

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Can Adults Revise Their Core Beliefs about Objects?

Permalink

<https://escholarship.org/uc/item/3d12384w>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Liu, Rongzhi

Xu, Fei

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Can Adults Revise Their Core Beliefs about Objects?

Rongzhi Liu (rongzhi_liu@berkeley.edu), Fei Xu (fei_xu@berkeley.edu)

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

Abstract

A set of fundamental principles governs our reasoning about objects since infancy. Studies have shown that adults are surprised when they observe apparent violations of these principles, which might prime them to learn from the violations and update their beliefs. However, little is known about whether these principles can be revised given counterevidence. In the present experiments, we demonstrate that although adults have strong prior beliefs about these principles, they can revise these beliefs in a specific, virtual world, when they observe multiple pieces of counterevidence.

Keywords: belief revision; core knowledge; intuitive physics

Introduction

The Core Knowledge view (Spelke, 1988, 2000; Spelke & Kinzler, 2007) argues that human infants are endowed with core knowledge systems – a small number of systems of domain-specific knowledge, each accompanied by a set of principles. Later in development, infants and children construct intuitive theories based on these systems, such as intuitive physics and intuitive psychology (e.g., Carey, 1985; Gopnik & Meltzoff, 1997; Wellman & Gelman, 1992).

One of the core knowledge systems guides how we represent and reason about objects. A set of principles in this system emerges by 2.5 to 6 months of age. These principles include solidity – objects cannot occupy the same space as other objects (Spelke et al., 1992), continuity – objects exist and move continuously in time and space (Anguiar & Bailargeon, 1999), cohesion – objects move as connected and bounded wholes (Anguiar & Baillargeon, 1999), and contact – objects do not interact at a distance (Leslie & Keeble, 1987). These principles support further learning in the physical domain (see Baillargeon, 2008 for a review). These principles also persist into adulthood. Adults' ability to track multiple, independently moving objects is disrupted when the objects violate the principle of continuity (Scholl & Pylyshyn, 1999) or cohesion (Scholl et al., 2001; vanMarle & Scholl, 2003).

Thus, these core principles about objects are early emerging or innately given, they support learning about the physical domain throughout development, and they continue to guide reasoning about the physical world in adulthood. Yet one of the hallmarks of human learning is that beliefs can be revised given new evidence (Chater & Oaksford, 2008; Gopnik & Wellman, 2012; Tenenbaum et al., 2011; Xu, 2019; Ullman & Tenenbaum, 2020). Are these earliest-emerging core principles about objects also subject to revision once we acquire them? If adults are given enough evidence that violates these principles, will they rationally update their beliefs? This is the focus of the current studies.

Previous studies have provided many demonstrations that adults are sensitive to core knowledge principles about

objects and belief revision is a powerful learning mechanism for humans. For example, adults are surprised by apparent violations of the core physical principles. In one study (Smith et al., 2020), participants observed events that violated the continuity principle (e.g., objects disappearing and appearing; objects moving from behind one screen to another without traveling through the gap between the two screens), and the solidity principle (e.g., objects moving pass other objects). Participants reported that these events were more surprising compared to control events that did not violate any physical principles. When they were asked why these events were surprising, they referred to the violations of the corresponding principles.

Surprise provides opportunities for learning. Studies with infants and children have shown that surprise following violations of the core physical principles leads to exploration and learning. Observing an object violate a physical principle prompted infants to explore that object (Stahl & Feigenson, 2015) and search for an explanation of the violation (Perez & Feigenson, 2022). Infants and children showed enhanced learning for properties and novel words related to the object that violated physical principles (Stahl & Feigenson, 2015; 2017). These findings suggest that when adults are surprised by violations of the core physical principles, they might also learn from the violations, and update their beliefs about the physical principles.

In more complex domains that go beyond the basic physical principles, adults learn from new evidence and update their beliefs. For instance, adults' misconceptions about a specific domain (i.e., the variables that affect the pendulum period) can be revised when they were given evidence that contradicted their initial beliefs (Masnick et al., 2017). A set of studies has shown that adults' causal learning rationally integrates the strength of their prior knowledge about the probability of different forms of causal relationships and the strength of the evidence (Lucas & Griffiths, 2010; Griffiths et al., 2011; Lucas et al., 2014). Lastly, adults can reason about complex scenes of object interactions using probabilistic mental simulations based on intuitive physics (Battaglia et al., 2013), and they update their beliefs about object properties given new observations and simulation outcomes (Hamrick et al., 2016; Allen et al., 2020).

However, one question still remains. Are our most fundamental beliefs about objects, already present in infancy, revisable given counterevidence? Past studies have shown that preschoolers can revise their beliefs about the contact, continuity, and solidity principles when given counterevidence (Liu & Xu, 2021; Kushnir & Gopnik, 2007). Adults might have stronger prior beliefs about these principles since they have observed more evidence consistent with these principles. Can adults also revise their beliefs about these

principles when given counterevidence? In two experiments, participants observed events that supported or violated these principles, or they did not receive any new evidence about these principles. Then, they made predictions about the outcomes of new events. We hypothesized that compared to those who saw the belief-consistent evidence and those who did not receive new evidence, participants who saw the belief-violating evidence would be more likely to predict outcomes that are inconsistent with the principles.

Experiment 1

Method

Participants Forty-seven adults (mean age = 30 years; range = 18 to 55; $SD = 9.2$; 25 females) participated in the experiment on Prolific. Participants provided written informed consent prior to participating in the experiment. They completed a 20-minute survey for which they were paid \$3.20.

Stimuli and Procedure Participants were instructed to watch some videos and observe how objects move. Then, they were asked to make predictions about new events.

Participants were randomly assigned to one of the two conditions, the Belief Consistent (BC) condition ($n = 23$) and the Belief Violation (BV) condition ($n = 24$). They were tested on 3 principles, Contact, Continuity, and Solidity, in counter-balanced orders. For each principle, there were 4 familiarization trials and 4 test trials (2 easy test trials and 2 hard test trials; order counterbalanced). The familiarization trials in the BC condition displayed events that were consistent with the principle and those in the BV condition displayed events that violated the principle. In test trials, participants chose between the *Belief Consistent (BC) response* and the *Belief Violation (BV) response*. They never received feedback about whether their choices were correct or incorrect.




Figure 1: Events shown in the familiarization trials and test trials for the Continuity principle in Experiment 1.

Continuity principle. In the familiarization trials, two orange screens appeared side by side, with a gap in between. An object disappeared behind one of the screens. Then, the screens were removed. The object was either at the location of the screen that the object disappeared behind (BC condition) or at the location of the other screen (BV condition) (Figure 1). The object was different in each trial.

In the easy test trials, a new object disappeared behind one of the orange screens. A blue triangle and a green triangle indicated the screen that the object disappeared behind (the *BC response*) and the other screen (the *BV response*) (Figure 1). Participants chose the location they believed they would find the object. In the hard test trials, a red door and a yellow door appeared. A new object disappeared behind one of the doors (Figure 1). Participants chose the location they believed they would find the object, either the door that the object disappeared behind (the *BC response*) or the other door (the *BV response*).

Solidity principle. In the familiarization trials, a dark grey wall appeared and rotated 180 degrees to show that there was no hole in the wall. A green screen was placed in front of the wall and occluded the lower half of the wall. An object moved behind the screen. Then, the screen was removed. The object was either on the side of the wall that it went behind (BC condition) or on the other side of the wall (BV condition) (Figure 2). A different object was used in each trial.




Figure 2: Events shown in the familiarization trials and test trials for the Solidity principle in Experiment 1.

In the easy test trials, a new object moved behind the green screen. A purple heart and an orange heart indicated the side of the wall that the object went behind (the *BC response*) and the other side of the wall (the *BV response*) (Figure 2). Participants chose the location they believed they would find the object. In the hard test trials, two doors (side by side, with no gap in between) were placed in front of the wall and occluded the lower half of the wall. A new object moved behind the doors (Figure 2). Participants chose the location they believed

they would find the object, either the side of the wall that the object went behind (the *BC response*) or the other side of the wall (the *BV response*).

Contact principle. In the familiarization trials, participants were shown a blue box that can play music. In each trial, an object was placed either on the toy (BC condition) or above the toy (BV condition), and immediately the toy lit up and played music for 5 seconds (Figure 3). A different object was used to activate the toy in each trial.




Figure 3: Events shown in the familiarization trials and test trials for the Contact principle in Experiment 1.

In the easy test trials, a new object was placed next to the blue box. Participants were told that the object can activate the toy. A red star and a yellow star indicated the location on the toy (the *BC response*) and the location above the toy (the *BV response*) (Figure 3). Participants chose the location where they would place the object to activate the toy. In the hard test trials, a new, brown box and a new object appeared (Figure 3). The participants were told that the brown box is another music toy, and the object can activate the toy. Again, participants chose the location where they would place the object to activate the toy.

Results

The proportion of *BV response* by condition and principle is shown in Figure 4. We used logistic regression to predict participants' binary choice (*BV response* = 1, *BC response* = 0) from condition, principle, and test trial type, with random intercepts for participants. The best-fitting model included condition as the only predictor. Participants were more likely to choose the *BV response* in the BV condition than in the BC condition ($\beta = 14.79$, $SE = 3.08$, $p < .001$). This model outperformed the null model ($AIC_{\text{condition}} = 305.09$, $AIC_{\text{null}} = 411.40$, $p < .001$), and more complex models that included principle and trial type (easy vs. hard test trials).

Next, we looked at participants' choice for each principle separately, and found that the main effect of condition was significant for each principle. Participants were more likely to select the *BV response* in the BV condition than in the BC condition for Contact ($\beta = 18.97$, $SE = 3.73$, $p < .001$), Continuity ($\beta = 14.50$, $SE = 3.92$, $p < .001$), and Solidity ($\beta = 15.15$, $SE = 3.68$, $p < .001$).




Figure 4: The proportion of trials that participants selected the *BV response* by condition and principle, in Experiment 1. The dashed line indicates chance selection (.5), and the error bars indicate bootstrapped 95% CIs.

Discussion

These results suggest that adults can revise their beliefs about the most fundamental principles that govern object reasoning in a specific context. After observing evidence consistent with these principles, adults predicted that the outcomes of new events would be consistent with the principles. But after observing evidence violating these principles, they were more likely to predict outcomes inconsistent with the principles. Moreover, adults' performance did not differ in the easy and hard test trials, suggesting that they also generalized their revised beliefs to new situations.

In the next experiment, we aimed to replicate these findings. In addition, we measured adults' prior beliefs about these principles and tested the effect of belief-violating evidence on their prior beliefs; we increased the strength of the belief-violating evidence and tested its effect on their beliefs. We also assessed participants' interpretations of the evidence to see if they accepted the counterevidence.

Experiment 2

Method

Participants Sixty adults (mean age = 33 years; range = 18 to 54; $SD = 9.41$; 35 females) participated in the experiment on Prolific. Participants provided written informed consent prior to participating in the experiment. They completed a 25-minute survey for which they were paid \$4.

Stimuli and Procedure The procedure of Experiment 2 was similar to that of Experiment 1, with a few modifications. First, the events used in the Contact and the Solidity principles were slightly modified (see below for details). Second, we added a third, Baseline condition, where participants did not receive any new evidence that supported or violated the principles. Participants were randomly assigned to the Baseline condition ($n = 20$), the BC condition ($n = 19$), and the BV

condition ($n = 21$). Third, we increased the strength of the evidence; participants were shown 6 familiarization trials for each principle. Last, at the end of the study, participants in the BC and BV conditions were asked to explain one of the familiarization events for each principle.

Continuity principle. The events used in the familiarization trials in the BC and BV conditions and the test trials in all 3 conditions were the same as in Experiment 1. In the familiarization trials of the Baseline condition, the screens were not removed after the object disappeared behind one of the screens (Figure 5). Thus, participants did not observe the location of the object.

For the explanation question, participants were asked to explain why the object appeared at the respective locations when the screens were removed.




Figure 5: Events shown in the familiarization trials and test trials for the Continuity principle in Experiment 2.

Solidity principle. The events used for the Solidity principle were slightly modified so they were more similar to previous infant studies (e.g., Stahl & Feigenson, 2015). Two dark grey walls appeared and rotated 180 degrees to show that there was no hole on the walls. An object went down a ramp and moved behind a screen. In familiarization trials, when the screen was removed, the object was either stopped before the first wall (BC condition) or it went past the first wall and appeared between the two walls (BV condition) (Figure 6). In the Baseline condition, the screen was not removed so that participants could not observe the location of the object. In the test trials, participants chose the location they believed they would find a new object that went down the ramp. The location before the first wall was the *BC response* and the location between the two walls was the *BV response* (Figure 6).

For the explanation question, participants were asked to explain why the object appeared at the respective locations when the screen was removed.




Figure 6: Events shown in the familiarization trials and test trials for the Solidity principle in Experiment 2.

Contact principle. Object launching events were used for the Contact principle, making these more similar to previous infant studies (e.g., Leslie & Keeble, 1987). In the familiarization trials, participants were told that a yellow car would launch various objects. In each trial, the yellow car moved toward an object and launched the object either by contacting it (BC condition) or at a distance (BV condition). In the Baseline condition, a screen blocked the view between the yellow car and the object so that participants could not see whether the yellow car contacted the other object or not (Figure 7). The yellow car launched a different object in each trial.




Figure 7: Events shown in the familiarization trials and test trials for the Contact principle in Experiment 2.

In the easy test trials, participants were told the yellow car can launch a new object. A red star and a yellow star indicated the location right next to the object (the *BC response*) and the location at a distance (the *BV response*) (Figure 7). Participants chose the location where the yellow car should stop to launch the new object. In the hard test trials, participants were

told a new wheeled toy (e.g., the helicopter) can launch an object. Again, participants chose the location where the wheeled toy should stop to launch the object.

For the explanation question, participants were asked to explain why the yellow car launched the other object.

Results

The proportion of *BV response* by condition and principle is shown in Figure 8. We used logistic regression to predict participants' binary choice (*BV response* = 1, *BC response* = 0) from condition, principle, and test trial type, with random intercepts for participants. The best-fitting model included the interaction between condition and principle as predictors. Participants were more likely to choose the *BV response* in the BV condition than in the Baseline condition ($\beta = 8.79$, $SE = 1.38$, $p < .001$) and the BC condition ($\beta = 10.68$, $SE = 1.75$, $p < .001$); their choices did not differ between the Baseline and the BC conditions ($\beta = -0.17$, $SE = 1.49$, $p = .91$). In the Baseline condition, compared to the Continuity principle, participants were more likely to choose the *BV response* for the Contact principle ($\beta = 2.47$, $SE = 0.91$, $p = .007$) and the Solidity principle ($\beta = 1.93$, $SE = 0.92$, $p = .04$). In the BV condition, participants were less likely to choose the *BV response* for the Contact principle compared to the Continuity principle ($\beta = -3.58$, $SE = 0.69$, $p < .001$) and the Solidity principle ($\beta = -4.00$, $SE = 0.77$, $p < .001$). This model outperformed the null model ($AIC_{\text{condition} \times \text{principle}} = 333.91$, $AIC_{\text{null}} = 462.91$, $p < .001$), the model that did not include the interaction term ($AIC_{\text{condition} + \text{principle}} = 375.46$, $p < .001$), as well as more complex models that included trial type (easy vs. hard test trials).




Figure 8: The proportion of trials that participants selected the *BV response* by condition and principle, in Experiment 2. The dashed line indicates chance selection (.5), and the error bars indicate bootstrapped 95% CIs.

Next, we looked at participants' choice for each principle separately. For all 3 principles, participants were more likely to choose the *BV response* in the BV condition than in the Baseline condition (Continuity: $\beta = 15.84$, $SE = 4.20$, $p < .001$; Solidity: $\beta = 14.45$, $SE = 3.45$, $p < .001$; Contact: $\beta =$

18.27, $SE = 3.48$, $p < .001$), and their choices did not differ between the Baseline and the BC conditions (Continuity: $\beta = 0.56$, $SE = 2.49$, $p = .82$; Solidity: $\beta = -0.81$, $SE = 1.92$, $p = .67$; Contact: $\beta = -1.18$, $SE = 3.58$, $p = .74$).

Then, we compared participants' performance across the two experiments. We used logistic regression to predict participants' binary choice (*BV response* = 1, *BC response* = 0) from condition, principle, test trial type, and experiment, with random intercepts for participants. The best-fitting model included the interaction between condition and principle as predictors. Participants were more likely to choose the *BV response* in the BV condition than in the Baseline condition ($\beta = 8.93$, $SE = 1.27$, $p < .001$) and the BC condition ($\beta = 8.38$, $SE = 0.99$, $p < .001$); their choices did not differ between the Baseline and the BC conditions ($\beta = 0.81$, $SE = 1.21$, $p = .50$). In the Baseline condition, compared to the Continuity principle, participants were more likely to choose the *BV response* for the Contact principle ($\beta = 2.43$, $SE = 0.91$, $p = .007$) and Solidity principle ($\beta = 1.95$, $SE = 0.92$, $p = .03$). In the BV condition, participants were less likely to choose the *BV response* for the Contact principle compared to the Continuity principle ($\beta = -1.84$, $SE = 0.35$, $p < .001$) and the Solidity principle ($\beta = -2.37$, $SE = 0.39$, $p < .001$). This model outperformed the null model ($AIC_{\text{condition} \times \text{principle}} = 701.65$, $AIC_{\text{null}} = 882.65$, $p < .001$), the model that did not include the interaction term ($AIC_{\text{condition} + \text{principle}} = 739.06$, $p < .001$), as well as more complex models that included experiment or trial type (easy vs. hard test trials).

Table 1: Coding criteria and examples for explanations

Category	Criterion	E.g., Solidity
Accept Evidence	Accepted the violation of the target principle in the counterevidence.	"The car can go through the wall."
Explain Away	Explained the counterevidence with reasons that would not involve any violations of the target principle.	"The first wall was further away towards the back."
Pattern	Noted the pattern in the evidence.	"Because it will follow the sequence of events previously seen."
Other	Explanations that cannot be categorized into the first three categories.	"I don't know."

Lastly, for the explanation questions, 2 researchers coded participants' responses into different categories (the interrater reliability was excellent, Cohen's Kappa = .94; disagreements were resolved through discussion). In the BC condition, most responses (98.2%) referred to the principle itself to explain the evidence (the only response not referring to the principle was incomprehensible). In the BV condition, we categorized participants' explanations into four categories. The criteria for categorization and examples are shown in Table 1.

Table 2 shows the number of responses coded within each category for each principle. Logistic regression revealed that participants were more likely to choose the *BV response* for the principle if they provided “accept evidence” ($\beta = 1.61$, $SE = 0.52$, $p = .002$) or “pattern” ($\beta = 2.06$, $SE = 0.68$, $p = .003$) explanations for that principle, compared to if they provided “other” explanations. There was no difference between participants who provided “explain away” and “other” explanations ($\beta = 0.66$, $SE = 0.44$, $p = .14$). In addition, participants who provided “accept evidence” explanations for all 3 principles were more likely to choose the *BV response* overall compared to participants who provided “accept evidence” explanation for none of the principles ($\beta = 1.95$, $SE = 0.80$, $p = .015$). There was no difference between participants who provided “accept evidence” explanation for zero and one principle ($\beta = 0.64$, $SE = 0.79$, $p = .42$), or for zero and two principles ($\beta = 0.22$, $SE = 0.66$, $p = .73$).

Table 2: Number of responses by category and principle

	Contact	Continuity	Solidity
Accept Evidence	6	7	7
Explain Away	9	5	6
Pattern	2	8	3
Other	4	1	5

Discussion

Experiment 2 showed that adults have strong prior beliefs about the contact, continuity, and solidity principles, but they can revise their prior beliefs in a specific context given counterevidence. When adults did not receive any new evidence about these principles and when they received evidence supporting these principles, they predicted that the outcomes of new events would be consistent with the principles. But when they received evidence violating these principles, they were more likely to predict outcomes inconsistent with the principles. Similar to Experiment 1, adults’ performance did not differ in the easy and hard test trials, suggesting that they were willing to generalize their revised beliefs to new situations. Participants’ explanations of the belief-violating evidence suggest that only a small portion of participants had accepted the counterevidence and revised their beliefs, and these participants were indeed more likely to predict outcomes that violated the principles. However, other participants simply learned from the statistical pattern of the evidence or explained away the counterevidence.

We also discovered some interesting differences across principles. First, when given no new evidence, participants were more likely to predict outcomes inconsistent with the principles for the contact and solidity principles compared to the continuity principle. This might suggest that adults have weaker prior beliefs about the contact and solidity principles. However, this might also be due to perceptual ambiguities in the stimuli (e.g., for solidity, some participants thought that the first wall was further towards the back).

Second, across two experiments, we found that participants were less likely to revise their beliefs about the contact principle compared to the other two principles. One possible explanation for this result is that negative evidence (e.g., placing an object on the toy does not activate it) is necessary to revise the contact principle. Indeed, Kushnir & Gopnik (2007) found that preschoolers can revise their beliefs about the contact principle when shown contrasting evidence (i.e., placing an object above the toy made it go but placing an object on the toy did not). Moreover, we might have observed more violations of the contact principle in real life (e.g., magnets, remote controls). This leads to 2 possibilities. First, we might have weaker prior beliefs about the contact principle, making this principle easier to be revised in a specific context. Second, we might believe that the contact principle is more probabilistic, making this principle harder to revise because the counterevidence could be seen as exceptions and dismissed. Our results are consistent with the second possibility. Future studies can probe whether there are indeed more violations of the contact principle in the real world, and whether the contact principle is believed to be more probabilistic than others.

General Discussion

The present study reports the first systematic investigation of whether adults can revise their beliefs about the most fundamental principles governing object reasoning. We found that although adults have strong prior beliefs about these principles, they were able to revise these beliefs in a specific, virtual world with just a few pieces of counterevidence, and generalize their revised beliefs to new contexts in this virtual world.

Adults and children often immerse themselves in fictional worlds (e.g., movies, novels, magic shows), where they suspend their beliefs about the core physical principles. Past research showed that our judgments about the fictional worlds (e.g., adults’ estimation of the efforts required to cast spells that cause physical violations; McCoy & Ullman, 2019) are still guided by our intuitive theories about physics. The present study is the first to suggest that adults have the potential to entertain a system of physical principles that is completely different from our intuitive theories in the real world.

In the present study, only a small portion of adults had accepted the counterevidence and revised their beliefs in this virtual world, while others responded based on statistical learning or explained away the counterevidence. In an ongoing study, we examine whether adults are more likely to revise their beliefs when given more compelling counterevidence (i.e., photorealistic, three-dimensional stimuli made with Blender). The new stimuli also rule out many perceptual ambiguities (e.g., for solidity, it is clear that the first wall is not further towards the back). In the present study, adults generalized their revised beliefs to slightly different contexts. We further probe the extent to which adults are willing to generalize the revised beliefs in the ongoing study by asking them to make predictions about events that are more different from the original events.

Acknowledgments

We thank the participants who participated in this study, as well as members of the Berkeley Early Learning Lab for their help with data collection and thoughtful discussions.

References

- Aguiar, A., & Baillargeon, R. (1999). 2.5-Month-Old infants' reasoning about when objects should and should not be occluded. *Cognitive Psychology*, 39(2), 116–157.
- Allen, K. R., Smith, K. A., & Tenenbaum, J. B. (2020). Rapid trial-and-error learning with simulation supports flexible tool use and physical reasoning. *Proceedings of the National Academy of Sciences*, 117(47), 29302–29310.
- Baillargeon, R. (2008). Innate ideas revisited: For a principle of persistence in infants' physical reasoning. *Perspectives on Psychological Science*, 3(1), 2–13.
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110(45), 18327–18332.
- Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press/Bradford Books.
- Chater, N., & Oaksford, M. (Eds.). (2008). *The probabilistic mind: Prospects for Bayesian cognitive science*. Oxford University Press, USA.
- Gopnik, A., & Meltzoff, A. (1997). *Words, thoughts and theories*. Cambridge, MA: MIT Press.
- Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, 138(6), 1085–1108.
- Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2011). Bayes and Blickets: Effects of knowledge on causal induction in children and adults. *Cognitive Science*, 35(8), 1407–1455.
- Hamrick, J. B., Battaglia, P. W., Griffiths, T. L., & Tenenbaum, J. B. (2016). Inferring mass in complex scenes by mental simulation. *Cognition*, 157, 61–76.
- Kushnir, T., & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology*, 43(1), 186–196.
- Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, 25(3), 265–288.
- Liu, R. & Xu, F. (2021). Revising core beliefs in young children. In T. Fitch, C. Lamm, H. Leder, & K. Tessmar (Eds.), *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*. Cognitive Science Society.
- Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*, 131(2), 284–299.
- Lucas, C. G., & Griffiths, T. L. (2009). Learning the form of causal relationships using Hierarchical Bayesian Models. *Cognitive Science*, 34(1), 113–147.
- Masnack, A. M., Klahr, D., & Knowles, E. R. (2017). Data-driven belief revision in children and adults. *Journal of Cognition and Development*, 18(1), 87–109.
- McCoy, J., & Ullman, T. (2019). Judgments of effort for magical violations of intuitive physics. *PLOS ONE*, 14(5), e0217513.
- Perez, J., & Feigenson, L. (2022). Violations of expectation trigger infants to search for explanations. *Cognition*, 218, 104942.
- Scholl, B. J., & Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive Psychology*, 38(2), 259–290.
- Scholl, B. J., Pylyshyn, Z. W., & Feldman, J. (2001). What is a visual object? Evidence from target merging in multiple object tracking. *Cognition*, 80(1–2), 159–177.
- Smith, K. A., Mei, L., Yao, S., Wu, J., Spelke, E., Tenenbaum, J. B., & Ullman, T. D. (2020). The fine structure of surprise in intuitive physics: When, why, and how much? In S. Denison, M. Mack, Y. Xu, & B. C. Armstrong (Eds.), *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*. Cognitive Science Society.
- Spelke, E. S. (1988). The origins of physical knowledge. In L. Weiskrantz (Ed.), *Thought without language*. Clarendon Press/Oxford University Press.
- Spelke, E. S. (2000). Core knowledge. *American Psychologist*, 55(11), 1233–1243.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, 10(1), 89–96.
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological review*, 99(4), 605.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348(6230), 91–94.
- Stahl, A. E., & Feigenson, L. (2017). Expectancy violations promote learning in young children. *Cognition*, 163, 1–14.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to Grow a Mind: Statistics, Structure, and Abstraction. *Science*, 331(6022), 1279–1285.
- Ullman, T. D., & Tenenbaum, J. B. (2020). Bayesian models of conceptual development: Learning as building models of the world. *Annual Review of Developmental Psychology*, 2, 533–558.
- vanMarle, K., & Scholl, B. J. (2003). Attentive Tracking of Objects Versus Substances. *Psychological Science*, 14(5), 498–504.
- Wellman, H. M., & Gelman, S. A. (1992). Cognitive development: Foundational theories of core domains. *Annual review of psychology*, 43(1), 337–375.
- Xu, F. (2019). Towards a rational constructivist theory of cognitive development. *Psychological Review*, 126(6), 841–864.