative strategies and experimental design. Our first step in designing the experiment was to formalize four alternative strategies by probabilistic computational models. According to the difference strategy (m_6), people stochastically choose the option for which the *surplus of positive ratings* (i.e. the difference between the number of positive ratings and the number of negative ratings) is largest:

$$P(C = 1 | \beta, m_6) = \Phi\left(\beta \cdot \frac{(n_{1,+} - n_{1,-}) - (n_{2,+} - n_{2,-})}{w \cdot \sqrt{n_{1,+}^2 + n_{2,+}^2 + n_{1,-}^2 + n_{2,-}^2}} \right),$$

where the numerator is the difference between the first option's and the second option's surplus of positive ratings and the denominator is people's perceptual uncertainty about this quantity. The prior over the choice variability parameter is again uniform on the interval from zero to one. Strategy 7 probabilistically chooses the option with less negative ratings:

$$P(C = 1 | m_7) = \Phi\left(\beta \cdot \frac{n_{2,-} - n_{1,-}}{w \cdot \sqrt{n_{1,-}^2 + n_{2,-}^2}} \right). \quad (7)$$

The prior on the choice variability parameter is again $P(\beta) =$ Uniform($[0, 1]$). Strategy 8 determines if there is an option that has only positive ratings. If exactly one such option exists, strategy 8 chooses it, else it picks one of the options at random. Strategy 9 chooses the only option that has no negative ratings, if one exists, and otherwise chooses at random.

We used these probabilistic models to select trials that jointly discriminate between all 45 pairs of the nine strategies. For each pair of strategies, we picked at least two trials

Table 2: Design and results of Experiment 2.

Trial	Option 1 Ratings #Pos.	#Neg.	Option 2 Ratings #Pos.	#Neg.	Choice Freq. Option 2 Adults
1	6	24	1	0	85%
2	9	21	2	0	92%
3	1	0	29	1	97%
4	0	1	6	24	59%
5	7	23	1	0	83%
6	2	28	0	1	53%
7	30	0	0	30	1%
8	10	5	5	10	0%
9	3	11	9	21	87%
10	10	3	21	9	35%
11	8	22	18	12	99%
12	3	3	7	23	10%
13	22	8	9	1	65%
14	1	10	7	23	96%
15	9	1	24	6	38%
16	1	1	5	25	10%
17	4	26	1	1	91%
18	3	1	9	21	8%
19	7	23	2	1	93%
20	0	0	15	15	80%
21	16	14	0	0	15%
22	1	0	27	1	96%
23	26	1	1	0	6%

for which the two strategies make opposite predictions, and the predictions differ by at least 20%.

Participants. One hundred and sixty participants were recruited on Amazon Mechanical Turk. Each participant was paid $0.20 for about three minutes of work.

Materials and Procedure. The experiment comprised the 23 trials shown in Table 2. The trial order was randomized between subjects and two spatial arrangements of the ratings were counterbalanced.

Results and Discussion

Four participants were excluded due to incomplete data. The choice frequencies of the remaining participants are summarized in Table 2.

We analyzed our participants individual choices using random effects Bayesian model selection as described above. According to this analysis, 77.5% of adults use the Bayesian sampling strategy with a neutral prior (95% CI: 71.1% − 83.7%); see Figure 3. We can be more than 99.99% confident that this is the most frequently used strategy among adults. By contrast, the difference strategy is used by only about 8% of adults (95% CI: 4.90% − 13.05%), and only 6% use strategy 1 (95% CI: 2.66 − 9.69%). We re-analyzed the data under the assumption that each person might use multiple strategies and the results led to the same conclusion. In conclusion,
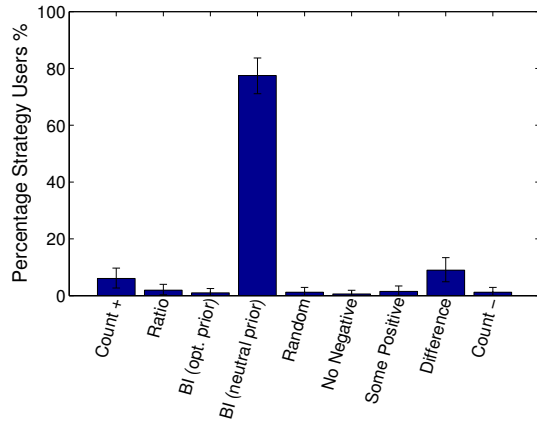
Figure 3: Proportion of adults using each of the nine social learning strategies according to Experiment 2.

we found that adults' choices are indeed best explained by Bayesian inference. This consolidates our interpretation of the results of Experiment 1.

General Discussion

Both adults and children need to learn from the testimony of others, but may use different strategies to make sense of that testimony. In Experiment 1 we found that the judgments of adults were consistent with Bayesian inference, while children seemed to rely on simpler heuristics more often than adults. A more detailed examination of adult performance in Experiment 2, using a wider range of alternative models, provided further support for the conclusion that adults' choices are consistent with a Bayesian sampling model. Our results raise interesting questions about the development of probabilistic reasoning.

Superficially, our results appear to contradict previous findings that infants' inferences conform to the principles of Bayesian inference (Xu & Kushnir, 2013) and that infants' decisions reflect those rational inferences (Denison & Xu, 2010). Note, however, that Xu and Kushnir (2013) describe modal or average responses across all participants and a small number of inference problems, demonstrating that behavior *at the group level* can be characterized as consistent with principles of Bayesian inference, whereas Experiment 1 analyzed the responses of individual participants across a large number of inference problems to determine strategy use *at the individual level*. Moreover, most previous infant studies have not contrasted specific competing models. By considering a wide range of problems we found that more than 50% of 5-6 year-old children employ simple heuristics rather than a sampling strategy to learn from statistical social information. Although the inferences drawn by the heuristic strategies often coincide with Bayesian inference, Experiment 1 included trials in which they do not and neither did the choices of our child participants.

These results suggest that most children use simple strategies whose inferences will sometimes (but not always) coincide with those of Bayesian inference. However, the conclusion that children do not use a rational sampling strategy should be interpreted with caution, as we cannot be sure that our task analysis is correct. For instance, we assumed that the total number of ratings is independent of the object's quality, but participants might assume that good options receive more ratings. Adults might be better able to override this intuitive (but wrong) interpretation in our experiment because their executive functions are superior to those of children.

The differences between children and adults in Experiment 1 raise the question how the rational sampling strategy is learned and how and when children transition to using it. Is this strategy domain-general or domain-specific, and are there individual differences with regards to its uptake? Our results suggest the following answers to these questions (Bonawitz et al., 2014): First, cognitive strategies for reasoning under uncertainty or the mechanism by which they are chosen develop during childhood. According to our data, the transition to a rational sampling strategy is still incomplete at age 6. Thus either there is no innate rational sampling strategy, or it takes children more than six years to learn when to use it. Second, rather than using a single, universal sampling strategy, children appear to use multiple specialized strategies. The strategies children use for some problems, such as inferring others' preferences in the experiments reviewed by Xu and Kushnir (2013) and the causal inference tasks used by Denison et al. (2013), accomplish sampling, but the strategies they used to make choices based on conflicting social information in our task did not. Third, our results indicate qualitative inter-individual differences: While some children appeared to use a rational sampling strategy, others appeared to rely solely on the number of positive ratings. Similarly, while most adults employed a rational sampling strategy, others appeared to use simpler heuristics. Future experiments will rigorously test alternative explanations of children's responses and zoom in on the strategic development during childhood.

The interpretation that children's social learning strategies change during development is consistent with Siegler's observations that children discover new arithmetic strategies and gradually transition to using the most effective strategy (Siegler, 1999). Such findings illustrate that cognitive development is not limited to the acquisition of knowledge and representations but also includes learning cognitive strategies. While there has been substantial progress in characterizing the acquisition of knowledge and representations using rational models of cognition (Xu & Griffiths, 2011), theoretical frameworks for the development of cognitive strategies and strategy selection are still in their infancy (Shrager & Siegler, 1998; Lieder et al., 2014; Lieder & Griffiths, 2015). Developing a rational framework for the development of cognitive strategies and the metacognitive mechanisms of strategy selection is an important goal for future research.

collecting the children's data. Part of this work will be presented as a poster at the Biennial Meeting of the Society for Research in Child Development in March 2015.

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.

Benedicktus, R. L. (2011). The effects of 3rd party consensus information on service expectations and online trust. *Journal of Business Research*, *64*(8), 846–853.

Benedicktus, R. L., & Andrews, M. L. (2006). Building Trust with Consensus Information : The Effects of Valence and Sequence Direction. *Journal of Interactive Advertising*, *6*(1), 17–29.

Benedicktus, R. L., Brady, M. K., Darke, P. R., & Voorhees, C. M. (2010). Conveying Trustworthiness to Online Consumers: Reactions to Consensus, Physical Store Presence, Brand Familiarity, and Generalized Suspicion. *Journal of Retailing*, *86*(4), 310–323.

Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends in Cognitive Sciences*, *18*(10), 497–500.

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.

Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in cognitive sciences*, *13*(4), 148–53.

Csibra, G., & Gergely, G. (2011). Natural pedagogy as evolutionary adaptation. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *366*(1567), 1149–57.

Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The sampling hypothesis. *Cognition*, *126*(2), 285–300.

Denison, S., & Xu, F. (2010). Twelve- to 14-month-old infants can predict single-event probability with large set sizes. *Developmental Science*, *13*(5), 798–803.

Gelman, S. A. (2009). Learning from others: Children's construction of concepts. *Annual review of psychology*(60), 115–140.

Halberda, J., & Feigenson, L. (2008). Developmental change in the acuity of the "number sense": The approximate number system in 3-, 4-, 5-, and 6-year-olds and adults. *Developmental Psychology*, *44*(5), 1457–1465.

Harris, P. L., & Koenig, M. A. (2006). Trust in testimony: How children learn about science and religion. *Child development*, *77*(3), 505–524.

Haun, D. B. M., Rekers, Y., & Tomasello, M. (2012). Majority-biased transmission in chimpanzees and human children, but not orangutans. *Current Biology*, *22*(8), 727–731.

Jimnez, F. R., & Mendoza, N. A. (2013). Too popular to ignore: The influence of online reviews on purchase intentions of search and experience products. *Journal of Interactive Marketing*, *27*(3), 226 - 235.

Katz, E., & Lazarsfeld, P. (1955). *Personal influence*. New York, NY: Free Press.

Lieder, F., & Griffiths, T. L. (2015). When to use which heuristic: A rational solution to the strategy selection problem. In *Proceedings of the 37th annual conference of the cognitive science society.* Austin, TX: Cognitive Science Society.

Lieder, F., Plunkett, D., Hamrick, J. B., Russell, S. J., Hay, N., & Griffiths, T. (2014). Algorithm selection by rational metareasoning as a model of human strategy selection. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (Eds.), *Adv. neural inf. process. syst. 27* (pp. 2870–2878). Curran Associates, Inc.

Lim, Y.-S., & Van Der Heide, B. (2014). Evaluating the wisdom of strangers: The perceived credibility of online consumer reviews on Yelp. *Journal of Computer-Mediated Communication*. doi: 10.1111/jcc4.12093

Rakoczy, H., Warneken, F., & Tomasello, M. (2008). The sources of normativity: young children's awareness of the normative structure of games. *Developmental psychology*, *44*(3), 875–881.

Sherif, M. (1936). *The psychology of social norms.* New York, NY: Harper.

Shrager, J., & Siegler, R. S. (1998). SCADS: A model of children's strategy choices and strategy discoveries. *Psychological Science*, *9*(5), 405–410.

Siegler, R. S. (1999). Strategic development. *Trends in Cognitive Sciences*, *3*(11), 430–435.

Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, *46*, 1004–1017.

Xu, F., & Griffiths, T. L. (2011). Probabilistic models of cognitive development: Towards a rational constructivist approach to the study of learning and development. *Cognition*, *120*(3), 299 - 301.

Xu, F., & Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science*, *22*(1), 28–32.