

Is that your final answer? The effects of neutral queries on children's choices

Aaron Gonzalez¹, Patrick Shafto², Elizabeth Bonawitz¹, & Alison Gopnik¹

¹Department of Psychology, University of California, Berkeley

² Department of Psychological and Brain Sciences, University of Louisville

Abstract

Preschoolers often switch a response on repeated questioning, even though no new evidence has been provided (Krahenbuhl, Blades, & Eiser, 2009). Though apparently irrational, this behavior may be understood as children making an inductive inference based on their beliefs about whether initial responses were correct and the knowledgeability of the questioner. We present a probabilistic model of how the questioners' knowledge and biases to be positive should affect inferences. The model generates the qualitative prediction that an ideal learner should switch responses more often following a "neutral query" from a knowledgeable questioner than following queries from an ignorant questioner. We test predictions of the model in an experiment. The results show that four-year-old children are sensitive to questioners' knowledge when responding to a neutral query, demonstrating more switching behavior when the query is provided by a knowledgeable questioner. We conclude by discussing the practical and theoretical implications for cognitive development.

When should a learner abandon their current hypothesis in favor of a new one? It is becoming clear that even preschool-aged children rationally update beliefs and generate new explanations following informative evidence (Gopnik, Glymour, Sobel, Schulz, & Danks, 2004; Schulz, Bonawitz, & Griffiths, 2007; Denison, Bonawitz, Gopnik, & Griffiths, 2010; Bonawitz, Schijndel, Friel, & Schulz, 2012; Bonawitz, Fisher, & Schulz, in press). These tasks often involve assessing children's starting belief state, either presenting the child with new evidence or allowing the child to generate their own evidence, and then eliciting an explanation or judgment. What constitutes "evidence" in these tasks is fairly intuitive; for example, children may be presented with a storybook containing information about two variables that tend to co-occur (Schulz et al., 2007) or they may observe a toy that reacts when certain objects are placed on top of it (Bonawitz, Denison, Gopnik, & Griffiths, 2011).

In addition to revising beliefs following correlational evidence and interventions, children also learn from others. Interestingly, even preschool-aged children do not do so indiscriminately; they use information about informants' knowledge and intent to guide who they trust. For example, children do not trust informants who label familiar objects incorrectly (Koenig & Harris, 2005), who dissent from groups (Corriveau, Fusaro, & Harris, 2009), and even with familiar informants, children update social inferences (Corriveau & Harris, 2009). Similarly, recent research suggests that children use information about informants' intent to guide inferences (Mascaro & Sperber, 2009; Shafto, Eaves, Navarro, & Perfors, in press). Other research has suggested that even four-year-old children make different causal inferences depending on the social context when evidence is presented: direct instruction by a knowledgeable and helpful informant

provides strong constraints on likely hypotheses as compared to accidental information by a not-knowledgeable informant. Even when contrasted with intentional (but not instructional) actions and interrupted demonstrations, preschoolers make stronger inferences about causal structure from direct instruction by leveraging the informant's knowledge and intent (Bonawitz, Shafto, et al., 2011; Buchsbaum, Gopnik, Griffiths, & Shafto, 2011; Butler & Markman, 2010). These pedagogical assumptions have been captured by probabilistic models (e.g. Shafto & Goodman, 2008; Bonawitz, Shafto, et al., 2011), which suggest a rational account of how learners update their beliefs following evidence generated in the context of teaching.

These literatures suggest that preschool children are sophisticated and powerful social learners; they revise their beliefs when evidence is sufficient and use social or instructional contexts to help interpret the strength of the evidence. However, other research suggests that children may abandon hypotheses too capriciously. For example, the effects of repeated questioning on eyewitness testimony in young children have been studied extensively in the context of the judicial system, and work done by Poole and White (1991) found that, in contrast to adults, four-year-olds were more likely to change their responses to repeated yes or no questions. More recently Krahenbuhl et al. (2009) found that children as young as four changed over a quarter of their responses to repeated questions, resulting in a decline in accuracy. In Howie, Sheehan, Mojarrad, and Wrzesinska (2004) four-year-old children were more likely than older children to change responses toward an incorrect answer on repeated questioning. That is, although no additional evidence was provided, by simply asking children the same question a second time, children were very likely to switch their predictions. How might we reconcile these findings with those suggesting that children should only rationally update beliefs following informative evidence?

One explanation for this apparently irrational switching behavior is that seemingly neutral questions such as "Is that your final answer?" may provide strong information in certain social contexts. If preschoolers are sensitive to the potential communicative intent behind such queries, the question itself may be a source of evidence about whether an initial guess was correct. Consider a game in which a sticker is hidden under one of two cups. Suppose an informant asks the child which cup they believe the sticker is under and after the child's initial guess, the teacher asks, "Is that your final answer? Would you like to change your guess?" In which contexts does such a question provide information about the true location of the sticker? Intuitively, it seems obvious that if the

questioner does not know the actual location of the sticker, then the repetition of the question provides little additional evidence; however, a *knowledgeable* questioner might have a good reason for giving the learner with a second chance at answering the question. In these cases, this apparently neutral query is not neutral at all; it is a strong cue to the learner that they should change their answer.

In what follows we will explore the idea that even a “neutral” query carries information about the state of the world when a questioner is knowledgeable. We present a probabilistic model that demonstrates how an ideal learner might evaluate such “neutral” queries in scenarios in which the questioner is knowledgeable and scenarios in which she is ignorant. With the model, we evaluate the conditions under which switching guesses is the rational choice for the learner. We then test the basic prediction with preschoolers in an experiment in which the informant’s knowledge or ignorance is made explicit. We suggest that repeated questioning does indeed lead a learner to switch responses and that even preschoolers are sensitive to the knowledge state of others when making such inferences.

Modeling learners’ responses to neutral queries

Here we consider a model of “neutral queries”. Bayesian inference provides a natural framework in which to consider how an ideal learner should update her beliefs following this kind of information. In Bayesian inference, the learner’s goal is to update their beliefs about hypotheses, h , given data, d , where the degree of belief in a hypothesis given data is denoted $P(h|d)$. These updated posterior beliefs are determined by two factors: the learner’s prior beliefs in hypotheses, $P(h)$, and the probability of sampling the observed data, assuming the hypothesis is true, $P(d|h)$. Specifically, the updated posterior belief in a particular hypothesis is proportional to the product of the prior belief in that hypothesis and the probability of sampling the data given that hypothesis, $P(h|d) \propto P(d|h)P(h)$.

Because we are considering only two hypotheses, we can use Bayes Odds to simplify the problem:

$$\frac{P(h_1|d)}{P(h_2|d)} = \frac{P(d|h_1)P(h_1)}{P(d|h_2)P(h_2)}, \quad (1)$$

where $P(h_1|d)$ is the probability that the sticker is in the first location, given the statement from the informant (“correct”, “incorrect”, “is that your final answer”) and $P(h_2|d)$ is the probability that the sticker is in the second location given the statement. It is reasonable to assume that the learner has no a priori assumptions about either location, which allows us to cancel out the prior beliefs (i.e. $P(h_1) = P(h_2)$), thus $\frac{P(h_1)}{P(h_2)} = 1$.

The main issue is the probability of the statement given the location of the sticker, $P(d|h)$. We can model this likelihood with a simple causal graphical model (see Figure 1). Causal graphical models consist of a structure indicating the causal relationships among a set of variables, where nodes are




Figure 1: Graphical model depicting dependencies in cases when the informant is *knowledgeable* and *ignorant*.

variables and dependence relationships are indicated by arrows from causes to effects. To complete a graphical model, conditional probability distributions give the probability that each variable takes on a particular value given the value of its causes (Pearl, 2000; Spirtes, Glymour, & Schienens, 1993, see Table 1).

In our model of the problem, the variables include the actual state of the world (“World”, i.e. location of the sticker), the intention of the speaker (to provide helpful feedback, to avoid negative feedback, etc.), and the possible statements the informer can make (“Correct”, “Incorrect”, “Is that your final answer?”). In the first model (*Knowledgeable*) the informant is aware of the actual state of the world. As a result, both the true location of the sticker and the intention of the speaker influence the statement given to the learner. In the second model (*Ignorant*), the informant is not aware of the actual state of the world. As a result, the true state of the world does not influence the statement given to the learner. That is, the state of the world and the information provided are causally independent of each other.

The dependence assumptions captured by the graphical model generate predictions about the behavior of learners. In the case of a knowledgeable informant, given information about the actual state of the world and the informant’s statement, a learner can infer something about the informant’s goals (e.g. to provide positive or negative feedback). Given information about the state of the world and the informant’s goals, a learner could also predict (with some probability) the likelihood that the informant would produce different statements. In our problem, given the informant’s goals and the statement provided, the learner can make an inference about

Table 1: Conditional probability table for knowledgeable graphical model.

Guess	“Correct”	“Incorrect”	“Final answer?”
Correct	α_c	$\beta_c \approx 0$	$1 - (\alpha_c + \beta_c) \approx 1 - \alpha_c$
Incorrect	$\alpha_i \approx 0$	β_i	$1 - (\alpha_i + \beta_i) \approx 1 - \beta_i$

the state of the world. A learner could also make more abstract inferences: given information about the goals of the informant, the statement provided, and the actual location of the sticker, a learner could infer which model (*Knowledgeable* vs *Ignorant*) best captures the knowledge state of the informant.

The specification of the conditional probability distribution provides additional qualitative predictions. In the *knowledgeable* graphical model, we might reasonably argue that if the child chooses the correct location initially, the speaker is very unlikely to say “incorrect”; in this case $P(\text{“incorrect”}|\text{correct choice}) = \beta_c \approx 0$. Similarly, if the child chooses the incorrect location initially the speaker is very unlikely to say “correct”, $P(\text{“correct”}|\text{incorrect choice}) = \alpha_i \approx 0$. Given these intuitive assumptions, we can compare three possible ways biases about the goals of the teacher might play out in the model’s predictions.

The first possibility is that the informant is *unbiased*. Let us consider the case when the informant is knowledgeable. In this case, the *unbiased informant* is just as likely to say “correct” when the initial guess is correct as “incorrect” when the initial guess is incorrect, $\alpha_c \approx \beta_i$. If this is the case, then the learner can not infer whether their initial guess is correct or not if they hear the statement “is that your final answer”. This is because $P(\text{“final answer?”}|\text{correct}) = P(\text{“final answer?”}|\text{incorrect})$. That is, the statement “is that your final answer” provides no additional information about the location of the sticker (Equation 1 is approximately equal to 1). Now consider the case where the informant is ignorant. In this case, because the informant has no information about the actual state of the world, the true location is conditionally independent of the statements made by the informant, and the learner cannot make any inferences about the state of the world. Thus, assuming unbiased informants, learners should make the same inferences if asked “is that your final answer” in a *knowledgeable* condition as if asked “is that your final answer” in an *ignorant* condition. This model does not predict the degree to which the learners should switch responses, but it does predict no difference between conditions.

A second possibility is that the informant is *positively biased*. In this model, the knowledgeable informant may be inclined to want to say “correct” following correct initial guesses, but would be reluctant to say “incorrect” following an incorrect initial guess, $\alpha_c > \beta_i$. If this is the case, then the statement “is that your final answer” provides support for the hypothesis that the learner’s initial guess was incorrect because she is more likely to hear “is that your final answer” given an incorrect guess than “is that your final answer” given a correct guess (Equation 1 > 1). Thus, the *positively biased* model predicts that a learner should show increased switch-

ing in a *Knowledgeable* condition as compared to an *Ignorant* condition (in which the state of the world is still conditionally independent of the statements and thus does not provide additional information).

The third possibility is that the informant is *negatively biased*. In this model, the informant may be inclined to say “incorrect” following an incorrect initial guess, but would be comparatively reluctant to say “correct” following a correct initial guess, $\alpha_c < \beta_i$. If this is the case, then the statement “is that your final answer” provides support for the hypothesis that the learner’s initial guess was correct (Equation < 1) and the learner should show a decrease in switching responses in the *knowledgeable* condition as compared to an *Ignorant* condition.

Note that the precise values of α and β are not important for the predictions of this model, but the relationship between these variables drives the predictive differences. We investigate three implications of this model: first, do we replicate the finding that preschoolers tend to switch responses following what might be considered a “neutral query”; second, do preschoolers take the knowledge state of the informant into account when inferring whether or not to switch hypotheses; third, do preschoolers assume that the informant is neutral, positively, or negatively biased when they provide a query?

Experiment: Preschoolers’ switching behavior in response to a neutral query

To investigate the predictions of our model we invited preschoolers to participate in a game where the goal was to guess the location of a sticker under one of two cups. After their initial guesses, children were given some feedback and the opportunity to change their guess. After two training trials in which the experimenter told the child that their first guess was either correct or incorrect, children were given three test trials. In the test trials the experimenter asked the child “Is that your final guess?” after children’s initial guesses. Some children participated in a condition in which the experimenter looked under the cups before generating the query and others participated in a condition in which the experimenter did not look before the query. The critical measure is simply on what percentage of trials children switched their prediction to the other cup by condition.

Method

Participants Thirty-two preschoolers (mean age: 58.6 months; range: 48-79 months) were recruited from local preschools and museums for participation.

Design Preschoolers were randomly assigned to either the *Knowledgeable* condition or *Ignorant* condition. Four children were dropped and replaced for demonstrating a side bias (see results) in the *Knowledgeable* condition and five children were dropped and replaced for side biases in the *Ignorant* condition. There were no differences in ages between condition, $t(30) = 0.35, p = ns$.

Materials 5 pairs of colored cups (pink, blue, yellow, green, orange) were used in the conditions. An animal sticker was placed on the inside of one of the cups in each set.

Procedure The experimenter began by pulling aside a pair of cups and showing the child that there was a sticker inside one and no sticker inside the other. The experimenter then said, “In this game, it is going to be your job to guess which cup has the sticker inside. For each set of cups I’m going to ask you twice which cup you think has the sticker inside. After you make your first guess I will ask you once more and you can either keep your guess the same or guess the another cup. If your second guess is right then you get a point and for every point you get we will play a game at the end. So remember you want to try and get your second guess right so you can get a point!” The experimenter then proceeded to take a different pair of cups and said, “Let’s take a look at these two cups. One of them has the sticker inside and it’s going to be your job to guess which one. I’m going to look inside so I know which cup has the sticker.” The experimenter then looked inside and then asked the child which cup they believed had the sticker. Regardless of the accuracy of the children’s guess, the experimenter randomly responded either “Yes that’s right!” or “Hmmm that’s not right” and then asked the child again which cup they believed had the sticker inside. Children did not see the contents of the cup immediately after the trials, so they did not receive feedback as to whether their guesses were correct. The second training trial was the same as the first with the exception that the experimenter reversed the response provided after the child’s initial guess.

The experimenter then began the test trials. In the *Knowledgeable* condition the experimenter said, “I’m going to look so I know which cup has the sticker inside” making their knowledge state explicit. She then proceeded to ask the child which cup they believed had the sticker inside; when the child responded, the experimenter provided no explicit feedback as in the training trials, but instead said, “Okay, you said this cup had the sticker inside. Is that your final guess; which cup do you think has the sticker inside?” Children did not see the location of the sticker. In the *Ignorant* condition the experimenter said, “I’m not going to look inside so I don’t know which cup has the sticker either” making their ignorance explicit; the rest of the condition proceeded as with the *Knowledgeable condition*. At the end of the experiment, the experimenter brought back all the pairs of cups and let the children discover which cups contained the stickers.

Results

Coding Children’s responses were video taped and recoded by an assistant blind to condition; seven children were coded live because either no video consent was provided by the parents or because the view of the children’s pointing response was obstructed. For the remaining 25 children, there was 92% agreement; the errors were caused by obvious Left/Right coding errors and were resolved by a third coder.

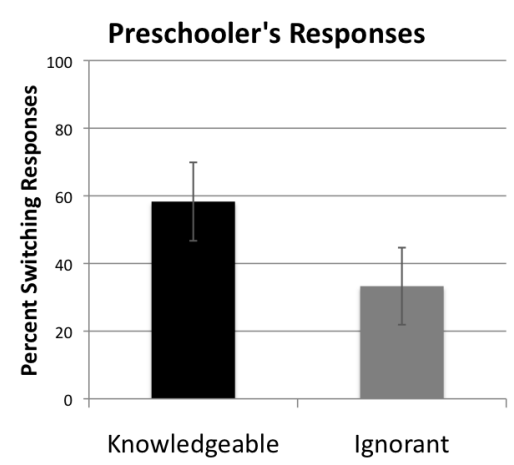


Figure 2: The percent of total test trials on which preschoolers switched their guesses in the *Knowledgeable* and *Ignorant* conditions following the neutral query.

Inclusion Criteria Because the order of “correct” and “incorrect” feedback was randomly assigned and counterbalanced, and because the experimenter could not control which cup the children would initially guess, some children were unintentionally “trained” to believe that the sticker always was located in the cup to one side. We classify two scenarios in which this side-bias occurs as follows: (1) During the first training trial the participant selects side A and is told that he is incorrect so he switches his guess to side B. On the second training trial, the participant chooses side B again and is told he is correct and so he continues to select side B for all subsequent guesses, both before and after the query. (2) During the first training trial the participant selects side A and is told he is correct and stays with side A. On the second training trial, the participant then chooses side B and is told he is incorrect, which leads him to switch his guess to side A and continue to select side A without switching on all subsequent guesses. In these two cases we have (by virtue of our random design) given the participant evidence that the sticker will always be on the same side. As discussed above nine children were dropped and replaced for this reason.

Test Results. Children’s performance on the initial feedback trials was nearly perfect, with children appropriately switching responses following the feedback that they were incorrect on training trials and appropriately staying with their response following the feedback that they were correct (94%). Our measure of interest was whether children switched their predictions in the test trials following the neutral query of “Is that your final answer; which cup has the sticker inside?” On aggregate, there were no differences between the three test trials within condition (*knowledgeable* Fisher Exact $p = 1$; *Ignorant* Fisher Exact $p = 1$), so the switching response of each child was scored as a total of switching responses by condition; children were more likely to switch their responses in the *Knowledgeable* condition (58% of the trials) than were

children in the *Ignorant* condition (33% of the trials; Pearson $\chi^2 = 6.04, p = .01$, see Figure 2). Children in the *Ignorant* condition were less likely to switch responses than would be predicted by chance (Binomial, $p = .01$). No age effects were observed within or across conditions.¹

Taken together, these results show that children take the knowledge state of the questioner into account; they are more likely to switch responses when the questioner is knowledgeable than when she is ignorant. These results are consistent with the *positive bias* model, suggesting that in addition to taking the knowledge state of the informant into account, preschoolers also have the assumption that the informant prefers to provide positive (over neutral) feedback following correct guess outcomes and a preference to provide a neutral query (over negative) following incorrect guess outcomes.

General Discussion

Recent research investigating children's learning of causal relationships and inferences about knowledgability suggests that while children's behavior is captured by rational models, in certain cases children's behavior appears less sensible, almost capricious. We have focused on children's seemingly irrational tendency to switch their responses when subjected to repeated questioning, proposing that this behavior may be the rational consequence of reasoning about the questioner's knowledge.

In our example, we presented a relatively simple model that accounts for switching behavior by considering the relationship between the kinds of responses informants may give and the relationship of those responses to the true state of the world. The main prediction of the model is that children's tendency to switch in response to questioning should depend on the epistemic state of the questioner. We presented a simple experiment which showed that preschool children switch their guesses more in response to neutral queries by a knowledgeable informant than by neutral queries by an ignorant informant.

Given that we told participants that they would be asked to guess twice, one might wonder whether the neutral queries should provide any information. By virtue of our design, we made it clear to children that the experimenter could also provide positive ("that's right") or negative ("that's wrong") feedback to children. Thus, by instead asking "is that your final answer" on test trials, children could reason about the implications of *not* hearing one of the alternatives. By providing neutral feedback when other feedback was available, the experimenter provided implicit information to the learner.

Our results leave open the question of why children should switch their hypotheses at all in the *Ignorant* condition. The switching pattern observed in the *Ignorant* condition suggests that social inferences about informant's intentions are not the

only explanation for children's switching behavior. One possibility is that regardless of the epistemic state of the questioner, the opportunity to re-evaluate an answer leads children who have a degree of uncertainty about their answers to reconsider their response. If guessing the outcome of a coin toss, there is no necessary reason to guess heads repeatedly. Similarly, in the ignorant condition, children's switching behavior may simply reflect their uncertainty about the correct answer. Recent research suggests that questioning may simply present the opportunity to reconsider hypotheses (Abbot & Griffiths, 2011), which can lead to a new response; deciding to stay or switch are both rationally viable (Denison et al., 2010; Bonawitz, Denison, et al., 2011).

The model highlights the kinds of assumptions a learner must bring to the table. Our empirical results suggest that preschool children must (1) take knowledge state of informant into account, and (2) assume a positive bias in informants. A large and growing body of evidence suggests that children keep track of others' knowledge and use it for various reasoning problems (e.g. Corriveau et al., 2009; Corriveau & Harris, 2009; Koenig & Harris, 2005). To our knowledge, the argument that children assume this positive bias in informants is novel and has interesting implications for learning and cognitive development. For instance, one might reasonably ask whether this bias is a product of experience and whether such experiences would affect the interpretation of neutral queries.

Our model makes additional predictions that may be explored in future work. For example, training children that the questioner has negative biases should lead to a decrease in switching responses. Switching behavior should also be dependent on the degree to which the informant has information about the actual location; neutral queries from an informant with partial knowledge may lead children to switch responses less often than an informant with complete knowledge. Another possibility is to explore the model's predictions over time to see how children learn about the intentions and knowledge of the informant over a series of trials. Finally, older children and adults may bring different assumptions about informants, so finding age-dependent changes may help characterize how social-causal inferences change with development.

Our results raise important practical considerations in science and education. Often times researchers, teachers, and parents ask children questions as a means of evaluating what the child believes. These questions are assumed to be inert, providing no additional information to the child, and the responses are treated as windows into the child's thinking. Our results call both of these assumptions into question. Even seemingly neutral queries provide information to children when posed by someone who can reasonably be assumed to be knowledgeable. And, because they are savvy social operators, children may change their minds in response to these neutral queries. In our experiments, we intentionally chose situations in which children would have weak beliefs about the location of the sticker, because previous research has sug-

¹Comparing the average number of switches per child across young and old (by median split), we find no differences (Knowledgeable: $t(14) = -.35, p = .73$; Ignorant: $t(14) = 0, p = 1$, Aggregate: $t(30) = -.25, p = .80$

gested that children's prior beliefs play an important role in how they interpret ambiguous evidence (Schulz et al., 2007) and how they account for evidence that conflicts with strongly held beliefs (Bonawitz et al., 2012). We do not know the degree to which children's prior beliefs will interact with inferences in these settings; However, regardless of prior beliefs, neutral queries may act as prod for the child to reconsider what she previously believed and thus may not provide a true window into the child's thinking. Thus, children's responses following repeated questioning in research and education settings should be interpreted with caution.

Conclusions

Children use various assumptions about a teacher's knowledge and intent to guide their reasoning. Our work takes this idea and expands on it, suggesting that questions originally assumed to be inert and provide no feedback may in fact serve as cues for children to draw inferences from. Learning about the world is an immensely daunting task for young minds, yet children are able to learn rapidly and accurately even in light of limited information. In school settings, at home, and even in eyewitness testimony, we employ repeated questioning as a means to assess a child's beliefs. When asked by a knowledgeable informant, such as a teacher, parent, or attorney, these questions may not simply elicit information about what the child believes, but instead may give the child reason to change their beliefs all together.

Acknowledgements: Thanks to Sophie Bridgers and Jaelyn Harris for data collection, to Swe Tun, Alvin Chan, and Sonia Spindt for coding, and to participating daycares, parents and children. This research was additionally supported by grant IIS-0845410 from the National Science Foundation and the James S. McDonnell Foundation.

References

- Abbot, J., & Griffiths, T. L. (2011). Exploring the influence of particle filter parameters on order effects in causal learning. In *Proceedings of the 33rd annual conference of the cognitive science society*.
- Bonawitz, E., Denison, S., Gopnik, A., & Griffiths, T. (2011). A simple sequential algorithm for approximating bayesian inference. In *Proceedings of the 33rd annual conference of the cognitive science society*.
- Bonawitz, E., Fisher, A., & Schulz, L. (in press). Teaching three-and-a-half year olds to reason from ambiguous evidence. *Journal of Cognition and Development*.
- Bonawitz, E., Schijndel, T. van, Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, 64(4), 215-234.
- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N., Spelke, E., & Schulz, L. (2011). The double-edge sword of pedagogy: Teaching limits children's spontaneous exploration and discovery. *Cognition*.
- Buchsbaum, D., Gopnik, A., Griffiths, T., & Shafto, P. (2011). Childrens imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, 120(3), 331 - 340.
- Butler, L., & Markman, E. (2010). Pedagogical cues influence children's inductive inference and exploratory play. In *Proceedings of the 32nd annual conference of the cognitive science society*.
- Corriveau, K. H., Fusaro, M., & Harris, P. L. (2009). Going with the flow: Preschoolers prefer non-dissenters as informants. *Psychological Science*, 20, 372-377.
- Corriveau, K. H., & Harris, P. L. (2009). Choosing your informant: Weighing familiarity and past accuracy. *Developmental Science*, 12, 426-437.
- Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. (2010). Preschoolers rationally sample hypotheses. In *Proceedings of the 32nd annual conference of the cognitive science society*.
- Gopnik, A., Glymour, C., Sobel, D., Schulz, L., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, 111, 1-31.
- Howie, P., Sheehan, M., Mojarrad, T., & Wrzesinska, M. (2004). Undesirable and desirable shifts in children's responses to repeated questions: age differences in the effect of providing a rationale for repetition. *Applied Cognitive Psychology*, 18(9), 1161 - 1180.
- Koenig, M., & Harris, P. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, 76, 1261-1277.
- Krahenbuhl, S., Blades, M., & Eiser, C. (2009). The effect of repeated questioning on children's accuracy and consistency in eyewitness testimony. *Legal and Criminological Psychology*, 14(2), 263-278.
- Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition*, 112, 367-380.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge, UK: Cambridge University Press.
- Poole, D., & White, L. (1991). Effects of question repetition on the eyewitness testimony of children and adults. *Developmental Psychology*, 27(6), 975 - 986.
- Schulz, L. E., Bonawitz, E. B., & Griffiths, T. L. (2007). Can being scared make your tummy ache? naive theories, ambiguous evidence, and preschoolers' causal inferences. *Developmental Psychology*, 43, 1124-1139.
- Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (in press). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*.
- Shafto, P., & Goodman, N. (2008). Teaching games: Statistical sampling assumptions for pedagogical situations. In *Proceedings of the 30th annual conference of the cognitive science society*.
- Spirtes, P., Glymour, C., & Schienens, R. (1993). *Causation prediction and search*. New York: Springer-Verlag.