

# Learning About Causal Systems Through Play

Zi L. Sim\* (zi@berkeley.edu)

Kuldeep K. Mahal\* (kmahal@berkeley.edu)

Fei Xu (fei\_xu@berkeley.edu)

Department of Psychology, University of California, Berkeley  
Berkeley, CA 94720 USA

## Abstract

It is commonly believed that children are able to learn through play. Recent studies have found that children are able to learn causal rules through free play (Sim & Xu, in press). One such study found that children learned how to correctly activate machines, using either a block that was the same shape or the same color as the machine, when given five minutes to play with them. However, would children be able to learn a more complex causal rule through free play as well and would their performance be comparable to children who were didactically taught the same causal rule? In the current study, we show that children are able to learn more complex causal rules through free play. We also show that children perform significantly better when learning these rules through free play or by first engaging in free play and then observing, as opposed to solely through observation.

**Keywords:** free play; causal learning; generalization

## Introduction

It is widely accepted that play is important for young children. Most elementary schools have a designated play time, where children are free to play and socialize with their peers. Studies have shown that on average children spend 51 hours a week, or 30% of their week, engaging in free play (Hofferth & Sandberg, 2001). Play is also commonly encouraged by parents and educators, and access to play has been recognized by the United Nations Convention on the Rights of the Child (UNCRC) as a fundamental human right (Davey & Lundy, 2011). It is evident that children spend a lot of time playing and are encouraged to do so, but it remains unclear why this is the case. Why do children spend such a large portion of their time playing, and why do adults actively encourage play behavior?

One explanation is that play has the potential to result in better learning than direction instruction because play provides learners with the opportunity to choose what they want to do (Whitebread, Coltman, Jameson, & Lander, 2009; Weisberg, Hirsh-Pasek, & Golinkoff, 2013; Weisberg, Hirsh-Pasek, Golinkoff, & McCandliss, 2014). By doing so, learners may better encode new information in their memory (Metcalf & Kornell, 2005); process the problem structure more deeply (Sobel & Kushnir, 2006); they may pay more attention and are more motivated (Corno & Mandinach, 1983); or they may be able to focus on acquiring data that can address gaps in their knowledge (Markant & Gureckis, 2013). There is some empirical evidence that adults benefit from the opportunity to select the information that they want to learn. In fact, research has shown that adults learn better when they

engage in active hypothesis testing, where they are able to select the data they observe to test their hypotheses, as compared to those who engaged in reception learning, where they observed data generated by another adult (Castro et al., 2009; Markant & Gureckis, 2014; Sobel & Kushnir, 2006).

For example, in one study, adults were shown a spectrum of sixteen “alien eggs” on a computer that went from spiky to smooth. Students were told that spiky eggs most likely hatched into alien snakes while smooth eggs most likely hatched into alien birds. They were asked to determine the boundary between the two types of eggs, so they had to determine the point at which the eggs shifted from hatching one species to the other species. To do this, subjects either selected a sequence of eggs to see which animal hatched from them (active learning condition) or observed randomly selected eggs being hatched (random condition). The study found that participants in the active learning condition generally performed better than those in the random condition (Castro et al., 2009), which suggests that adults experience benefits in learning when they play an active role in gathering information.

Within the developmental literature, several studies have established that children learn successfully when they have the opportunity to choose what they want to do as well. For example, researchers have found that children as young as five years were able to use self-generated evidence to learn about an ambiguous causal system (McCormack, Bramley, Frosch, Patrick, & Lagnado, 2016; Schulz, Gopnik, & Glymour, 2007). In another study, Sim and Xu (in press) found that three-year-olds were capable of forming higher-order generalizations about a causal system after a short play period. In this study, children were presented with a causal learning task in which blocks would activate machines either based on a shape-rule (a block that matched the machine in shape activated the machine), or a color-rule (a block that matched the machine in color activated the machine). Children were randomly assigned to a didactic condition, such that an experimenter showed the children how to activate the machines, or a free play condition, such that the children were given the opportunity to play freely with these machines and blocks. The children were then tested using a first-order and a second-order generalization task. In the first-order generalization task, children were asked to activate a familiar machine, and in the second-order generalization task, children were asked to activate a novel machine. The study found that both groups performed at levels well above chance, and there was no significant difference between the accuracy between the two conditions (Sim & Xu, in press).

Similarly, Smith and Dutton (1979) compared the performance of children in a play condition to both a training condition, where children were taught to use the materials by an experimenter, as well as a control condition, where children were neither taught by an experimenter nor played with the materials. Although children in the play and training condition performed significantly better in a problem-solving task compared to the control condition, there was no significant difference between the play and training conditions.

Although all of these studies suggest that children are able to learn through free play, there is as yet little evidence that young children's learning and generalizations under free play conditions would actually differ from those in didactic conditions. Learning in the free play condition has repeatedly been found to be *comparable* to a didactic condition, but has not been found to be *different* from a didactic condition.

Are there conditions under which young children may benefit more from the free play that they engage in independently, as compared to a training that is directed by an adult? In the current study, we examined this question by presenting children with a causal learning task in which the generalization to be acquired was more complicated than that examined in Sim and Xu (in press). To do so, machines in our study were activated when either two blocks that were the same shape as the machine, or two blocks that were the same color as the machine, were placed on the machine. Prior research have indicated that young children can learn causal rules of a similar form: Walker and Gopnik (2014) showed that after a short demonstration by an experimenter, 18- to 24-month-olds were able to learn a "same" or "different" causal rule, i.e. that a machine was activated only by placing two identical blocks on the machine, or that a machine was activated only by placing two dissimilar blocks on the machine. Given that the rule used in the current study is of a more complex form – the two blocks had to match each other *and* the machine on a specific dimension (i.e., shape or color) – we chose to test 3- and 4-year-olds in the current study.

Children were randomly assigned to a free play condition, where they were presented with six machines (three categories of machines, with two identical machines within each category) and twelve blocks to play with for approximately 10 minutes, or the didactic condition, where they observed an experimenter activate each machine once. To further examine any potential benefits of play, children who did not activate the machines at least once in the free play condition were placed in a third condition, the free-play-first condition, where they observed the experimenter showing them how to activate the machines after they played by themselves for approximately 10 minutes. Similar to the study by Sim and Xu (in press), the children's ability to learn the correct rule was measured using a first- and second-order generalization test, where they were asked to activate both a familiar and a novel machine respectively.

## Method

### Participants

Sixty-one three- to five-year-old English-speaking children (29 boys and 32 girls) with the mean age of 48.4 months (range= 36.3 months to 59.1 months) were tested. All were recruited from Berkeley, California and its surrounding communities. Children were tested either in a small testing room at our lab or in a small quiet room in a preschool. Each child was randomly assigned to the didactic (N=24) or free play condition (N=22). Children assigned to the free play condition but did not activate the machines at least once during the free play phase were placed in the free play first condition (N=15). The mean ages in the didactic, free play, and free play first condition were 48.1 months, 49.7 months, and 47.1 months respectively. An additional 9 children were tested but were excluded due to parent interference (N=5) and experimenter error (N=4).

### Materials

Five different types of machines were constructed for this experiment. Each type of machine made a distinct sound when activated. This activation was completed using a foot pedal connected to a remote that activated a doorbell that was placed inside the machines. There were two blue rectangle machines, two red triangle machines, two green circle machines, and one orange L-shaped machine. In addition, there was a colorful felt-covered plus-shaped machine (demonstration machine) that looked considerably different from the other four types of machines. Each machine was approximately 20cm x 12cm x 10cm.

A variety of small blocks (approximately 7cm x 5cm x 1cm) were used to activate the machines. The activator blocks were of different shapes and colors. Some matched the machines in shape but not in color, or in color but not in shape, while other distractor blocks matched in neither shape nor color. In total, twenty-two blocks were used.

### Procedure

Each child was tested individually. For children tested in our lab, parents sat next to the child during the procedure and were asked not to interact with their child. For children tested at preschools, an observer watched the procedure through a one-way mirror.

Both the didactic and free play conditions consisted of three phases: a demonstration phase, a training/free play phase, and a testing phase. For half the children, each machine was activated by placing two blocks on the machine that matched the machine in shape (shape rule). For the other half, each machine was activated by placing two blocks on the machine that matched the machine in color (color rule).

Children in the free play condition who did not activate the machines at least once during the free play phase formed a separate group: the free-play-first condition. This condition consisted four phases: a demonstration phase, a free play phase, a training phase, and a testing phase. In other words, children who did not activate any machines during free play

were then trained by the experimenter using the procedure of the didactic condition. For four of these children, machines were activated by the shape rule, and for eleven children machines were activated by the color rule.




Figure 1: Schematic of materials and procedure.

**Didactic Condition** For the didactic condition, the children sat at a small table across from an experimenter. The demonstration phase began with the experimenter showing the child 12 blocks and pointing out that the blocks had different shapes and different colors. The blocks were then removed and the demonstration machine was placed on the table along with one block. The experimenter then showed the child how to make the machine go by placing the block on the machine and pressing down. After the machine made a sound, the experimenter noted that the blocks made the machine go and then allowed the child to try. Two new blocks were then placed in front of the child. The child activated the machine by pressing down with these *two* blocks and the experimenter stated that the blocks made the machine go. This was repeated next with three new blocks. The blocks and machine were then removed, ending the demonstration phase. The duration of this phase was around four minutes.

The experimenter then told the child that she had some new machines to show them. She also emphasized that these new machines were much pickier than the demonstration machine, so only some blocks would make them go.

The training phase that followed began with the experimenter presenting the first machine (e.g., blue

rectangle machine). The experimenter placed four activator blocks on the table next to the machine (e.g., the red, green, yellow, and purple rectangle blocks if the machine was activated by the shape rule; the blue triangle, blue circle, blue heart, and blue star blocks if the machine was activated by the color rule). The experimenter then stated, “Let me show you how to make this machine go,” and placed two of the four activator blocks (e.g., purple rectangle block and yellow rectangle block, or blue circle block and blue star block) on the machine and pressed both blocks down, activating the machine. She then exclaimed, “The machine made a sound! It played music!” The experimenter then told the child that she had another machine that was identical to the one in front of them. This machine was placed on the table, and the other two activator blocks were now placed on the new machine, activating it. The machines and blocks were then removed and the process was repeated with the remaining two sets of machines. Once the child had seen all six machines activated one time each, the training phase was complete. The order of the presentation for the types of machines was counterbalanced. The duration of the training phase was approximately five minutes.

The testing phase consisted of both a first-order generalization test and a second-order generalization test. The order of the tests was counterbalanced. For the first-order test, the children were first presented with six separate blocks: two blocks that matched the machine in shape, two blocks that matched the machine in color, and two distractor blocks that did not match the machine in shape or color (see Figure 1). The experimenter presented the blue rectangle machine from the training phase, and said, “Remember this machine? Remember that I made this machine go just now? Can you show me how to make this machine go?” If the child placed the correct blocks on to the machine, the machine activated and the experimenter neutrally stated that the machine made a sound. If the child did not place the correct blocks, the machine did not activate and the experimenter neutrally stated that the machine did not make a sound. For the second-order generalization, the child was presented once more with six new blocks (see Figure 1). The child was then shown a novel machine (the orange L-shaped machine) and told that the machine “was a picky machine too”. The experimenter then asked the child, “Can you show me how to make this machine go?” Once again, the experimenter neutrally stated that the machine made a sound if the child was correct, or did not make a sound if the child was incorrect.

**Free Play Condition** For the free play condition, the children sat on a blanket on the floor. The demonstration phase in the free play condition was identical to that in the didactic condition.

The free play phase began with the experimenter saying that she needed to check that all the machines worked (pilot testing suggested that this step was necessary in order to encourage children to keep playing even if they did not activate any machines after a few attempts). The machines were taken behind a table and activated so that the child could hear them

activate by the sound they made, but could not see how they were activated. The experimenter then placed all six machines (two blue rectangle machines, two red triangle machines, and two green circle machines) in front of the child and noted that these machines were pickier and that only some blocks made them go. The twelve activator blocks shown in the demonstration phase were placed in front of the child as well. The experimenter then told the child, “I just remembered I have to do some work now, but while I work, you can play with these machines and these blocks.” The children were given approximately ten minutes to play with the machines and blocks. If the child did not make any attempt to activate the machines for one minute, the experimenter prompted the child by saying, “Why don’t you try to make the machines go?” After ten minutes, the blocks and machines were removed, ending the free play phase.

The testing phase was identical to that of the didactic condition.

**Free-Play-First Condition** For the free-play-first condition, the children sat on a blanket on the floor. The demonstration phase and free play phase were identical to that in the free play condition. Children were then moved to a table and the training phase was identical to the training phase in the didactic condition. The testing phase was identical to that of the didactic and free play conditions.

## Coding

For children exposed to the shape rule, selecting the two blocks that matched the machine in shape was scored as one point. On the other hand, for children exposed to the color rule, selecting the two blocks that matched the machine in color was scored as one point. Since each child completed two tests, the maximum score that a child could receive was two points.

## Results

We analyzed effects on children’s responses with generalized linear mixed effects models in R, using an alpha level of 0.05 for all analyses. Children’s responses were coded as a binary variable where correct responses were coded as 1 and incorrect responses were coded as 0. In the model, subjects were specified as a random factor since this was a repeated measures task; each subject gave two responses. Although children in the didactic condition and the free play first condition all saw a total of six activations, the same cannot be said for children in the free play condition. In the free play condition, the total number of activations ( $M= 17.6$ ,  $SD= 12.6$ ) each child saw as well as the amount of time that each child played for ( $M= 7.91$ ,  $SD= 2.34$ ) varied. Although the didactic condition did not receive any negative evidence (i.e., observing unsuccessful activations), this was not the case for both the free play condition and the free-play-first condition. The amount of negative evidence generated by each child in the free play condition varied ( $M= 29.5$ ,  $SD= 21.8$ ). Children in the free-play-first condition also generated a varied amount of negative evidence ( $M= 28.1$ ,  $SD= 16.1$ ). Preliminary analysis showed no significant effect of age, sex, or

presentation order of the machines and testing phases. Additionally, there was no significant difference in the children’s performance between first-order and second-order generalization tests in each of the three conditions.

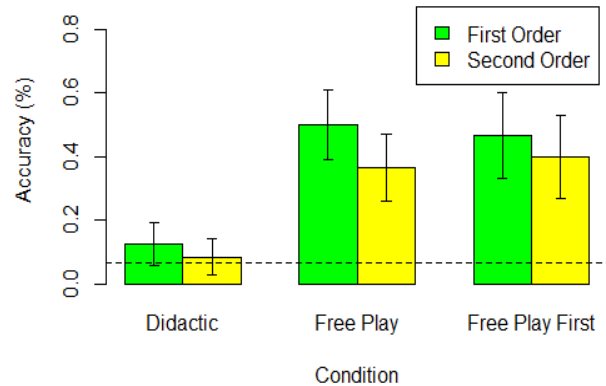


Figure 2: Percent accuracy for the conditions. Dashed line represents a conservative calculation of chance. Error bars represent standard error.

There was, however, a statistically significant difference between the performance of children who were in the free play condition ( $M= 0.4318$ ,  $SD= 0.4168$ ) and children who were in the didactic condition ( $M= 0.10$ ,  $SD=0.21$ ), as shown in Figure 2. More specifically, our analysis showed that the free play condition performed significantly better than the didactic condition ( $\beta= 2.259$ ,  $SE= 0.787$ ,  $p= 0.004$ ). Analysis of the exponentiated coefficients revealed that being in the free play condition increased the children’s odds of being correct by 856%. The free-play-first condition ( $M= 0.4333$ ,  $SD= 0.4169$ ) also performed significantly better than the didactic condition ( $\beta= 2.267$ ,  $SE= 0.840$ ,  $p= 0.007$ ), and being in the free play first condition increased the children’s odds of being correct by 865%. However, there was no significant difference between the performance of children in the free-play-first condition and the free play condition. In addition, we analyzed whether the amount of negative evidence received was predictive of performance in the free play and free-play-first conditions. It was found that the amount of negative evidence did not have a significant effect on performance ( $\beta= -0.005$ ,  $SE= 0.018$ ,  $p= 0.787$ ), indicating that the extent to which children received negative evidence during play did not influence their performance during the generalization test trials.

The best fit model was also found by comparing various models that included potential predictors of performance such as condition, sex, age, rule, and amount of negative and positive feedback. Through model comparisons, it was found that the best fit model predicted accuracy from condition ( $\chi^2= 13.72$ ,  $df= 2$ ,  $p= 0.001$ ). This model outperformed all other models, including the null model.

A conservative value for chance was also calculated by considering all possible two block combinations that could be placed on the machine and then calculating the probability of

placing the correct blocks. This resulted in a value of 0.067 for chance performance. Children in the free play condition were significantly more likely to choose the correct blocks compared to chance,  $t(43) = 4.834$ ,  $p = 0.00002$ , as were children in the free-play-first condition,  $t(29) = 3.985$ ,  $p = 0.0004$ . Children in the didactic condition, however, were not significantly more likely to choose the correct blocks over chance,  $t(47) = 0.842$ ,  $p = 0.404$ .

## Discussion

In the current study, we demonstrate that 3- and 4-year-old children can successfully acquire fairly complex causal generalizations through free play. They independently generated evidence that allowed them to understand the causal system that they were presented with, and they formed higher-order generalizations at a level above chance. More strikingly, children's learning in the free play condition and the free-play-first condition was superior to that of children in the didactic condition. Just as in the study by Sim and Xu (in press), children were equally successful in learning first-order and second-order generalizations during the course of free play, and there was no difference in their performance when it came to learning the shape or the color rule. It is interesting to note that there were more children in the free-play-first condition who had been exposed to the color rule rather than the shape rule, which may indicate the potential influence of a shape bias. However, we did not find an overall difference in performance between children who were exposed to the color rule vs. the shape rule.

We also sought to determine if there were scenarios in which children benefited more from learning through free play than through direction instruction by an experimenter. We found that children who engaged in free play performed significantly better at test than children in the didactic condition. Even children who were unable to activate the machines during play but who later observed an experimenter doing so performed significantly better than those assigned to the didactic condition, suggesting that the former group also benefited from engaging in free play. To the best of our knowledge, this study presents the first evidence that children can learn about a causal system more effectively through play than through training. So why was learning more effective under free play?

One possible reason for this difference between the two conditions is that children in the free play condition were able to engage actively with the materials, whereas children in the didactic condition played a passive role in learning about the machines, observing an experimenter activate them but never activating the machines themselves. Sobel and Sommerville (2010) found that four-year-old children learned a causal structure more accurately when they were given some time to engage with a causal system. Likewise, McCormack et al. (2016) showed that children who acted out interventions on a causal system following specific directions from an experimenter performed better than children who witnessed the same interventions but watched as the interventions were performed by an experimenter. Together, these studies

suggest that children may benefit more from intervening on a causal system, rather than observing an experimenter do so. This may explain why the free-play-first condition performed significantly better than the didactic condition, even though the children never successfully activated the machines while playing.

It is also possible that the children in the didactic condition struggled because they had different hypotheses about how the machines worked which were not contradicted by the evidence they witnessed (e.g., they may have thought that each of the two blocks would activate the machine by itself). Children in the free play condition and in the free-play-first condition, in contrast, had the opportunity to carefully test their own hypotheses, particularly through the generation of negative evidence. In other words, children in these conditions were able to see which combinations of blocks would make the machine go, *as well as* which combinations of blocks would not make the machine go. However, we did not find in our additional analyses that the amount of negative evidence that children generated was predictive of their performance at test, suggesting that there was something more to the evidence that the children saw during free play that assisted them in forming the correct generalizations. Further research is still necessary to understand the differences we found between the free play and didactic condition in our study. One worthwhile direction is to conduct an additional "yoked" didactic condition, where the experimenter presents children with evidence that was generated by children from the free play condition. This additional condition will clarify whether the differences found in the present study can be attributed solely to the difference in the quality of evidence between conditions.

We note that the current findings also appear to differ from those of other studies comparing the performance of children in free play and training/didactic conditions for other kinds of tasks. For example, Klahr and Nigam (2004) compared discovery learning, which they defined as learning that children engaged in by themselves without the assistance or feedback from a teacher, to direct instruction in third- and fourth-grade children for designing unconfounded experiments, which are experiments that clearly reveal the effect of a particular variable. The researchers found that children in the direct instruction condition performed significantly better than children in the discovery learning condition. However, in this particular study, children in the direct instruction condition were engaged in designing and manipulating variables during training as well, while children in the didactic condition in the current study engaged with the materials more passively, observing as an experimenter taught them how the machines worked. Another difference between the two studies is that in the study by Klahr and Nigam (2004), children in the discovery learning condition were given time to explore the ramp and marbles and design experiments, however they were not provided with any negative or positive feedback. In our study, on the other hand, children in the free play condition were provided with both negative and positive evidence, since the machine only made

a sound when activated correctly. Although the finding from Klahr and Nigam (2004) is sometimes used as evidence for the benefits of direct instruction for teaching children science, our study suggests that children have the potential to learn causal systems effectively through play when provided with useful feedback, even if the feedback does not come from an instructor.

Previous work has demonstrated that children attending child-centered preschools, where free play and child initiative are highly encouraged, were more motivated to learn, showed more pride in their accomplishments, and claimed to be less worried than children attending didactic, highly academic preschools (Stipek, Feiler, Daniels, & Milburn, 1995). The current study extends these results by showing that children can learn effectively through free play, and under some conditions, learning within a free play context may be better than learning in a didactic context. However, it is important to note that free play is just one aspect of the child-centered instructional approach in preschools, and our results cannot speak directly to any potential learning differences between the two instructional approaches.

In summary, the present study provides evidence that young children can learn effectively through free play, and this learning might be better than the learning achieved through direct demonstration. Our results provide one source of empirical evidence on why play is important for children and why unstructured play should be incorporated into school curriculums. This study also suggests that there is merit to child-centered learning in preschools, as it appears that children are able to learn through play and that they are able to successfully engage in active learning.

### Acknowledgments

We thank the Berkeley Early Learning Lab (BELL) for their help in testing and recruitment, as well as the parents and children for their participation. Mahal and Sim are joint first authors.

### References

- Chalnick, A., & Billman, D. (1988). Unsupervised learning of correlational structure. *Proceedings of the tenth annual conference of the cognitive science society* (pp. 510-516). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Castro, R. M., Kalish, C., Nowak, R., Qian, R., Rogers, T., & Zhu, X. (2009). Human active learning. In *Advances in neural information processing systems*, 241-248.
- Corno, L., & Mandinach, E. B. (1983). The role of cognitive engagement in classroom learning and motivation. *Educational Psychologist*, 18(2), 88-108.
- Davey, C., & Lundy, L. (2011). Towards Greater Recognition of the Right to Play: An Analysis of Article 31 of the UNCRC. *Children & Society*, 25, 3-14.
- Hofferth, S. L., & Sandberg, J. F. (2001). How American Children Spend Their Time. *Journal of Marriage and Family*, 63, 295-308.
- Klahr, D., & Nigam, M. (2004). The equivalence of learning paths in early science instruction: Effects of direct instruction and discovery learning. *Psychological Science*, 15(10), 661-667.
- Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, 143(1), 94-122.
- McCormack, T., Bramley, N., Frosch, C., Patrick, F., & Lagnado, D. (2016). Children's use of interventions to learn causal structure. *Journal of Experimental Child Psychology*, 141, 1-22.
- Metcalf, J., & Kornell, N. (2005). A region of proximal learning model of study time allocation. *Journal of Memory and Language*, 52(4), 463-477.
- Pepler, D. J., & Ross, H. S. (1981). The Effects of Play on Convergent and Divergent Problem Solving. *Child Development*, 52(4), 1202-1210.
- Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, 10(3), 322-332.
- Sim, Z., & Xu, F. (in press). Learning Higher-Order Generalizations through Free Play: Evidence from Two- and Three-Year-Old Children. *Developmental Psychology*.
- Smith, P. K., & Dutton, S. (1979). Play and training in direct and innovative problem solving. *Child Development*, 50(3), 830-836.
- Sobel, D. M., & Kushnir, T. (2006). The importance of decision making in causal learning from interventions. *Memory & Cognition*, 34(2), 411-419.
- Sobel, D. M., & Sommerville, J. A. (2010). The importance of discovery in children's causal learning from interventions. *Front. Psychol.*, 1, 176.
- Stipek, D., Feiler, R., Daniels, D., & Milburn, S. (1995). Effects of different instructional approaches on young children's achievement and motivation. *Child Development*, 66(1), 209-223.
- Walker, C. M., & Gopnik, A. (2014). Toddlers infer higher-order relational principles in causal learning. *Psychological science*, 25(1), 161-169.
- Weisberg, D.S., Hirsh-Pasek, K., & Golinkoff, R. M. (2013). Guided play: Where curricular goals meet a playful pedagogy. *Mind, Brain, and Education*, 7(2), 104 - 112.
- Weisberg, D.S., Hirsh-Pasek, K., Golinkoff, R. M., & McCandliss, B. D. (2014). *Mise en place*: Setting the stage for thought and action. *Trends in Cognitive Sciences*, 18(6), 276 - 278.
- Whitebread, D., Coltman, P., Jameson, H., & Lander, R. (2009). Play, cognition and self-regulation: What exactly are children learning when they learn through play? *Educational and Child Psychology*, 26(2), 40-52.