The Double-edged Sword of Pedagogy: Modeling the Effect of Pedagogical Contexts on Preschoolers' Exploratory Play

Elizabeth Bonawitz[1*]**, Patrick Shafto**[2*]**, Hyowon Gweon**[1]**, Isabel Chang**[1]**, Sydney Katz**[1]**, & Laura Schulz**[1]

[1] **(liz_b@mit.edu)** Department of Brain and Cognitive Sciences, MIT, Cambridge, MA 02139 USA
[2] Department of Psychological and Brain Sciences, University of Louisville, Louisville, KY 40292 USA

Abstract

How does explicit instruction affect exploratory play and learning? We present a model that captures pedagogical assumptions (adapted from Shafto and Goodman, 2008) and test the model with a novel experiment looking at 4-year-olds' exploratory play in pedagogical and non-pedagogical contexts. Our findings are consistent with the model predictions: preschool children limit their exploration in pedagogical contexts, spending most of their free play performing only the demonstrated action. By contrast, children explore broadly both at baseline and after an accidental demonstration. Thus pedagogy constrains children's exploration for better and for worse; children learn the demonstrated causal relationship but are less likely than children in non-pedagogical contexts to discover and learn other causal relationships.

Keywords: Exploratory Play; Pedagogy; Bayesian Model

Learning from Play

In a preschool classroom, there are few sayings more ubiquitous than 'children learn from play'. Indeed, since Piaget (1929), parents, teachers, and scientists alike have argued that self-guided play serves as an important vehicle for learning both inside and outside the classroom. However, research in the Vygotskyean tradition (1978) has placed relatively less emphasis on children's self-directed exploration and more emphasis on how children learn from social interactions and cultural transmission. Surprisingly, few studies have looked at how exploratory learning and direct instruction interact.

Specifically, research on children's spontaneous exploratory play suggests that children's play is affected by both the evidence children observe and the children's prior beliefs. For example, children play more when evidence is ambiguous than when it is unambiguous, and play more when evidence violates their causal beliefs than when it is unsurprising (Schulz & Bonawitz, 2007; Bonawitz, Lim, & Schulz, 2007; Gweon & Schulz, 2008). However, this research has not looked at whether children's exploration differs in pedagogical and non-pedagogical contexts.

By contrast, research investigating learning through social interaction suggests that young children are sensitive to whether information was generated intentionally or accidentally (Xu & Tenenbaum, 2007), by a reliable or unreliable teacher (Koenig & Harris, 2005; Kushnir, Wellman, & Gelman, 2008), or in a neutral verses a pedagogical setting (Gergely, Kiraly, & Egyed, 2007; Tomasello & Barton, 1994). However, these projects have not looked at how children's inferences affect their exploratory behavior.

Given that children learn both from spontaneous exploration and explicit instruction, how does explicit instruction affect exploratory play and learning? In pedagogical situations, it is reasonable for learners to expect that the teacher is helping them learn; this expectation may facilitate learning in novel situations (Klahr & Nigam, 2004; Csbra & Gergely, in press). Learning the affordances of a novel artifact is challenging because for any object, there are an unknown, and potentially large, number of causal properties. If a knowledgeable teacher explicitly demonstrates one action and a novel effect results, a learner might reasonably infer that that there is a causal relationship between the action and the effect. Additionally, if the teacher demonstrates only the single action/outcome relation, the learner might infer that other potential actions afforded by the object are less likely to generate novel or interesting effects. Thus in this example, teaching informs the learner both about the existence of demonstrated causal relationships and the non-existence of other relationships.

In this paper, we present a computational model of reasoning in pedagogical situations, which predicts decreased exploration in pedagogical situations. We test this prediction using a novel toy exploration paradigm. The toy was created to have many different, not immediately obvious causal properties. We contrast exploratory play in three conditions: pedagogical demonstration, accidental demonstration, and no demonstration. If children's play is sensitive to pedagogical sampling assumptions, then in the pedagogical condition, children should be more likely than children in the accidental condition to assume that the demonstrated action is the only causal property in the toy; thus, children in the pedagogical condition should be less likely to discover other causal properties of the toy.

Pedagogical Model

Recent research has contrasted models of pedagogical and non-pedagogical settings on learners' inferences. Shafto and Goodman (2008) formalize pedagogical learning as Bayesian inference based on the assumption that the teacher is being helpful. The learner expects that the teacher chooses data, $p(d|h)$, that tend to increase the learner's belief in the true hypothesis

$$p(d|h) \propto p(h|d)$$

---

[*] The first two authors contributed equally to this work.

where the p(h|d) represents what the learner's beliefs would be after having observed the data. The learner is thus inferring why the teacher provided this data.

Pedagogical sampling can be contrasted with more standard random sampling assumptions, which assume that which data are chosen do not provide any information about the true hypothesis (e.g. Fried & Holyoak, 1984). Random sampling contributes only a multiplicative constant to inferences,

$$p(d|h) = 1/n$$

where n is the number of possible sets of examples. Because n does not depend on the hypotheses, random sampling does not differentially prefer any particular hypothesis.

Learning in pedagogical and non-pedagogical settings can be formalized as Bayesian inference using the appropriate sampling model. Learners update their beliefs in different hypotheses based on the product of the probabilities of the data given the hypothesis and the prior probability of the hypothesis,

$$p(h|d) \propto p(d|h)p(h)$$

where the appropriate p(d|h) depends on whether the setting is pedagogical or accidental, and p(h) specifies the learner's prior beliefs about possible hypotheses. These two models formalize the computational problems that children face in different situations, and can therefore be used to make predictions about how behavior will differ as a consequence of the different learning situations. Our argument is that children understand the inferential implications of these situations, and we make no claims about the underlying process that generates these inferences.

In this paper, we assume that the child is inferring how many different possible actions on a novel toy have effects. In principle, there may be an unbounded number of cause-effect pairs on the toy. However, for most toys a relatively small number of possible actions have novel effects. We formalize this intuition using a poisson prior on the number of causes with effects in a causal graphical model. Thus, learners are inferring which possible causes have effects, and how many cause-effect pairs there are, with the prior belief that there are a relatively small number of active causes.[1]

We model the causal relationships using a noisy-or parameterization. Noisy-or models are parameterized with a background rate and a transmission rate. The background rate specifies the probability that a cause or effect spontaneously activates itself and was set to zero to capture the intuition that the toy cannot spontaneously activate itself. The transmission rate is the probability of an effect given a cause. A deterministic relationship would have a value of 1. Causes may be perceived as non-deterministic

---

[1] Though the space of possibilities is in principle very large, because prior probability drops off rapidly with increasing numbers of causes, we truncate the hypothesis space to only graphs with up to four causes and four effects.

for a number of reasons, including if children are unable to reliably elicit the effect. To set a realistic transmission rate, we coded the number of times that children successfully and unsuccessfully could generate the effect when they acted on the cause (83% successful actions) and accordingly set the value of the transmission rate to 0.83. (However, note that results are robust across a range of values.)

Imagine a child observing either a pedagogical or accidental demonstration. Should the child be inclined to explore more? The two models generate predictions for these situations. The random sampling model captures the case of an accidental intervention. In this case, it seems plausible that there may be more cause-effect relationships to be found, given that a random intervention generated an effect. More formally, the fact that the data are sampled independently from the hypothesis means that the only information that is gained is contained in the results of the demonstration – there is at least one cause-effect pair.

In the pedagogical situation, because the demonstrator knows about the toy, the demonstration could be understood to be teaching the child about the toy. More formally, the demonstration implies information about the hypothesis that is being taught. If the correct hypothesis was that there are two causes, then the teacher would have demonstrated both to maximize the learner's chances of converging on the correct hypothesis. The demonstration of only one cause implies that there are no other cause-effect pairs. More generally, according to the pedagogical model, the absence of evidence for a hypothesis is taken as evidence against that hypothesis. Thus, after pedagogical demonstrations, children should expect that additional causes are less likely, and therefore be less inclined to explore.

Experiment

If children are sensitive to the differences in pedagogical and non-pedagogical situations, and the sampling assumptions of the pedagogical and random sampling models are correct, then children should infer that there are few or no other potential causal relationships to learn when they are only 'taught' one cause-effect pair. As such, we predict that children's exploratory play and learning will be affected by manipulations in conditions. We expect that children who observe pedagogical demonstrations will spend a larger percentage of time exploring the demonstrated action than children who observe the same information generated accidentally. Similarly, children in the accidental condition should perform more different types of interventions on the toy than children in the pedagogical condition. Perhaps most interestingly, because of the differences in exploration of the toy, we should find differences in learning. Children in the pedagogical condition, who have limited their exploration, should be less likely than children in the accidental condition to learn the other causal properties of the toy.

However, differences between conditions may also be caused because 'accidentally' discovering a causal property motivates exploration. That is, pedagogical assumptions

may not constrain children's exploration; rather, the 'surprise' in the Accidental condition may increase children's exploration. To address the alternative account, we ran a third condition called the *No Demo* condition. In the *No Demo* condition, children were never shown any of the toy's causal affordances when the toy was first introduced, and we expect children to explore readily in this condition as well.

Methods

Participants Forty-Eight preschoolers (mean age: 58 months; range: 46 – 74 months) were recruited in a metropolitan Science Museum. Sixteen children were tested in each of three conditions: *Pedagogical, Accidental,* and *No Demo*. There were approximately equal number of boys and girls in each condition.

Materials A novel-looking toy was created using colored PVC pipes attached to a board (see Figure 1). The toy was approximately 18" x 6" x 15". The toy had four different non-obvious causal affordances: the toy made a squeak sound when a yellow-colored tube was pulled out from inside a larger purple tube; one end of a blue tube lit up when a small button hidden inside the other end was pressed; a small yellow pad attached to the plastic board played music notes when different parts of the pad were pressed; there were two adjoining black tubes with mirrors inside so that a mirror image of the observer's face was visible. All other aspects of the toy were inert.

Procedure Children were tested in a quiet corner in the museum. The experiment included three phases, the introduction phase, the play phase, and the question phase.

Introduction Phase. In both conditions, the experimenter brought the toy out from under the table and introduced the toy to the child. In the *Pedagogical* condition, the experimenter said, "Look at my toy! This is my toy. I'm going to show you how my toy works. Watch this!" The experimenter then pulled the yellow tube out from the purple tube to produce the squeak sound. She said, "Wow, see that? This is how my toy works!" and demonstrated the same action again. In the *Accidental* condition, the experimenter said, "Look at this toy! See this toy?" However, as she brought out the toy from underneath the table, she pulled the yellow tube out from the purple tube as if she did so by accident. Then she said, "Wow that was weird. Did you see that? Let me try to do that!" and performed the same action to produce the squeak sound. In the *No Demo* condition, the experimenter did not initially demonstrate the squeaking property of the toy. After she brought out the toy from underneath the table, she picked up the toy and said "Wow, see this toy? Look at this!" She looked at the toy for about 2 seconds (to match the other conditions for amount of familiarization time), and then put it back down on the table.

Play Phase. In both conditions, after the child observed that pulling the tube made the squeaking sound, the experimenter said "Wow, isn't that cool? I'm going to let


Figure 1: Stimuli with labeled causes.

you play and see if you can figure out how this toy works. Let me know when you're done!" and left the child to play. She returned to the table when the child said that he or she was done. If the child stopped interacting with the toy for more than 5 consecutive seconds without indicating completion, the experimenter prompted the child by saying "Are you done?" and returned to the table if the child answered "Yes". Otherwise, she let the child continue to play and then returned to the table if the child stopped interacting with the toy a second time for more than 5 seconds.

Question Phase. After returning to the table, the experimenter hid the toy behind an occluder and assessed whether the child had discovered the four causal affordances. First, she produced the squeaking sound behind the occluder. She handed the toy back to the child and asked, "Did you hear that? Can you show me how to make that sound?" The child was allowed only one attempt to answer each question. The experimenter placed the toy behind the occluder again and showed the top end of the blue tube lighting up while the rest of the toy was hidden, and asked the child to show her how to make the toy light up. For the third question, she played the music notes behind the occluder and asked to play the music sounds. For the fourth question, she ducked behind the occluder and said, "Wow, I can see myself! Can you show me how to see yourself in the toy?" The experimenter concluded the experiment by letting the children volunteer any additional information about the toy that they wanted to communicate.

Results of Experiment

There were no differences in age between children in the conditions (*Pedagogical, Accidental: t*(30) = .4, *p = NS; Pedagogical, No Demo: t*(30) = -.48, *p = NS; Accidental, No Demo: t*(30) = -.12, *p = NS*). Children played for the same amount of time in *Pedagogical* and *Accidental* conditions (*Pedagogical* Mean = 120 seconds; *Accidental* Mean = 94 seconds; *t*(30) = 1.04, *p = NS*). However, children in the *No Demo condition* played longer (Mean = 192 seconds) on average than children in the *Accidental* condition (*t*(30) = 2.76, *p* < .01) and marginally longer than children in the *Pedagogical condition* (*t*(30) = 1.71, *p* < .09).

Figure 2: Model predictions and children's responses. The model (top) predicts probability that of each number of causes (0-4) based on the demonstration. The children's results (bottom) show the percentage of children who explored exactly 0, 1, 2, 3, and 4 of the cause-effect pairs.

We coded every action the child took on the toy. (That is, we coded not only children's exploration of the four built-in affordances but also other actions – shaking the entire toy; turning the toy upside down, etc.). We also coded the number of the built-in cause-effect relationships that the child could correctly demonstrate at the end of the experiment. All actions were coded by the fourth author and blind coded by the fifth author; reliabilility was high (r = .96) and all discrepancies on the final learning measure were resolved by the first author, blind to condition. To compare the children's actions to predictions of the model, we counted the number of children who explored just one, any two, any three, or all four built-in causal relations during the free-play period. Correlation between the model and child results was high (r = .90, see Figure 2).

Play Results. Children in the *Pedagogical* condition spent significantly more of their play time exploring the squeaker toy than children in the *Accidental* condition (*Pedagogical* Mean = 68%; *Accidental* Mean = 43%; $t(30) = 2.73$, $p < .01$). Of the children who discovered the 'squeaker' cause during the course of free play (N = 7), children in the *No Demo* condition spent significantly less of their play time with the squeaker toy (Mean = 22%, starting from point of discovery) than children in the *Pedagogical* condition ($t(30) = 4.27$, $p < .01$), but only marginally less time than children in the *Accidental* condition ($t(30) = -1.86$, $p = .08$).

Children in the *Pedagogical* condition also performed significantly fewer other types of actions on the toy (*Pedagogical* Mean = 3.7 actions; *Accidental* Mean = 5.1 actions; *No Demo* Mean = 5.8 actions; *Pedagogical,*

Accidental: $t(30) = -1.77$, $p < .05$; *Pedagogical, No Demo:* $t(30) = 2.47$, $p < .05$). However there were no differences in the number of types of actions performed on the toy between children in the *No Demo* and *Accidental* conditions ($t(30) = .79$, $p = NS$). (See Figure 3).

To make sure that differences were not driven by the slightly longer average playtime in the *No Demo* condition, we also compared the number of different actions children discovered in the first minute of play across conditions. The pattern of results held. Children in the *No Demo* condition explored more different actions in the first minute of play[2] (M = 4.33) than children in the *Pedagogical* condition (M = 2.69), ($t(26) = 2.23$, $p < .05$), but there were no differences between *No Demo* and *Accidental* conditions (*Accidental* Mean = 5.09; $t(24) = -.88$, $p = NS$).

Learning Results: As a consequence of the limited variable exploration in the *Pedagogical* condition, we also found differences in learning by children in the *Pedagogical* condition. While children in both *Pedagogical* and *Accidental* conditions were able to replicate the demonstrated action on the squeaker (*Pedagogical* Mean = 100% correct; *Accidental* Mean = 88% correct; $\chi^2 = 2.13$, $p = NS$), children in the *Pedagogical* condition were significantly less likely to have learned the other causal properties of the toy, (*Pedagogy* Mean = 20% correct; *Accidental* Mean = 46% correct; $\chi^2 = 6.75$, $p < .01$). Unsurprisingly, children in the *No Demo* condition were significantly less likely to be able to demonstrate the 'squeaker' action at the end of the experiment than children who had the benefit of the initial demonstration (Mean = 25% correct; *Pedagogical:* $\chi^2 = 19.2$, $p < .001$; *Accidental:* $\chi^2 = 12.7$, $p < .001$). However, children in the *No Demo* condition were significantly more likely to learn the other, non-demonstrated causal affordances of the toy (mean = 40% correct) than children in the *Pedagogical* condition ($\chi^2 = 4.00$, $p < .05$), but were equally as likely to learn as children in the *Accidental* condition ($\chi^2 = .383$, $p = NS$).

Discussion of Experiment

Our results are consistent with both the quantitative results of the model and our qualitative predictions. Even children as young as four seem to be sensitive to pedagogical information: they limit their exploration in pedagogical contexts by spending more time on the demonstrated action and performing fewer other actions on the toy; they are thus less likely to discover and learn new causal affordances of the toy. Surprisingly, children in the *Pedagogical* condition constrained their exploration relative to children who did not get any direct teaching (as in the *Accidental* and *No Demo* condition) and were thus less likely to learn the other causal properties of the toy.

While there were no differences in learning the demonstrated action (the 'squeaker') in *Pedagogical* and

---

[2] For comparison on this measure, we excluded children who played for less than 55 seconds total: 1 child removed in the *No Demo* condition, 3 in the *Pedagogical*, and 5 in the *Accidental*. Results are not affected if these children are included.

Figure 3. Mean percent of children's total play time spent playing with the squeaker toy in each condition; in the second graph, mean number of distinct types of actions (including squeaking the squeaker toy) that children in each condition perform; in the third graph mean number of other causal relations (besides the squeaker) that children learned.

Accidental conditions, the non-differences is perhaps not surprising. Generating the squeaking action just a minute after the demonstration is relatively simple task, so children in both groups should succeed in replicating this action. It remains an open question whether pedagogical information is more memorable than accidental information after a longer delay; or whether pedagogical information is more informative than accidental information when the target action is more complicated.

Discussion

Inspired by the ideas of Piaget and Vygotsky, we set out to investigate the implications of explicit instruction on exploratory play. We presented a formal model that captures our intuitions about how sampling assumptions (in pedagogical and non-pedagogical contexts) could influence our hypothesis space about possible causal models of the world. Our empirical results suggest that children do indeed constrain their exploratory interventions on a toy in pedagogical contexts. As a result, children in the pedagogical contexts replicate and learn the demonstrated causal relation, but are less likely to discover and learn other affordances of a novel toy.

These results are also consistent with, and can help interpret discrepancies in, the literature on direct instruction and discovery learning. For example Klahr and Nigam (2004) find that direct instruction helps children learn the control-of-variables strategy more effectively than discovery play. However, Dean and Kuhn (2006) found that over longer term learning, self discovery is an important factor in children's learning; in fact, they found that direct instruction is neither a necessary nor sufficient condition for full conceptualization. If, as our results suggest, direct instruction limits children's exploration to and therefore learning about the demonstrated concepts, then the conflict between these projects is not surprising. Initially, direct instruction offers a fast strategy for concept learning.

However, over time, children who receive only direct instruction will be less likely to explore and discover relevant strategies, and thus less able to acquire and consolidate the relevant concepts.

Although we have discussed this study specifically with reference to pedagogical contexts, teaching is just one of the contexts in which children are likely to obtain data from a knowledgeable intentional agent. Children may also constrain their exploration in a condition where they are not being explicitly instructed, but instead just witness a knowledgeable agent intentionally acting on a toy (see Goodman, Baker, & Tenenbaum, in press, for discussion). Although we believe pedagogical contexts are a particularly important means by which children learn and the implications of these situations for learning are potentially quite different (Csibra & Gergeley, in press), our experiment was not designed to highlight these differences. Future research will focus on contrasting model predictions and children's inferences in these situations.

It may be interesting to compare this study with research on 'functional fixedness' (Duncker, 1945; German & Defeyter, 2000). Our study resembles the functional fixedness studies in that demonstration of a given property of the toy impeded children's ability to discover additional properties of the toy. However, our study contrasts with the functional fixedness literature in a number of respects. First, we used a complex artifact with many plausibly functional parts. It is not clear whether, even in adults, functional fixedness would obtain for complex objects: many complex objects (e.g., computers) are clearly designed to have more than a single intended function. Second, children were never asked to generate more than one function for any given part of the object. Thus the type of difficulty involved in functional fixedness (overcoming the known function of an object) seems less readily applicable here. Third, although preschoolers are not subject to functional fixedness (the literature suggests difficulties with functional fixedness do not emerge until age six), three, four

and five-year-olds in our study were sensitive to the pedagogical demonstration. The functional fixedness literature has generally looked at learners' latency to solve problems of insight learning; that is, learners' have been asked to use an artifact with a known function to solve a problem requiring a different function. Here children did not have to engage in any problem solving; we simply looked at the number of different actions children took in free exploration. Our account also differs theoretically from the functional fixedness account: although we accept that preschoolers might not adopt a design stance with respect to artifacts (German & Defeyter, 2000), we believe that younger children can nonetheless make a rational inference about object properties given the evidence of a teacher's demonstration. If children can infer that a teacher demonstrated one and only one property because only the single property exists, this would lead to limited exploration. It would be interesting for future work to explore how children's assumptions about pedagogy and sampling interact with their development of a design stance towards artifacts and the onset of functional fixedness.

Finally, we note that for the purpose of the study, we deliberately misled the children: the presumably rational teacher (benevolent and knowledgeable) presented only a single function of a multi-functional toy. One might wonder, are these results simply attributable to poor teaching? We think this is not the case. In this experimental context we have predetermined the aspects that are to be learned, and this set is fixed and finite. However, in most natural contexts, the set of concepts to be learned is neither fixed nor finite. Indeed, this is why Dean and Kuhn (2006) focus on longer-term learning: education requires modifying learning goals with time and understanding how concepts build on each other. Thus, in natural learning contexts, pedagogical demonstrations cannot demonstrate all there is to know, and teaching will necessarily be limited. Understanding how to combine the efficiency of pedagogical knowledge transmission while encouraging curiosity and exploratory play is an important direction for future work.

More generally, this research presents a step toward reconnecting two historically divergent research traditions. Unquestionably, Piaget was correct in emphasizing the role of exploratory play in children's learning. Similarly, Vygotsky was also correct to emphasize the importance of social learning and cultural transmission of knowledge. Our research is an attempt to reconnect these traditions, both formally and empirically, by asking how social transmission affects exploratory play. We believe a more complete understanding of development depends on understanding how these processes interact.

References

Bonawitz, E.B., Lim, S., & Schulz, L.E. (2007) Weighing the Evidence: Children's theories of Balance affect play. *Proceedings of the Twenty-Ninth Annual Conference of the Cognitive Science Society.* Nashville, Tennessee.

Csibra, G. & Gergely, G. (in press). Natural Pedagogy. *Trends in Cognitive Sciences*

Dean, D. & Kuhn, D. (2006) Direct Instruction vs. Discovery: The Long View. *Science Ed., 91.* 384-397.

Duncker, K. (1945). On problem solving. *Psychological Monographs, 58:5* (Whole No. 270)

Gergely, G., Király, I., & Egyed, K. (2007). On pedagogy. *Developmental Science*, 10:1, 139-146.

German, T.P., & Defeyter, M.A. (2000). Immunity to functional fixedness in young children. *Psychonomic Bulletin & Review*, *7(4),* 707-712

Goodman, N., Baker, C., & Tenenbaum, J. (in review) Cause and Intent: Social Reasoning in Causal Learning. *Proceedings of the 31st annual meeting of the Cognitive Science Society.*

Gweon, H., & Schulz, L. (2008) Stretching to learn: Ambiguous evidence and variability in preschoolers' exploratory play. *Proceedings of the 30th annual meeting of the Cognitive Science Society.*

Klahr, D., & Nigam, M. (2004). The equivalence of learning paths in early science instruction: Effects of direct instruction and discovery learning. *Psychological Sciences, 15,* 661-667.

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

Kushnir, T., Wellman, H. M. & Gelman, S. A.(2008). The role of preschoolers social understanding in evaluating the informativeness of causal interventions. *Cognition. 107 (3),* p. 1084-1092

Piaget, J. (1929) *The Child's Conception of the World.* New York: Harcort, Brace.

Tomasello, M. and Barton, M. (1994). Learning words in nonostensive contexts. Journal of Developmental Psychology, 30(5):639--650.

Schulz, L., & Bonawitz, E.B. (2007) Serious fun: Preschoolers play more when evidence is confounded. *Developmental Psychology, Jul Vol 43(4)* 1045-1050.

Shafto, P. & Goodman, N. (2008) Teaching games: Statistical sampling assumptions for pedagogical situations. *Proceedings of the 30th annual meeting of the Cognitive Science Society.*

Vygotsky, L. (1978) *Mind in Society.* Cambridge, MA. Harvard University Press.

Xu, F. and Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review 114*(2).