Expectancy Violations Promote Learning in Young Children

Aimee E. Stahl¹ & Lisa Feigenson²

¹ The College of New Jersey, 2000 Pennington Road, Ewing, NJ 08628; stahla@tcnj.edu

² Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218; feigenson@jhu.edu

Corresponding Author: Aimee E. Stahl, Psychology Department, The College of New Jersey, 2000 Pennington Road, Ewing, NJ 08628; stahla@tcnj.edu

Abstract

Children, including infants, have expectations about the world around them, and produce reliable responses when these expectations are violated. However, little is known about how such expectancy violations affect subsequent cognition. Here we tested the hypothesis that violations of expectation enhance children's learning. In four experiments we compared 3- to 6-year-old children's ability to learn novel words in situations that defied versus accorded with their core knowledge of object behavior. In Experiments 1 and 2 we taught children novel words following one of two types of events. One event violated expectations about the spatiotemporal or featural properties of objects (e.g., an object appeared to magically change locations); the other event was almost identical, but did not violate expectations (e.g., an object was visibly moved from one location to another). In both experiments we found that children robustly learned when taught after the surprising event, but not following the expected event. In Experiment 3 we ruled out two alternative explanations for our results. Finally, in Experiment 4, we asked whether surprise affects children's learning in a targeted or a diffuse way. We found that surprise only enhanced children's learning about the entity that had behaved surprisingly, and not about unrelated objects. Together, these experiments show that core knowledge – and violations of expectations generated by core knowledge – shapes new learning.

Keywords: children, word learning, surprise, expectations, object knowledge

Expectancy Violations Promote Learning in Young Children

1. Introduction

Humans are promiscuous predictors – we readily form expectations about countless aspects of the world around us, and note when these expectations are violated. Indeed, the concept of expectancy violation has been central in the effort to understand the developmental roots of human cognition. By documenting whether, and at what ages, children respond to events that adults find surprising, researchers have made great progress in characterizing the nature and trajectory of children's knowledge. The logic is that if infants or young children differentiate a surprising event from an expected one (assuming perceptual factors are controlled for), they must already have had expectations in place to drive the differentiating behavior. In this sense, children's responses to surprising events have been an invaluable methodological tool for uncovering early knowledge.

But what else might surprise reveal about cognition? Here we explore the idea that violations of expectation not only help to identify children's extant knowledge, but also shed light on when and how new learning occurs.

1.1 Violation of Expectation in Infants and Children

The strategy of using surprising events as a tool to characterize thinking has revealed much about infant cognition. Young infants, despite being unable to verbally express their knowledge, look reliably longer at events that defy expectations than at similar events that accord with expectations. Perhaps the earliest in-depth use of expectancy violations was to characterize infants' knowledge about objects. Experiments using violation-of-expectation paradigms revealed that even before they have experience grasping or manipulating objects,

infants have a suite of expectations about how objects can (and cannot) behave. For example, infants as young as 2 months old expect objects to exist continuously – they look longer when objects appear to violate continuity by disappearing from a hiding location, or vanishing from one location and reappearing in another, compared to when objects remain where they were hidden (e.g., Baillargeon, 1986; Baillargeon & Graber, 1987; Baillargeon, Spelke, & Wasserman, 1985; Spelke, Breinlinger, Macomber, & Jacobson, 1992; Spelke, Kestenbaum, Simons, & Wein, 1995; Wilcox, Nadel, & Rosser, 1996; Wynn, 1992). Infants also expect objects to fall if unsupported – they look longer when objects hover in mid-air than when objects rest on a supporting surface (Baillargeon & Hanko-Summers, 1990; Needham & Baillargeon, 1993). And infants expect objects to affect each other through contact – they look longer when one object launches another without touching than when one launches another with contact (Leslie & Keeble, 1987; Spelke, 1990).

While infants' knowledge of objects is arguably the most thoroughly studied domain of core knowledge, infants' looking patterns have also revealed sophisticated knowledge about other aspects of the world. In the numerical domain, for example, infants look longer when a collection of five objects added to another collection of five objects yields only five objects as opposed to ten (McCrink & Wynn, 2004; see also Wynn, 1992). In the social domain, infants look longer when social agents approach individuals who were previously antisocial over those who were previously prosocial (e.g., Kuhlmeier, Wynn, & Bloom, 2003; Meristo & Surian, 2013), when agents switch goals as opposed to maintain consistent goals (e.g., Woodward, 1998), when agents behave in irrational compared to rational ways (e.g., György, Nádasdy, Csibra, & Biró, 1995), and when people appear to have knowledge of events they did not witness, compared to when they are ignorant of unseen events (e.g., Luo & Baillargeon, 2007;

Onishi & Baillargeon, 2005). Because this “core knowledge” of objects, quantities, and social agents can sometimes be observed in newborn infants (e.g., Izard, Sann, Spelke, & Streri, 2009), appears to be universal across human cultures (e.g., Dehaene, Izard, Pica, & Spelke, 2006; Everett 2005; Gordon, 2004), is shared with other species (e.g., Kunder, De Los Reyes, Taglang, Baruch, & German, 2010; Santos & Hauser, 2002), and emerges in newborn animals under controlled rearing conditions (e.g., Chiandetti & Vallortigara, 2010), it is often described as having evolutionary origins and being independent of specific learning experience (see Spelke & Kinzler, 2007 for review).

Findings like these highlight the sense in which infants’ longer looking to surprising events has been a major driver in the effort to characterize core knowledge. In addition, expectancy violations have been detected using other measures. Surprising events induce changes in infants’ facial expression (Camras et al., 2002), increase social referencing behavior (Walden, Kim, McCoy, & Karass, 2007), and trigger changes in brain activity (Berger, Tzur, & Posner, 2006; Wilcox, Bortfeld, Woods, Wruck, & Boas, 2005). Furthermore, in addition to violations of core knowledge (which are in place prior to the experimental session), infants detect violations to associative expectations acquired on-line over the course of a laboratory task. Infants not only look longer when these associative expectations are violated (e.g., Saffran, Aslin, & Newport, 1996; Stahl, Romberg, Roseberry, Golinkoff, & Hirsh-Pasek, 2014), but also exhibit differential brain responses to events they have recently established as surprising (Emberson, Richards, & Aslin, 2015; Kouider et al., 2015).

1.2 Violation of Expectation as an Opportunity for Learning

Taken together, this growing body of evidence shows that infants exhibit behavioral, social, and physiological responses when there is a mismatch between the expected and the observed. These responses have empowered researchers to draw conclusions about what expectancies infants have, and what the origins of these expectancies might be. Yet it is noteworthy that, despite their importance for understanding the nature of early knowledge, little is known about whether and how expectancy violations might also affect children's subsequent thinking.

Why might infants look longer, and experience physiological changes, in response to violations of expectation? One possibility is that these responses reflect active changes in children's cognitive processing, triggered by surprising events. Such processing changes could be useful for the daunting learning problem that infants (and indeed, all learners) face. Given the overwhelming complexity of the natural environment – the number of objects, features, and events present in any given scene, coupled with the dynamic moment-by-moment changes in these – the learning space must somehow be constrained; otherwise the learner confronts an infinite number of possible updates to prior knowledge. Existing knowledge – and the predictions it generates – could help constrain this learning space. If learners allocate cognitive resources to entities or events that failed to behave as predicted, learning efficiency will increase.

Several studies have shown that, consistent with this prediction, older children use expectations to guide their explanations and direct the targets of their exploration (e.g., Bonawitz, van Schijndel, Friel, & Schulz, 2012; Chandler & Lalonde, 1994; Johnson & Harris, 1994; Phelps & Woolley, 1994; Rosengren & Hickling, 1994; Schulz, 2012). For instance, Bonawitz and colleagues (2012) asked how 5- to 7-year-old children's prior knowledge

influenced their behavior following belief-violating events. Children were categorized as either believing that objects should balance on their geometric center, or on their center of mass. They then saw an object that balanced in a way that either accorded with or violated their theory of balance, and were later given a choice to play with that very same apparatus, or an entirely novel toy. Children who saw the object balance in a way that was consistent with their beliefs opted to explore the novel toy, whereas children who saw the object balance in a way that was inconsistent with their beliefs preferentially explored the balancing apparatus. Moreover, children who had experienced evidence that violated their beliefs were more likely to appeal to a hidden cause for the surprising event (e.g., to suggest that a magnet was present). Thus, children's prior beliefs can mediate their exploratory choices (Bonawitz et al., 2012).

Other studies show that children can use beliefs constructed on-line over the course of an experiment to drive their exploration and explanations (e.g., Cook, Goodman, & Schulz, 2011; Legare, 2012; Legare, Gelman, & Wellman, 2010; Legare, Schult, Impola, & Souza, 2016; Schulz & Bonawitz, 2007; Schulz, Standing, & Bonawitz, 2008; van Schijndel, Visser, van Bers, & Raijmakers, 2015). In one study, children saw a novel object that, when placed on a box, always made the box light up; a different novel object never made the box light up. Next children saw each toy placed on a separate box, with neither box lighting up. When asked by the experimenter, "Why did this happen?" children were more likely to offer explanations for the event that was inconsistent with their new knowledge (i.e., children tried to explain the behavior of the toy that unexpectedly failed to activate the box) (Legare et al., 2010). In addition, children explored the toys in ways that reflected their explanations. For example, children who explained the inconsistent event by appealing to causal explanations (e.g., the toy or its batteries being broken) often tested hypotheses about the object's failure: they attempted to open the object to

see whether it had batteries, or combined the broken object with a functioning one (Legare, 2012).

Infants, too, selectively and rationally explore toys that fail to produce expected effects (Baldwin, Markman, & Melartin, 1993; Gweon & Schulz, 2011). Sixteen-month-olds who saw an adult successfully activate a toy, but then failed to activate that same toy themselves, frequently sought parents' help rather than choose to play with another toy. But when infants saw an adult successfully activate a toy, and then received a different toy that did not activate, they were more likely to choose another toy with which to play. These findings suggest that infants rationally generated an appropriate explanation for why the toy failed to activate as expected, and used this explanation to guide their behavior (Gweon & Schulz, 2011).

Finally, recent work shows that infants can use pre-existing core knowledge to guide exploration and motivate hypothesis-testing behaviors. We showed 11-month-old infants a single event that accorded with core expectations about object behavior (e.g., a moving ball was stopped by a wall), or an event that was nearly identical but violated core expectations (the ball appeared to pass through the wall). When infants were then given the opportunity to explore the object from the event (ball) and an entirely new object (e.g., a car), infants who had seen the expected outcome showed a novelty preference (i.e., they played mostly with the car). In contrast, infants who saw the surprising outcome explored the object that had violated their expectations over the novel object (i.e., they played mostly with the ball). Infants also engaged in hypothesis testing behaviors that reflected the particular kind of violation seen. Infants who saw an object appear to violate solidity tended to bang the object on a solid surface, whereas infants who saw an object appear to violate support by floating in mid-air repeatedly dropped the object.

Hence, preverbal infants selectively seek information from objects that violated core expectations, and explore the objects' properties in a targeted way (Stahl & Feigenson, 2015).

1.3 Do Violations of Expectations Enhance Children's Learning?

Although the above evidence shows that children and even infants change their exploratory actions in response to expectancy violations, it does not tell us whether expectancy violations also change learning itself, as opposed to changing the opportunities children have for learning (e.g., the exploratory behaviors they produce). Over and above guiding the objects children chose to explore and their interactions with those objects, seeing a violation of expectation might affect the efficiency with which new information is acquired. Children might be better at learning about an object that has behaved in surprising ways, relative to an object that behaved as expected.

In recent work we addressed this question with infants, using a looking-time paradigm. We showed 11-month-old infants either a surprising event (e.g., a ball passes through a wall) or a perceptually-matched expected event (e.g., the ball is stopped by the wall). Immediately afterward we taught infants new information about the object – that it had a hidden auditory property that they could not have known beforehand (e.g., it squeaked). When we later tested infants' learning of this new information, we found that infants had learned significantly more effectively following the surprising than the expected event. Control experiments confirmed that this learning enhancement was not due to longer perceptual exposure to objects in the surprising events, nor to a general preference for objects that violated expectations (Stahl & Feigenson, 2015). Thus, surprising events not only promote exploration, but also can change learning.

However, the scope of this surprised-induced enhancement of learning remains unknown. One important question is whether violations of core knowledge expectations boost learning only in preverbal infant learners, or only within the context of a laboratory looking-time task. For example, older children might be less reliant on core knowledge, having acquired a wider range of expectations through parents, teachers, and their own experiences. As such, subtle differences between possible and impossible events might not suffice to visibly affect their learning. In addition, it is possible that surprise measurably enhances learning only under rigorously controlled conditions, where a learner's attention is entirely focused on the event in question because there is little else to compete for their interest. Alternatively, surprise may benefit different types of learning across the lifespan, and in diverse environments. Here we tested these possibilities by asking whether surprise enhances older children's ability to learn, in the naturalistic learning setting of a busy science museum.

To test the effect of violations of expectation on learning, we chose a task that would be interactive and naturally engaging to children, but would also present a relatively difficult learning challenge. To this end, we asked whether surprise affects children's ability to learn a novel word from a single exposure. Young children are naturally engaged in learning new words, and can sometimes learn following a single stimulus presentation (e.g., Carey & Bartlett, 1978). Nonetheless, children vary in the trajectory of their vocabulary development (Fenson et al., 1994), and certain classes of words are more difficult for children to learn. Verbs in particular are more difficult to learn than nouns (e.g., Golinkoff, Jacquet, Hirsh-Pasek, & Nandakumar, 1996; Gentner, 1982; Bornstein et al., 2004), as the actions to which verbs refer are often ephemeral, and their meaning often relies on the syntactic structure of a sentence (Golinkoff & Hirsh-Pasek, 2008 for review). Children often require multiple exposures and contrastive examples to learn

novel verbs (e.g., Waxman, Lidz, Braun, & Lavin, 2009). Furthermore, distributed opportunities to learn a new verb aid children's learning (Childers & Tomasello, 2002); hence we anticipated that learning a new word from a single exposure would be challenging.

As a further test of the scope of surprise-induced learning, we tested children in a naturalistic learning context outside the laboratory. Children were taught new words within the dynamic setting of a busy science museum, with other museum-goers present and sights and sounds competing for children's attention. This allowed us to ask whether surprise shapes learning under natural learning conditions, as opposed to only in a tightly controlled lab environment.

In four experiments, we investigated whether 3- to 6-year-old children learn better from surprising events than from closely matched, unsurprising events. We taught children a series of novel words (verbs in Experiments 1 and 3, nouns in Experiments 2 and 4) following events that were novel but entirely possible, or events that were impossible in that they violated a principle of core object knowledge. We then measured children's learning of the new words. To preview, we found (in Experiments 1 and 2) that children learned significantly better following violations of expectation. In Experiment 3 we tested alternative explanations of our findings, and in Experiment 4 we tested whether surprise affects children's learning of any subsequent information, or only of information relevant to the surprising event.

2. Experiment 1

In Experiment 1 we compared children's ability to learn a new verb following a surprising event versus an expected event. To make the task dynamic and challenging, children had the opportunity to learn several new words within the experimental session. All children saw

two triads of events, with each of the six events containing a novel action. Following each event, children were taught a verb for the new action. The key feature of our experimental design was that the last event in each triad, the critical event, was either expected or surprising on the basis of core knowledge of object behavior. That is, it culminated in either an Expected outcome or in a Violation of spatiotemporal or featural continuity, violations to which even young infants are sensitive (e.g., Spelke et al., 1995; Wilcox et al., 1996; Wilcox, 1999; Wilcox & Baillargeon, 1998). The expected and surprising events were carefully constructed to be otherwise identical. Immediately following each of the two critical events, we tested children's identification of the novel verb's referent.

2.1 Method

2.1.1 Participants. Rather than targeting children in one narrow age window, we tested any children who were visiting the children's wing of the science museum and also were old enough to engage in our verbally-challenging task. This resulted in a sample of 38 children between 3 and 6 years old (range = 3 years, 1 month, 5 days – 6 years, 10 months, 6 days; mean = 4 years, 8 months, 3 days; 19 females). Four additional children were excluded for refusing to play the game (1) or experimenter error (3). All children in Experiments 1-4 were visitors to the Maryland Science Center and were invited to participate by a second experimenter who greeted families on the museum floor. Parents of all children provided written informed consent prior to children's participation.

2.1.2 Stimuli. Testing occurred in a corner of the museum exhibit hall. The testing area was separated from the rest of the exhibit hall by a waist-high wall; hence children could see and hear activity in the museum, and museum visitors could watch the testing session. Children sat in a small chair in front of a table with the experimenter sitting across from them. A small stuffed

bunny rabbit sat on the table facing the child, who was introduced as “Miss Bunny who has some toys and friends.”

2.1.2.1 Stimuli for Triad Ending in a Spatiotemporal Continuity Event. A black wooden stage (46 x 15 x 30 cm) sat on the table. The stage had two concealed circular trap doors that the experimenter could open and close with hidden knobs on the rear of the stage. When the knobs were turned, the trap doors allowed objects placed on top of them to be flipped out of view underneath the stage surface and, conversely, objects hidden underneath the stage (that were attached magnetically to the trap door’s surface) flipped up to become visible on the stage surface.

A green clay disc with a cross-hatched texture (4 cm diameter, 2 cm high) served as the target object throughout the Spatiotemporal Continuity triad of events. Three visually distinctive pairs of opaque plastic cups were used to perform the target actions: a pair of tall, thin red cups (9 cm diameter, 15.5 cm high) with vertical stripes made of red ribbon, a pair of shiny orange cups (11.5 cm diameter, 7 cm high) with beveled edges, and a pair of wide blue bell-shaped cups (14.5 cm diameter, 9 cm high) with a blue checkerboard pattern.

2.1.2.2 Stimuli for Triad Ending in a Featural Continuity Event. A green felt doll (7 x 13.5 cm) served as the target object throughout the Featural Continuity triad, and four other felt dolls (red, brown, and two orange) were used just during the critical third event. Three visually distinctive containers were used to perform the target actions: a woven yellow basket (12 x 9 x 14 cm) with a clear plastic container inside (6 x 7 x 6 cm); a translucent blue box (17 x 12 x 3 cm) containing a Velcro sticker; and a red velvet bag (12 cm diameter, 18.5 cm long) with a wooden handle. The bag was manufactured to be used in professional magic tricks. Its main pouch could be accessed through the top opening or through a visible zippered opening at the

bottom, but it also contained a secret hidden compartment that could only be accessed through the top opening when the experimenter surreptitiously moved a lever hidden underneath the handle. Moving the lever changed which compartment was accessible, so that an object placed in the bag could appear to vanish, or an empty bag could appear to suddenly contain something (both violations of continuity).

2.1.3 Procedure. All children saw two triads of events: one culminated in an Expected outcome and the other culminated in a Violation outcome. This within-subjects design allowed us to ask whether for a given child, word learning differed for surprising relative to expected events. The first two events in each triad were filler events that were always physically possible, but that were designed to be salient, novel, and to involve an action for which children did not have any existing verbal label. Children were taught a novel word after each of these possible events. The last event in each triad, the critical event, was either physically possible (i.e., ended in an Expected outcome) or physically impossible (i.e., ended in a Violation outcome); children were also taught a novel word following this critical event. Immediately after the third, critical event we tested children's learning of the novel verb from this third event. Half of the children saw the Spatiotemporal Continuity event culminate in a Violation outcome and the Featural Continuity event culminate in an Expected outcome, and the other half saw the reverse. Triad order was counterbalanced across children.

2.1.3.1 Triad Ending in a Spatiotemporal Continuity Event. The experimenter first told children that they were about to see some of Miss Bunny's toys. She retrieved the green toy and said, "This is one of Miss Bunny's toys!"

2.1.3.1.1 Filler Event 1. The experimenter said, "I'm going to put this toy right here," and placed the green toy on the left side of the stage. She then held up the two red cups, one in each

hand, so that children could see inside them. The left cup visibly contained a green cloth. The experimenter said, “See these? Do you see what’s in there? Watch this!” and lowered the cups facedown onto the stage, the left one covering the toy (and the right one covering an empty area of the stage). She then lifted the cups to reveal that the toy had been covered by the cloth, and then said, “Look what happened!” The experimenter pointed to the cups and said, “These *daxed* the toy!” She then cleared everything from the stage. The covering of a toy with a cloth was designed to be possible, but novel.

2.1.3.1.2 Filler Event 2. The experimenter placed the same green toy on the right side of the stage. She held up the two orange cups, one in each hand, so that children could see inside of them. The right cup visibly contained colored confetti. The experimenter said, “See these? Do you see what’s in there? Watch this!” and lowered the cups facedown onto the stage, the right one covering the toy (and the left one covering an empty area of the stage). She then lifted the cups to reveal that the toy had been covered with the confetti and said, “Look what happened!” The experimenter then pointed to the cups and said, “These *fepped* the toy!” She then cleared everything from the stage. The covering of a toy with confetti was designed to be possible, but novel.

2.1.3.1.3 Event 3 (Critical Event). The experimenter placed the same green toy on the left side of the stage. She held up the two blue cups, one in each hand, so that children could see inside and observe that both were empty. She said, “See these? Watch this!” and lowered the cups facedown onto the stage, the left one covering the toy (and the right one covering an empty area of the stage).

In the Expected outcome, the experimenter held both cups and, without lifting them, swapped their locations (as in a shell game) so that the cup that covered the toy was now on the

right. The experimenter momentarily removed her hands from the cups, then lifted the cups to reveal that the toy was now on the right side of the stage. This outcome was considered expected, because children had just seen the cups swap locations (Figure 1). As with the two filler events that had just occurred, this event was designed to be possible but novel.

In the Violation outcome, the experimenter left the cups in their original resting locations and momentarily removed her hands from the cups. As she did this, she secretly turned the knobs behind the stage to activate the trap doors. This made the object under the cup in the left-hand location flip out of view underneath the stage, and made the trap door in the right-hand location flip upward so that a duplicate toy rested on the stage under the right-hand cup. The experimenter then lifted the cups to reveal the toy on the right side of the stage, so that it appeared to have vanished from under the left cup and reappeared under the right cup (thus violating spatiotemporal continuity) (Figure 1). Critically, the final outcomes of the Expected and Violation events were perceptually identical (in both cases the toy was revealed under the cup on the right), but was expected in one case (because there was a physical explanation for the toy's new location) and surprising in the other (because there was no such explanation available). Immediately following either the Expected or Violation outcome, the experimenter said, "Look what happened!" then pointed to the cups and said, "These *blicked* the toy!" She then cleared everything from the stage.

2.1.3.1.4 Test trial. Immediately after this third critical event the experimenter brought out a tray holding one cup from each of the three preceding events (red, orange, blue), with the cups' spatial positions counterbalanced across children. She then probed children's learning of the verb that had been taught after the third, critical event by saying, "If I wanted to *blick* the toy again, which one should I use? Which one will *blick* the toy?" (Figure 1). Children were

prompted to point to or touch the object that could be used to *blick* the toy. If children did not respond immediately, the experimenter repeated the question (e.g., “Which one will *blick* the toy?”) until children answered. Across all experiments, children’s responses were coded as correct if they pointed to or touched the correct object; responses were coded as incorrect if they pointed to or touched either the wrong object, multiple objects (despite the experimenter’s encouragement to pick just one), or no objects at all.

Figure 1. Event sequence and test trial for the critical Spatiotemporal Continuity event. Children saw a toy hidden under the left cup. The experimenter then either swapped the cups' left-right locations and revealed the toy under the cup that was now on the right (Expected outcome), or she left the cups in place yet revealed the toy under the cup on the right (Violation outcome). She then taught children a novel verb to describe the action. In the test trial, children were presented with the cup from the critical event and the cups from the preceding filler events and were probed on the novel verb from the critical event.

2.1.3.2 Triad Ending in a Featural Continuity Event. The experimenter first told children that they were about to see some of Miss Bunny's friends. She held up the green felt doll and said, "This is one of Miss Bunny's friends!"

2.1.3.2.1 Filler Event 1. The experimenter showed children the yellow basket with the transparent container inside. She held the doll above the basket and said, "Watch this, I'm going to put him inside!" She shook the basket as she squeezed the doll so that it fit into the transparent container inside. The experimenter then removed the transparent container to reveal the doll tightly compressed inside and said, "Look what happened!" She wiggled the basket and said, "This *tizzed* him!" She cleared everything from the table. The squeezing of a toy into a tight container was designed to be possible, but novel.

2.1.3.2.2 Filler Event 2. The experimenter showed children the blue box with the Velcro sticker inside. She held the doll above the box and said, "Watch this, I'm going to put him inside!" She placed the doll in the box, closed the box and shook it, then removed the doll from the box, now with the Velcro sticker affixed to its front. She held up the doll and said, "Look what happened!" The experimenter wiggled the blue box and said, "This *zavved* him!" She cleared everything from the table. The sticker becoming stuck on the doll was designed to be possible, but novel.

2.1.3.2.3 Event 3 (Critical Event). The experimenter brought out the red magician's bag. In the Expected outcome, the experimenter first showed children that the bag was empty by unzipping its bottom opening, sticking her arm through it, and looking through one side while children looked through the other. Once children agreed that the bag was empty, the experimenter zipped the bag closed. She then showed them three more dolls (red, orange, and brown) introduced as Miss Bunny's other friends, and placed them in the bag, saying "Watch

this, I'm going to put these inside!" She then said, "I'm going to take just these two out" as she removed the red and brown dolls, showing them to children before placing them out of sight. She then said, "But there is still another one inside! Can you feel it?" She prompted children to feel the bulge in the bag made by the remaining orange doll. She then held up the original green doll and told children that she was going to put that doll inside the bag. She dropped the doll in and shook the bag. Finally, the experimenter reached in and pulled the orange doll out. This outcome was considered expected, because children had previously seen the orange doll hidden in the bag, then saw it removed (Figure 2).

For children who saw the Violation outcome, one orange doll was secretly pre-hidden in the bag's secret compartment. The experimenter started by showing children the bag and the three other felt dolls (red, orange, and brown). Children watched as she placed them inside the bag, saying "Watch this, I'm going to put these inside!" She prompted children to feel the bulge in the bag made by the dolls (to match children's experience with that in the Expected outcome). She said, "Now I'm going to take all of them out," removed all three dolls, and showed them to children before placing them out of sight. The experimenter then demonstrated that the bag was now empty, exactly as in the Expected outcome (by unzipping the bottom and prompting children to look through the apparently empty bag – note that there was still a duplicate orange doll inside, but surreptitiously hidden in the secret compartment), then zipped the bag closed. She held up the original green doll and told children that she was going to put it in the bag, dropped it inside and shook the bag. As she did this, she flipped the hidden lever on the bag's handle so that the secret compartment containing the hidden orange doll became accessible. The experimenter reached into the bag and pulled the orange doll out, so that it appeared that the green doll had turned into the orange doll (a violation of featural continuity) (Figure 2). As with

the Spatiotemporal Continuity event, the final outcome of the Expected and Violation events was identical (in both cases an orange doll was removed from the bag), but was expected in one case (because there was a physical explanation for the orange doll being inside) and surprising in the other (because there was no such explanation available). Immediately following either the Expected or Violation outcome, the experimenter said, “Look what happened!” then wiggled the bag and said, “This *moxed* him!”

2.1.3.2.4 Test trial. Immediately following this third critical event, the experimenter brought out a tray holding the container from each of the three preceding events (yellow basket, blue box, red bag), with the spatial positions of the containers counterbalanced across children. She then probed children’s learning of the verb that had been taught after the third, critical event by saying, “If I wanted to *mox* him again, which one should I use? Which one will *mox* him?” (Figure 2).

Figure 2. Event sequence and test trial for the critical Featural Continuity event. Children either saw a bag that already contained a doll inside (such that the hidden doll made a visible bulge in the bag), then saw a green doll placed into the bag and an orange doll removed (Expected outcome), or saw an empty bag and then saw the green doll placed inside and an orange doll removed (Violation outcome). The experimenter then taught children a novel verb to describe the action. In the test trial, children were presented with the bag from the critical event and the containers from the preceding filler events and were probed on the novel verb from the critical event.

2.2 Results

The goal of Experiment 1 was not to compare children's learning of the verb taught after the critical event to learning of the verbs taught after the filler events (but see Experiment 3). Rather, our aim was to compare learning following two nearly identical versions of a single event: learning following an event that ended in an Expected outcome, versus learning following the same event that ended in a Violation outcome.

We compared the proportion of children's correct responses on the test trial to chance (.33, as there were three objects to choose from) using a two-tailed binomial test. We found that children who had been taught a verb for an action that culminated in an Expected outcome chose at chance for both the Spatiotemporal Continuity event (6 out of 17 children, .35, $p = .51$) and the Featural Continuity event (7 out of 21 children, .33, $p = .57$). In contrast, children who had been taught the same verb for a nearly identical action – but one that culminated in a Violation outcome – exhibited significant learning. They chose the correct object at above chance levels following violations to Spatiotemporal Continuity (12 out of 21 children, .57, $p = .02$) and Featural Continuity (14 out of 17 children, .82, $p < .001$) (Figure 3).

We directly compared children's verb learning following surprising and expected events using a McNemar test (two-tailed, uncorrected). More children chose correctly following a Violation outcome and chose incorrectly following an Expected outcome than the reverse (14 vs. 1), $X^2(1) = 11.27$, $p < .001$.

Given the relatively wide age range in our sample, we asked whether younger children performed any differently than older children. Using a median split by age, a two-tailed Fisher's exact test revealed no significant difference in the performance of younger versus older children for either Expected events ($p = .17$) or Violation events ($p = 1.0$). We found that boys

outperformed girls on Expected events ($p = .04$), but found no gender differences for Violation events ($p = .73$).

Figure 3. Proportion of children in Experiment 1 who selected the correct object when tested on the novel verb from the critical event. Dashed line represents chance performance.

3. Experiment 2

We had anticipated that children in Experiment 1 would fail to learn novel verbs following expected events – the task was hard by design, involving exposure to six new words over the course of the testing session, with just a single exposure to each. As expected, children were at chance when tested on their knowledge of the new words – at least when the words had been taught immediately following novel but possible events. However, we observed markedly different performance when children were taught the same verb following a novel event with a surprising outcome. When children's expectations about physical object behavior were violated (i.e., when objects appeared to teleport from one location to another, or appeared to magically change their properties), children successfully learned the novel verbs that described the objects' behavior.

Given the striking nature of these findings, we next sought to replicate the pattern observed in Experiment 1, this time in the context of a slightly different word-learning task. In Experiment 2, we taught a new group of children nouns for the objects used in the events, rather than verbs for the actions themselves. Given that nouns are acquired developmentally earlier and are often easier to learn than verbs (e.g., Bornstein et al., 2004; Fenson et al., 1994; Gentner, 1982; Maguire, Hirsh-Pasek, & Golinkoff, 2006; Imai et al., 2008), this allowed us to ask whether children would still experience enhanced learning following surprising events, even within a somewhat less challenging task.

3.1 Method

3.1.1 Participants. Fifty-four children between 3 and 6 years old participated (range = 3 years, 7 days – 6 years, 10 months, 10 days; mean = 5 years, 24 days; 27 females). Sixteen additional children were excluded for refusing to play the game (5), experimenter error (1),

inattentiveness (6), parental/sibling interference (2), watching another child participate prior to testing (1), or missing birth date information (1).

3.1.2 Stimuli. The stimuli were identical to those in Experiment 1.

3.1.3 Procedure. The procedure was identical to that of Experiment 1 with one exception. Instead of labeling each of the three novel actions in each triad with a novel verb, the experimenter labeled each of the target objects (that had produced the novel action) with a novel noun. That is, instead of telling children that the cups in the Triad Ending in a Spatiotemporal Continuity Event had *daxed/fepped/blicked* the toy, she labeled the cups as *daxers/feppers/blickers*. For example, the critical third event in that triad involved the pair of blue cups. Just as in Experiment 1, children either saw the blue cups move the toy from one position to the other by tracing a spatiotemporally continuous path, as in a shell game (Expected outcome), or saw the blue cups magically “teleport” the toy from one location to the other (Violation outcome). Immediately following either the Expected or the Violation outcome, the experimenter pointed to the blue cups and said, “These are *blickers*!”

Similarly, instead of telling children that the containers in the Triad Ending in a Featural Continuity Event had *tizzed/zavved/moxed* the doll, the experimenter labeled each container as a *tizzer/zavver/moxer* (e.g., “This is a *moxer*!”). For example, the third critical event in that triad involved a red bag. Just as in Experiment 1, children either saw the experimenter place a green doll in the red bag and reach in to find an orange doll that children had seen placed inside earlier and had not seen removed (Expected outcome), or saw the experimenter reach into the red bag to find the orange doll that children had already seen removed (Violation outcome). Immediately following either the Expected outcome or the Violation outcome, the experimenter wiggled the bag and said, “This is a *moxer*!”

As in Experiment 1, children received one test trial after each triad of events. In the test trial the experimenter brought out a tray containing the three objects from the previous three events and asked children, “If I wanted to play with the *blicker/ moxer* again, which one should I use? Which one is the *blicker/ moxer*?”

3.2 Results

We again compared children’s learning following the critical Expected and Violation outcomes to chance. We found that when children were taught the novel noun following an event culminating in an Expected outcome, they showed no evidence of having learned the new word; they chose at chance for both the Spatiotemporal Continuity event (10 out of 27 children, .37, $p = .40$) and the Featural Continuity event (10 out of 27 children, .37, $p = .40$). In contrast, when children were taught the same novel noun following a nearly identical event that culminated in a Violation outcome, they exhibited significant learning; children chose the correct object at above chance levels for both the Spatiotemporal Continuity event (16 out of 27 children, .59, $p = .004$) and the Featural Continuity event (14 out of 27 children, .52, $p = .03$) (Figure 4). As in Experiment 1, more children chose correctly following a Violation outcome and chose incorrectly following an Expected outcome than the reverse (18 vs. 8), $X^2(1) = 3.85$, $p = .049$ (two-tailed uncorrected McNemar’s test).

Lastly, a Fisher’s exact test showed equivalent performance for the younger and older children for both Expected events ($p = .78$) and Violation events ($p = .17$). We also found no gender differences for either Expected events ($p = .78$) or Violation events ($p = 1.0$).

Figure 4. Proportion of children in Experiment 2 who selected the correct object when tested on the novel noun from the critical event. Dashed line represents chance performance.

4. Experiment 3

Together, Experiments 1 and 2 offer evidence that children learn words more effectively following surprising than expected events. However, there are at least two alternative explanations of our findings. The first is that children only appeared to succeed following Violation events because they were attracted to objects that behaved in impossible or surprising ways, not because they had learned a new word for those objects or for the action the objects produced. That is, if children in Experiments 1 and 2 detected the violations of spatiotemporal and featural continuity (as the results of those experiments suggest they did), they might have pointed to objects that had committed violations simply because they wanted to further explore objects that behaved in surprising ways. Indeed, in previous work we found that preverbal infants preferentially engaged with a toy that had just been seen to violate expectations about object behavior, compared to a novel toy (Stahl & Feigenson, 2015). On this alternative interpretation, the results of Experiments 1 and 2 show that children treat objects that violate expectations differently from objects that accord with expectations, but do not provide evidence of an effect of surprise on actual learning.

The second alternative explanation for our findings is that rather than learning being enhanced following the surprising event, learning was *impaired* following the expected event. On this account, the expected versions of the critical spatiotemporal and featural continuity events did not provide a sufficiently distinctive action onto which children could map the novel label. Whereas the filler events visibly transformed the object in some way (in one triad a cup covered a toy with a cloth, and another cup covered a toy with confetti; in the other triad a container squeezed a doll tightly, and another container pressed a sticker onto the doll), and the violation events also transformed the object in some way (making it teleport to a new location, or

causing it to change its features), the expected events may have been less distinctive – they merely involved an object moving to a new location (while tracing a spatiotemporally continuous path) or being placed into, swished inside of, and then removed from a bag. If there was ambiguity regarding the referent of the novel word following these Expected outcomes, then children may have shown poorer word learning about these events (whereas learning following the Violation outcome reflected baseline levels of successful learning, comparable to learning about any distinctive, but possible, event).

Experiment 3 tested these alternative accounts. Children saw the same series of events and were taught the same novel verbs as in Experiment 1. But this time, instead of testing children's learning of the verb that had been taught following the third, critical event, we tested children's learning of a verb taught after one of the filler events (which, as in Experiment 1, was always a novel but physically possible transformation). If children's performance in Experiments 1 and 2 reflected attraction to surprising objects rather than actual learning, then children in Experiment 3 should erroneously choose the object that had participated in a violation event (i.e., should choose similarly to children in Experiment 1). If children in Experiments 1 and 2 failed to identify the referent of the word taught following an Expected outcome because that word did not describe an identifiable action, then children should answer correctly when probed about a word used to describe the filler event (because the filler events always involved a clearly identifiable novel action).

4.1 Method

4.1.1 Participants. Thirty-four children between 3 and 6 years old participated (range = 3 years, 28 days – 6 years, 10 months, 21 days; mean = 5 years, 2 months, 5 days; 11 females).

Seven additional children were excluded for refusal to play (2), experimenter error (3), watching another child participate prior to testing (1), or missing birth date information (1).

4.1.2 Stimuli. The stimuli were identical to those in Experiment 1.

4.1.3 Procedure. The procedure was identical to that of Experiment 1, with one exception: in each of the two test trials, the experimenter probed children's learning of the novel verb that had been taught following the second filler event (which was always physically possible). In the Triad Ending in a Spatiotemporal Continuity Event, this corresponded to the event in which an orange cup covered a toy with confetti. As in Experiment 1, after the confetti-covered object was revealed the experimenter said, "Look what happened!" The experimenter then pointed to the cups and said, "These *fep*ped the toy!" The covering of a toy with confetti was designed to be possible, but novel. After this, children saw either the Expected or Violation outcome of the Spatiotemporal Continuity event (i.e., they saw a toy hidden under one blue cup swap locations with another blue cup, as in a shell game, or they saw a toy hidden under one blue cup appear to teleport to a new location), and the experimenter said, "Look what happened! These *blicked* the toy!" Immediately following the third event, the experimenter brought out the tray containing one cup from each of the three preceding events (red, orange, and blue, with spatial location counterbalanced across children) and said, "If I wanted to *fep* the toy again, which one should I use? Which one will *fep* the toy?"

As in Experiment 1, in the Triad Ending in a Featural Continuity Event, the second filler event involved a doll placed in a blue box and becoming stuck to a Velcro sticker inside. After the doll with the sticker was revealed the experimenter said, "Look what happened!" She then wiggled the blue box and said, "This *zavved* him!" The adhering of the Velcro sticker to the doll inside the box was designed to be possible, but novel. After this, children saw either the

Expected or Violation outcome of the Featural Continuity event (i.e., they saw the experimenter drop a green doll into the red bag and then retrieve an orange doll that had previously been hidden and was never seen removed, or they saw the experimenter reach into the bag to retrieve the orange doll that had previously been hidden but was seen already removed), and the experimenter said, “Look what happened! This *moxed* him!” Immediately following the third event, the experimenter brought out the tray containing the container from each of the three preceding events (yellow basket, blue box, red bag, with spatial location counterbalanced) and said, “If I wanted to *zav* him again, which one should I use? Which one will *zav* him?” As in Experiment 1, each child saw one Expected outcome and one Violation outcome; which of these occurred in the Triad Ending in a Spatiotemporal Continuity Event and which in the Triad Ending in a Featural Continuity Event was counterbalanced across children.

4.2 Results

The first possibility we tested is that children were simply drawn to an object that had violated their expectations. If so, then in the test trial they should choose the third object from the critical Violation outcome at above chance levels, even though they had been asked to indicate the object from the second filler event. In contrast to this prediction, we found that children who were asked for the object from the second filler event did not select the target object from the third, critical event at above chance levels following either a Violation to Spatiotemporal Continuity (4 out 18 children, .22, $p = .89$) or Featural Continuity (8 out of 16 children, .50, $p = .12$). Crucially, this pattern did not differ from that observed after children had seen the Expected version of the critical event – children did not select the target object from the third event following an Expected outcome to Spatiotemporal Continuity (5 out of 16 children, .31, $p = .65$) or Featural Continuity (7 out of 18 children, .39, $p = .38$) (Figure 5A). Equal numbers of children

chose the third target object following a Violation outcome but not following an Expected outcome as chose the third object following an Expected outcome but not following a Violation outcome (7 vs. 7), $X^2(1) = 0$, $p = 1.0$ (two-tailed uncorrected McNemar's test). Fisher's exact tests revealed that younger and older children were equally likely to erroneously select the target object following Expected ($p = .72$) and Violation events ($p = .72$), as were boys and girls for Expected ($p = .46$) and Violation events ($p = .46$).

Next we tested the possibility that rather than children's learning being enhanced following our Violation outcomes, it was impaired following our Expected outcomes. If so, then children should have exhibited successful learning following the (non-surprising) filler events, because these involved clearly novel actions that transformed the objects. In contrast to this prediction, we found that children were at chance at identifying the object labeled in the filler event. This was true for children who had experienced the filler event followed by the Expected outcome of a Spatiotemporal Continuity event (5 out of 16 children, .31, $p = .65$) and for those who had experienced the filler event followed by the Expected outcome of a Featural Continuity event (9 out of 18, .50, $p = .10$). Likewise, children were at chance at identifying the object labeled in the filler event when that filler event was followed by the Violation outcome of a Spatiotemporal Continuity event (7 out of 18 children, .39, $p = .38$), and the Violation outcome of a Featural Continuity event (5 out of 16 children, .31, $p = .65$) (Figure 5B). There was no difference in the number of children who answered correctly following a Violation outcome and answered incorrectly following an Expected outcome than the reverse (6 vs. 8), $X^2(1) = .29$, $p = .59$ (two-tailed uncorrected McNemar's test). As in the previous experiments, Fisher's exact tests showed that younger and older children performed equally for Expected ($p = .73$) and Violation events ($p = .72$), as did boys and girls, for Expected ($p = .29$) and Violation events ($p = .70$).

Finally, a Fisher's exact test confirmed that children from Experiment 1 who were asked to identify the referent of the verb taught after the Expected outcome of the third, critical event performed no differently than children in Experiment 3 who were asked to identify the referent of the verb taught after the second (filler) event (in Experiment 1: 13 out 38 children, .34; in Experiment 3: 14 out of 34 children, .41; $p = .63$). Thus, children's inability to learn the word taught after an Expected event in Experiments 1 and 2 was not due to the particular event we used.

Figure 5. Proportion of children in Experiment 3 who A. incorrectly selected the target object from the third, critical event when tested on the novel verb from the second, possible filler event; B. correctly identified the referent of the novel verb from the second, possible filler event. Dashed line represents chance performance.

5. Experiment 4

The results of Experiment 3 rule out the possibilities that children's choices in Experiments 1 and 2 merely reflected a preference to engage with objects that had behaved in surprising ways, and that learning following Expected events was impaired due to referent ambiguity. Taken together, Experiments 1-3 offer evidence that children's learning of novel words is enhanced following surprising events.

How focused is this effect of surprise on learning? It could be the case that when a learner's predictions are observed to be wrong, the learner hones in on the opportunity to revise their knowledge about the specific objects or events involved in the prediction. Alternatively, seeing a violation might alter subsequent learning in a more diffuse way, by increasing overall levels of attention or arousal (Nielson, Yee, & Erickson, 2005; Schwabe, Bohringer, Chatterjee, & Schachinger, 2008). In previous work with infants, we tested the specificity of the surprise-induced learning enhancement. We found that when infants saw an object violate core expectations of object behavior (e.g., by appearing to teleport from one location to another), they showed enhanced learning about that object, compared to infants who had seen the same object accord with expectations. But when instead infants who saw the same surprising event were subsequently taught about a new object – one that had not participated in the central event – they showed no learning benefit (Stahl & Feigenson, 2015). This suggests that, for infants at least, surprise exerts a targeted effect on learning, only boosting learning about entities relevant to the surprising event.

In Experiment 4 we asked whether surprise also enhances learning in a targeted way in young children. Following the design of Experiments 1-3, we showed children triads of events that either culminated in a Violation outcome or an Expected outcome. In each event we

included a novel “innocent bystander” object that was present during each event but did not participate in any actions. Immediately following each event we labeled the bystander object with a novel noun, and then asked whether children were better at learning the novel noun following a violation event than following an expected event. If surprise enhances learning in a diffuse way (e.g., by increasing attention or arousal), children should learn novel words for the bystander objects better following Violation events than Expected events.

5.1 Method

5.1.1 Participants. Thirty-eight children between 3 and 6 years old participated (range = 3 years, 1 month, 27 days – 6 years, 7 months, 14 days; mean = 4 years, 9 months, 6 days; 14 females). Eleven additional children were excluded for refusal to play (2), experimenter error (3), inattentiveness (4), or missing birth date information (2).

5.1.2 Stimuli. The same stimuli were used as in Experiment 1. Three additional unique novel objects were used in the Triad Ending in a Spatiotemporal Continuity Event – these were all differently shaped metal drain components, each measuring around 8 cm (Figure 6A). Three additional unique novel objects were used in the Triad Ending in a Featural Continuity Event – these were all plastic plugs, measuring around 15 cm, with different colors and patterns (Figure 6B). These stimulus objects were chosen to be entirely novel to children, and to be visually distinctive from one another, but to share some perceptual properties (as did the cups and containers that participated in the events and served as test objects in Experiments 1-3).

5.1.3 Procedure. The procedure was nearly identical to that of Experiment 2. As in our previous experiments, in the Triad Ending in a Spatiotemporal Continuity Event, the first two events were filler events in which one of a pair of red cups covered a green toy with a cloth, and then one of a pair of orange cups covered the green toy with confetti. These events were novel

but possible. In the third, critical trial, children either saw the green toy covered under one of two blue cups and then saw the cups' positions swapped (Expected outcome), or saw the green toy covered under one of two blue cups and then saw the cups remain in place and the toy “teleport” to the other location (Violation outcome). In the Triad Ending in a Featural Continuity Event, the first two events were filler events in which a green cloth doll was compressed into a tiny container, and then the green doll was closed inside a box and became stuck to a Velcro sticker. These events were novel but possible. In the third, critical trial children either saw a green doll and an orange doll placed into a bag and then the orange doll removed (Expected outcome), or saw an orange doll removed from a bag, and then saw a green doll placed into the bag but saw an orange doll removed (Violation outcome).

Critically, throughout each of the three events in each triad, a unique novel object was present. In the Triad Ending in a Spatiotemporal Continuity Event, the novel object rested in the center of the stage, behind the other objects being shown to children but in a position where it was very clearly visible. In the Triad Ending in a Featural Continuity Event, the object rested on the table, next to the other objects being shown to children – again in a visible and salient position. Although the novel objects were always clearly visible to children, they never participated in any of the events (i.e., they were “innocent bystanders”). Immediately after each event ended (e.g., the green toy was revealed under the right cup), the experimenter pointed to and labeled the bystander object with a novel noun (using the same nouns from Experiment 2; e.g., “This is a *blicker!*”) rather than labeling the object(s) that had participated in the event. Hence the temporal dynamics of the presentation were identical to those in Experiments 1 and 2. Just as in Experiments 1-3, each child saw one Expected outcome and one Violation outcome; which of these occurred following a Spatiotemporal Continuity event versus a Featural

Continuity event was counterbalanced across children. As in Experiments 1-3, children heard each object labeled just once.

At test, immediately following the last event in each triad, children saw the three novel objects presented on a tray (with spatial location counterbalanced) and were asked for the word that had been taught following the critical Expected or Violation outcome (i.e., “If I wanted to play with the *blicker/ moxer* again, which one should I use? Which one is the *blicker/ moxer*?”).

Figure 6. The novel “innocent bystander” objects from Experiment 4, as they appeared in the two test trials. A. Objects used in the Triad Ending in a Spatiotemporal Continuity Event. B. Objects used in the Triad Ending in a Featural Continuity Event.

5.2 Results

We found that, as expected on the basis of the difficulty of the task and the results of Experiments 1-3, children did not learn the word for the bystander object following Expected outcomes to the Spatiotemporal Continuity event (9 out of 18 children, .50, $p = .10$) or Featural Continuity event (7 out of 20 children, .35, $p = .51$). Notably, children also failed to learn the word for the bystander objects following Violations to the Spatiotemporal Continuity event (10 out of 20 children, .50, $p = .09$) and the Featural Continuity event (8 out of 18 children, .44, $p = .21$) (Figure 7). Most critically, children were equally likely to answer correctly when the bystander object had been labeled following a Violation outcome and incorrectly following an Expected outcome as the reverse (7 vs. 5), $X^2(1) = .33$, $p = .56$ (two-tailed uncorrected McNemar's test). Finally, Fisher's exact tests again found no difference in performance between younger and older children for either Expected events ($p = .74$) or Violation events ($p = 1.0$). We also found no gender differences for Expected ($p = .31$) or Violation events ($p = .33$).

The results of Experiment 4 therefore suggest that surprise enhances children's learning in a targeted way. Children reaped no benefit for learning about objects that were unrelated to a surprising event, even when taught about these bystander objects immediately following that surprising event.

Figure 7. Proportion of children in Experiment 4 who selected the correct “innocent bystander” object when tested on the novel noun from the critical event. Dashed line represents chance performance.

6. Discussion

The current experiments investigated the effect of surprise on learning in 3- to 6-year-old children. We measured learning using a challenging fast mapping task, administered in a busy science museum. In Experiment 1 we showed children events that either accorded with or violated core knowledge expectations about object behavior, and immediately afterwards taught them novel words. Importantly, we ensured that the outcomes of the events were perceptually identical. In Experiment 1 we found that children learned novel verbs significantly more effectively when taught following the surprising than the expected events. This learning enhancement occurred despite children having had just a single exposure to the novel verb, and despite verbs being notoriously difficult to learn (e.g., Golinkoff & Hirsh-Pasek, 2008). Experiment 2 replicated these results, this time using nouns that labeled the objects themselves rather than the objects' actions. Even though nouns are comparatively easier for children to learn than verbs (e.g., Imai et al., 2008), children still showed enhanced learning for nouns that labeled objects that violated their expectations, compared to learning of nouns that labeled objects that accorded with expectations.

Experiment 3 ruled out two alternative explanations for these findings. The first was that children selected the correct object following surprising outcomes simply because they were drawn to objects that had violated expectations. The second was that the difference in children's learning following surprising versus expected events came not from the surprising test events having enhanced learning, but rather from the expected test events having impaired learning. In contrast to the predictions of both of these alternative accounts, we found that when we tested children's learning of a word that labeled a novel but unsurprising (i.e., physically possible) action from earlier in the stimulus presentation, children did not systematically choose the object

that had just violated their expectations, and also were at chance in identifying the word's referent. This suggests that children were not simply drawn to interact with objects that behaved in surprising ways, and also that they found our word learning task difficult, even when shown salient novel events – in other words, learning was not specifically impaired by the expected version of our test event. Finally, in Experiment 4 we found that children showed no learning enhancement for objects that were peripheral to the violation event. Even following surprising events, children failed to learn the novel nouns that described these irrelevant “innocent bystander” objects.

Many decades of research have capitalized on infants' longer looking at surprising events to characterize the nature of early knowledge (e.g., Baillargeon et al., 1985; Kuhlmeier et al., 2003; McCrink & Wynn, 2004; Spelke et al., 1992). Other evidence from physiological responses (e.g., Berger et al., 2006; Emberson et al., 2015; Gredebäck & Melinder, 2010; Kouider et al., 2015), social responses (Walden et al., 2007), and emotional responses (Camras et al., 2002) also converge to show that infants detect discrepancies between expectations and observations. Our current findings suggest that beyond merely detecting surprising events, young learners also experience cognitive consequences to seeing something fail to accord with expectations. Our prior work with preverbal infants showed that core knowledge shapes learning – infants selectively explored and tested hypotheses about objects that behaved surprisingly, and more effectively acquired new information about the features of objects that behaved in surprising ways (Stahl & Feigenson, 2015). The present results suggest that this surprise-induced enhancement of learning extends past infancy, and occurs outside of the controlled laboratory setting. Taken together, the experiments with infants and with older children show that expectancy violations enhance learning across a range of ages (infants and children),

environments (a controlled laboratory and a busy science museum), types of surprising events (violations of expectations regarding solidity, support, spatiotemporal continuity, and featural continuity), learning measures (looking, exploration, and pointing), and learning content (object properties and novel words).

Although we find that a range of surprising events promote enhanced learning, it remains open what kinds of learning are enhanced following surprising events. As reported in our prior work (Stahl & Feigenson, 2015) and in the current experiments, domain-specific knowledge (i.e., core knowledge of objects) enhanced learning about objects – both objects’ auditory properties (Stahl & Feigenson, 2015) and objects’ verbal labels (the present experiments). Why would surprise aid learning about these features? Although we have not yet tested the boundaries of the surprise-induced learning enhancement, we speculate that surprise will enhance learning of any nonobvious property of the surprising entity. An object’s hidden properties (like the sounds it makes) might provide cues about its internal structure, which could conceivably help explain its behavior. A verbal label for a surprising action or object might help identify and categorize similar actions or objects in the future, thereby leading to more accurate future predictions. Given our assumption that surprising events trigger a search for explanatory information, we predict that any information relevant to an object’s causal properties will be better acquired following surprise. These predictions align with prior findings that children are more likely to generate targeted causal explanations when their expectations are violated (e.g., e.g., Bonawitz et al., 2012; Legare, 2012; Legare, 2014 for review; Legare et al., 2010). Seeking and generating explanatory information might also benefit formal learning as learners attempt to accommodate unexpected information into their theories (Legare, 2014; Valdesolo, Shtulman, & Baron, in press).

Research with nonhuman animals (e.g., rats, monkeys) has long suggested that surprising events catalyze some types of learning (e.g., Friston, 2005; Hayden, Heilbronner, Pearson, & Platt, 2011; Hollerman & Schultz, 1998; Kamin, 1969; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Roesch, Esber, Li, Daw, & Schoenbaum, 2012 for review; Schultz, Dayan, & Montague, 1997; Waeleti, Dickinson, & Schultz, 2001). In these cases, violation of expectation is operationalized as prediction error – a quantifiable difference between what the animal expected to happen and what it observed. When prediction errors are large, the animal modifies its behavior to more successfully predict future events; when prediction errors are small or zero, little or no modification occurs (see Schultz & Dickinson, 2000 for review). The highly influential Pearce Hall model of learning (Pearce & Hall, 1980) offers an account of this effect: it posits that an unexpected unconditioned stimulus increases an animal's attention to, and thereby the associability of, a conditioned stimulus that accompanied that surprising event. Therefore, large prediction errors lead to faster learning.

The relationship between this literature on animal learning and the enhanced learning we observed in infants and children merits further investigation. However, we note that the prediction errors in studies with nonhuman animals arise as a result of associative expectations accrued throughout the experimental session. These errors vary continuously in strength depending on the reinforcement history the animal previously experienced. It is an open question whether violations of infants' and children's core knowledge function similarly, and whether they are subserved by brain circuitry analogous to that identified in non-human animals. On a core knowledge view (e.g., Spelke & Kinzler, 2007), core knowledge is not acquired gradually through hundreds of associative experiences. Rather, it is the product of a long evolutionary history, and may produce categorical rather than graded expectations. Further work is needed to

reconcile (or differentiate) findings on surprise's effect on learning across species and under various learning conditions.

Of course, at the heart of this issue is the question of what mechanism drives the effect of surprise on learning. As mentioned above, one possibility is that seeing a surprising event causes observers to allocate more attention to the particular entity that behaved surprisingly (e.g., Pearce & Hall, 2008). It is possible that children's enhanced learning about surprising objects and actions is in part due to having allocated more attention to the target objects following violation outcomes than expected outcomes. Although children had limited time to fixate on the events (as the experiment proceeded with the same pacing for each child, and stimuli were removed promptly from view after each event), we do not know precisely to what aspects of the scene children attended throughout. For instance, when children viewed the expected outcomes, they may have attended to less relevant aspects of the scene (e.g., the distance of the objects from themselves; the color of the stage apparatus). In contrast, when children saw the violation outcomes, they may have attended more to the objects themselves, leading to better memory and learning about those objects. Critically, however, this potential boost in attention must be highly localized to only those objects that participated the violation event, as children in Experiment 4 failed to learn about objects that were present during but irrelevant to the surprising event. Another related possibility for how surprise affects children's learning concerns the depth with which they processed the violation events. In particular, children who saw an object behave surprisingly might have been impelled to think about the object more deeply – perhaps better encoding features including the objects' potential causal properties and verbal label.

It may also be informative to consider children's enhanced learning following surprising events from a Bayesian framework. On such a framework, infants and children are thought to

have priors about, for example, object behavior (e.g., that objects exist continuously); these priors might be set by evolution (i.e., are innate) or by experience gained in the first few months or years of life (i.e., are acquired). Children then observe evidence, which is weighed and used to update priors to posterior probabilities. A variety of kinds of information are relevant to this Bayesian process. For example, a learner can determine how expected the observation was given the strength of their prior hypotheses – on this view, surprise is viewed not as a binary dimension (with events being either surprising or expected), but as continuous (e.g., Perfors et al., 2011; Schulz, 2012; Téglás et al., 2011; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). According to Bayesian models of surprise (e.g., Baldi & Itti, 2010; Courville, Daw, & Touretzky, 2006), children should revise their posteriors to a greater degree upon receiving surprising evidence than receiving evidence that matches their expectations. This difference between the initial prior and the subsequent posterior signals a change and thus uncertainty in the prior, therefore leading to faster or better learning (which aligns with animal models of surprise-induced learning, such as Pearce and Hall, 1980). This framework is compatible with the aforementioned possibilities regarding increased attention to and/or deeper processing of surprising objects. It might be the difference between children's priors and posteriors that induces a change in attention (Baldi & Itti, 2010; Itti & Baldi, 2009) or depth of processing that, in turn, enhances learning following core knowledge violations.

One route for better characterizing how surprise shapes learning is to understand potential constraints on this phenomenon. In the present experiments we asked whether violations of core knowledge of object behavior affect children's learning. Violations of other types of knowledge may also drive learning. For example, our results open the door to asking whether violations outside the core domain of objects (e.g., violations of numerical or social knowledge) would also

enhance learning. In addition, it will be important to know whether, for infants and children, events that are statistically unlikely (as opposed to impossible, as in the events in our present experiments) also boost learning as impossible core violations do. Children discriminate between the improbable and the impossible in fictional stories (Weisberg & Sobel, 2012; but see Shtulman, 2009; Shtulman & Carey, 2007), and even young infants can use probabilistic information to make inferences about whether outcomes are statistically likely or unlikely (Téglás & Bonatti, 2016; Téglás et al., 2011; Téglás, Girotto, Gonzalez, & Bonatti, 2007; Xu & Denison, 2009; Xu & Garcia, 2008). Infants can also track auditory and visual statistical patterns (e.g., Saffran et al., 1996; Stahl et al., 2014) and are sensitive to the degree to which those patterns are surprising (Kidd, Piantadosi, & Aslin, 2012, 2014). It is not yet known whether the degree of children's surprise parametrically affects learning.

This issue – regarding the degree of surprise – highlights one striking aspect of our findings here and in our previous work (Stahl & Feigenson, 2015): we found that children and infants experienced better learning from *impossible* events – events that adults believe cannot happen in actuality. This might seem to call into question the utility of a mechanism that enhances learning about events that, in the natural world, can never occur. Yet although we showed children events designed to appear physically impossible (e.g., an object changing its features), these events might nonetheless be thought of as “minimally counterintuitive” (e.g., Barrett & Nyhof, 2001; Boyer, 1994; Boyer & Ramble, 2001). The outcomes of our violation events were always interpretable on the basis of children's initial parsing of the scene – for example, children saw a green doll hidden in a bag, but the doll was orange when it was retrieved. This event, though surprising, is still describable using the same conceptual vocabulary as one would use for the expected version of the event. In both cases the event involved the same

key entities (doll, bag) and actions (hiding, retrieving). The surprising aspect of the violation outcome was limited to just a single aspect of these (the doll's color changed after hiding). Hence children saw an "impossible" outcome, but one that was in fact tightly linked to the true nature of the objects and the event.

It may be that although children in our studies learned better from impossible events than possible ones, the degree of impossibility matters for learning. Previous research finds that children better remembered details for stories that contained up to two violations (e.g., a banana that felt angry and could turn invisible), compared to stories with no violations (e.g., a banana that was yellow and smelled good) or too many violations (e.g., a banana that felt angry, could turn invisible, and could live in outer space without oxygen) (Banerjee, Haque, & Spelke, 2013). The events we showed children contained a single violation of a core principle of object behavior. Events containing further violations, like an object changing its features, then splitting into dozens of other objects, *then* disappearing, might be too unpredictable to attend to and learn about. Similarly, surprise-induced enhancements of learning might require continuity of stimulus entities, such that learning *about these entities* can benefit. For example, an object that teleports from one location to another, or that changes its color while it is occluded, is still perceived as the self-same object ("How did *that object* get over there?" "How did *that object* change its color?"). When events in a scene prevent an observer from tracing the same entities over space and time, effective learning may be blocked. Future experiments should explore how these differences in the degree of surprise might affect learning. One plausible prediction is that optimal learning follows events that are neither too predictable, nor too surprising (e.g., containing so many violations that the event is hard to interpret) (Kidd et al., 2012; 2014).

In summary, several decades of research have shown that infants and children detect, generate explanations for, and preferentially explore objects involved in surprising events (e.g., Baillargeon et al., 1985; Berger et al., 2006; Bonawitz et al., 2012; Legare, 2012; Stahl & Feigenson, 2015). Here we offer new evidence that surprising events also can change children's learning. When children's basic expectations about object behavior were violated by impossible outcomes, children learned better than when their expectations were confirmed. This surprise-induced learning enhancement following violations of core knowledge may be one way in which the developing mind constrains the learning challenge it faces.

Acknowledgments:

This work was supported by the National Science Foundation (Graduate Research Fellowship DGE-1232825 to A.E.S, and NSF grant 1113648 to the Living Laboratory). We thank Justin Halberda for valuable discussion, Jessica Taggart, Bridget McGowan, and Julissa Veras for assistance with data collection and coding, and James Garmon for construction of the experimental apparatus. We also thank Miriam Krause, Stacey Prinzing, and the Maryland Science Center for their support, and Rebecca Kipling for enabling collaboration between the university and the museum.

References

- Baillargeon, R. (1986). Representing the existence and the location of hidden objects: Object permanence in 6- and 8-month-old infants. *Cognition*, 23, 21–41.
- Baillargeon, R. & Graber, M. (1987). Where's the rabbit? 5.5-month-old infants' representation of the height of a hidden object. *Cognitive Development*, 2, 375-392.
- Baillargeon, R. & Hanko-Summers, S. (1990). Is the top object adequately supported by the bottom object? Young infants' understanding of support relations. *Cognitive Development*, 5, 29-53.
- Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition*, 20, 191-208.
- Baldi, P. & Itti, I. (2010). Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Networks*, 23, 649-666.
- Baldwin, D. A., Markman, E. M., & Melartin, R. L. (1993). Infants' ability to draw inferences about nonobvious object properties: Evidence from exploratory play. *Child Development*, 64, 711-728.
- Banerjee, K., Haque, O. S., & Spelke, E. S. (2013). Melting lizards and crying mailboxes: Children's preferential recall of minimally counterintuitive concepts. *Cognitive Science*, 37, 1251-1289.
- Barrett, J. L., & Nyhof, M. A. (2001). Spreading non-natural concepts: The role of intuitive conceptual structures in memory and transmission of cultural materials. *Journal of Cognition and Culture*, 1, 69-100.
- Berger, A., Tzur, G., & Posner, M. I. (2006). Infant brains detect arithmetic errors. *Proceedings of the National Academy of Sciences*, 103, 12649-1653.

- Bonawitz, E. B., van Schijndel, T.J., Friel, D., & Schulz, L. (2012). Balancing theories and evidence in children's exploration, explanations, and learning. *Cognitive Psychology*, 64, 215-234.
- Bornstein, M. H., Cote, L. R., Maital, S., Painter, K., Park, S. Y., Pascual, L., Pecheux, M. G., Ruel, J., Venuti, P., & Vyt. A. (2004). Cross-linguistic analysis of vocabulary in young children: Spanish, Dutch, French, Hebrew, Italian, Korean, and American English. *Child Development*, 75, 1115-1139.
- Boyer, P. (1994). *The naturalness of religious ideas: A cognitive theory of religion*. Berkeley and Los Angeles, CA: University of California Press.
- Boyer, P., & Ramble, C. (2001). Cognitive templates of religious concepts: Cross-cultural evidence for recall of counter-intuitive representations. *Cognitive Science*, 25, 535-564.
- Camras, L. A., Meng, Z., Ujiie, T., Dharamsi, S., Miyake, K., Oster, H., Wang, L., Cruz, J., Murdoch, A., & Campos, J. (2002). Observing emotion in infants: Facial expression, body behavior, and rater judgments of responses to an expectancy-violating event. *Emotion*, 2, 179-193.
- Carey, S. & Bartlett, E. (1978). Acquiring a single new word. *Proceedings of the Stanford Child Language Conference*, 15, 17-29.
- Chandler, M. J. & Lalonde, C. E. (1994). Surprising, miraculous, and magical turns of events. *British Journal of Developmental Psychology*, 12, 83-95.
- Chiandetti, C. & Vallortigara, G. (2010). Experience and geometry: Controlled-rearing studies with chicks. *Animal Cognition*, 13, 463-470.

- Childers, J.B. & Tomasello, M. (2002). Two-year-olds learn novel nouns, verbs, and conventional actions from distributed or massed exposures. *Developmental Psychology*, 38(6), 967.
- Cook, C., Goodman, N., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, 120, 341-349.
- Courville, A. C., Daw, N. D., & Touretzky, D. S. (2006). Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences*, 10, 294-300.
- Dehaene, S., Izard, V., Pica, P., & Spelke, E. S. (2006). Core knowledge of geometry in an Amazonian indigene group. *Science*, 311, 381-384.
- Emberson, L. L., Richards, J. E., & Aslin, R. N. (2015). Top-down modulation in the infant brain: Learning-induced expectations rapidly affect the sensory cortex at 6 months. *Proceedings of the National Academy of Sciences*, 112, 9585-9590.
- Everett, D. L. (2005). Cultural constraints on grammar and cognition in Pirahã: Another look at the design features of human language. *Current Anthropology*, 46, 621-646.
- Fenson, L., Dale, P., Reznick, J., Bates, E., Thal, D., & Pethick, S. (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development*, 59, 1-185.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360, 815-836.
- Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. Kuczaj (Ed.), *Language development: language, cognition, and culture* (pp. 301-334). Hillsdale, NJ: Erlbaum.

- Golinkoff, R. M. & Hirsh-Pasek, K. (2008). How toddlers begin to learn verbs. *Trends in Cognitive Sciences*, 12, 397-403.
- Golinkoff, R. M., Jacquet, R., Hirsh-Pasek, K., & Nandakumar, R. (1996). Lexical principles may underlie the learning of verbs. *Child Development*, 67, 3101-3119.
- Gordon, P. (2004). Numerical cognition without words: Evidence from Amazonia. *Science*, 306, 496-499.
- Gredebäck, G. & Melinder, A. (2010). Infants' understanding of everyday social interactions: A dual process account. *Cognition*, 114, 197-206.
- Gweon, H., & Schulz, L. E., (2011). 16-month-olds rationally infer causes of failed actions. *Science*, 332, 1524.
- György, G., Nádasdy, Z., Csibra, G., & Biró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, 56, 165-193.
- Hayden, B. Y., Heilbronner, S. R., Pearson, J. M., & Platt, M. L. (2011). Surprise signals in anterior cingulate cortex: Neuronal encoding of unsigned reward predictions errors driving adjustment in behavior. *The Journal of Neuroscience*, 31, 4178-4187.
- Hollerman, J. R. & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, 1, 304-309.
- Imai, M., Li, L., Haryu, E., Okada, H., Hirsh-Pasek, K., Golinkoff, R. M., & Shigematsu, J. (2008). Novel noun and verb learning in Chinese-, English-, and Japanese-speaking children. *Child Development*, 79, 979-1000.
- Itti, L. & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research*, 49, 1295-1306.

- Izard, V., Sann, C., Spelke, E.S., & Streri, A. (2009). Newborn infants perceive abstract numbers. *Proceedings of the National Academy of Sciences*, 106(25), 10382-10385.
- Johnson, C. N., & Harris, P. L. (1994). Magic: Special but not excluded. *Developmental Psychology*, 12, 35-51.
- Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279-296). New York: Appleton-Century-Crofts.
- Kidd, C., Piantadosi, S. P., & Aslin, R. N. (2012). The Goldilocks Effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS ONE*, 7, 1-8.
- Kidd, C., Piantadosi, S. P., & Aslin, R. N. (2014). The Goldilocks effect in infant auditory cognition. *Child Development*, 85, 1795-1804.
- Kouider, S., Long, B., Le Stanc, L., Charron, S., Fievet, A.-C., Barbosa, L. S., & Gelskov, S. V. (2015). Neural dynamics of prediction and surprise in infants. *Nature Communications*, 6, 1-7.
- Kuhlmeier, V., Wynn, K., & Bloom, P. (2003). Attribution of dispositional states by 12-month-old infants. *Psychological Science*, 14, 402-408.
- Kundey, S. M., De Los Reyes, A., Taglang, C., Baruch, A., & German, R. (2010). Domesticated dogs' (*Canis familiaris*) use of the solidity principle. *Animal Cognition*, 13, 497-505.
- Legare, C. H. (2012). Exploration explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Development*, 83, 173-185.

- Legare, C.H. (2014). The contributions of explanation and exploration to children's scientific reasoning. *Child Development Perspectives*, 8, 101-106.
- Legare, C. H., Gelman, S. A., & Wellman, H. M. (2010). Inconsistency with prior knowledge triggers children's causal explanatory reasoning. *Child Development*, 81, 929-944.
- Legare, C. H., Schult, C., Impola, M., & Souza, A. L. (2016). Young children revise explanations in response to new evidence. *Cognitive Development*, 39, 45-56.
- Leslie, A.M. & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, 25(3), 265-288.
- Luo, Y. & Baillargeon, R. (2007). Do 12.5-month-old infants consider what objects others can see when interpreting their actions? *Cognition*, 105, 489-512.
- Maguire, M., Hirsh-Pasek, K., & Golinkoff, R. M. (2006). A unified theory of word learning: Putting verb acquisition in context. In K. Hirsh-Pasek & R. M. Golinkoff (Eds.), *Action Meets Word: How Children Learn Verbs* (pp. 364-391). New York: Oxford University Press.
- McCrink, K. & Wynn, K. (2004). Large-number addition and subtraction by 9-month-old infants. *Psychological Science*, 15, 776-781.
- Meristo, M., & Surian, L. (2013). Do infants detect indirect reciprocity? *Cognition*, 129, 102-113.
- Needham, A. & Baillargeon, R. (1993). Intuitions about support in 4.5-month old infants. *Cognition*, 47, 121-148.

- Nielson, K.A., Yee, D., & Erickson, K.I. (2005). Memory enhancement by a semantically unrelated emotional arousal source induced after learning. *Neurobiology of Learning and Memory*, 84, 49-56.
- Onishi, K. H. & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, 308, 255-258.
- Pearce, J.M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532-552.
- Perfors, A., Tenenbaum, J. B., Griffith, T. L., & Xu, F. (2011). A tutorial introduction to Bayesian models of cognitive development. *Cognition*, 120, 302-321.
- Phelps, K. E., & Woolley, J. D. (1994). The form and function of children's magical beliefs. *Developmental Psychology*, 30, 385-394.
- Rescorla, R. A. & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black and W.F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- Roesch, M. R., Esber, G. R., Li, J., Daw, N. D., & Schoenbaum, G. (2012). Surprise! Neural correlates of Pearce-Hall and Rescorla-Wagner coexist within the brain. *European Journal of Neuroscience*, 35, 1190-1200.
- Rosengren, K. S., & Hickling, A. K. (1994). Seeing is believing: Children's explanations of commonplace, magical, and extraordinary transformations. *Child Development*, 65, 1605-1626.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926-1928.

- Santos, L. R., & Hauser, M. D. (2002). A non-human primate's understanding of solidity: Dissociations between seeing and acting. *Developmental Science*, 5, F1-F7.
- Schultz, W., Dayan, P., & Montague, R. R. (1997). A neural substrate of prediction and reward. *Science*, 275, 1593-1599.
- Schultz, W. & Dickinson, A. (2000). Neuronal coding of prediction errors. *Annual Review of Neuroscience*, 23, 473-500.
- Schulz, L. E. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences*, 16, 382-389.
- Schulz, L. E. & Bonawitz, E. B. (2007). Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, 43, 1045-1050.
- Schulz, L.E., Standing, H., & Bonawitz, E. B. (2008). Word, thought and deed: The role of object labels in children's inductive inferences and exploratory play. *Developmental Psychology*, 44, 1266-1276.
- Schwabe, L., Bohringer A., Chatterjee, M., & Schachinger, H. (2008). Effects of pre-learning stress on memory for neutral, positive, and negative words: Different roles of cortisol and autonomic arousal. *Neurobiology of Learning and Memory*, 90, 44-53.
- Shtulman, A. (2009). The development of possibility judgment within and across domains. *Cognitive Development*, 24, 293-309.
- Shtulman, A., & Carey, S. (2007). Improbable or impossible? How children reason about the possibility of extraordinary events. *Child Development*, 78, 1015-1032.
- Spelke, E.S. (1990). Principles of object perception. *Cognitive Science*, 14, 29-56.
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of

- knowledge. *Psychological Review*, 99, 605-632.
- Spelke, E. S., Kestenbaum, R., Simons, D., & Wein, D. (1995). Spatiotemporal continuity, smoothness of motion and object identity in infancy. *British Journal of Developmental Psychology*, 13, 113-142.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, 10, 89-96.
- Stahl, A. E. & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348, 91-94.
- Stahl, A. E., Romberg, A. R., Roseberry, S., Golinkoff, R. M., & Hirsh-Pasek, K. (2014). Infants segment continuous events using transitional probabilities. *Child Development*, 85, 1821-1826.
- Téglás, E., & Bonatti, L. L. (2016). Infants anticipate probabilistic but not deterministic outcomes. *Cognition*, 157, 227-236.
- Téglás, E., Girotto, V., Gonzalez, M., & Bonatti, L. L. (2007). Intuitions of probabilities shape expectations about the future at 12 months and beyond. *Proceedings of the National Academy of Sciences*, 104, 19156-19159.
- Téglás, E., Vul, E., Girotto, V., Gonzalez, M., Tenenbaum, J. B., & Bonatti, L. L. (2011). Pure reasoning in 12-month-old infants as probabilistic inference. *Science*, 332, 1054-1058.
- Tenenbaum, J. B., Kemp, C., Griffiths, T.L., & Goodman, N.D. (2011). How to grow a mind: Statistics, structure and abstraction. *Science*, 331, 1279-1285.
- Valdesolo, P., Shtulman, A., & Baron, A. S. (in press). Science is Awe-some: Emotional antecedents of science learning. *Emotion Review*.

- van Schijndel, T. J., Visser, I., van Bers, B. M., & Raijmakers, M. E. (2015). Preschoolers perform more informative experiments after observing theory-violating evidence. *Journal of Experimental Child Psychology, 131*, 104-119.
- Waeleti, P., Dickinson, A., & Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. *Nature, 412*, 43
- Walden, T. Kim, G., McCoy, C., & Karass, J. (2007). Do you believe in magic? Infants' social looking during violations of expectations. *Developmental Science, 10*, 654-663.
- Waxman, S. R., Lidz, J. L., Braun, I. E., & Lavin, T. (2009). Twenty-four-month-old infants' interpretations of novel verbs and nouns in dynamic scenes. *Cognitive Psychology, 59*, 67-95.
- Weisberg, D. S., & Sobel, D. M. (2012). Young children discriminate improbable from impossible events in fiction. *Cognitive Development, 27*, 90-98.
- Wilcox, T. (1999). Object individuation: Infants' use of shape, size, pattern, and color. *Cognition, 72*, 125-166.
- Wilcox, T., & Baillargeon, R. (1998). Object individuation in infancy: The use of featural information in reasoning about occlusion events. *Cognitive Psychology, 37*, 97-155.
- Wilcox, T., Bortfeld, H., Woods, R., Wruck, E., & Boas, D. A. (2005). Using near-infrared spectroscopy to assess neural activation during object processing in infants. *Journal of Biomedical Optic, 10*, 011010-1-011010-9.
- Wilcox, T., Nadel, L., & Rosser, R. (1996). Location memory in healthy preterm and full-term infants. *Infant Behavior & Development, 19*, 309-323.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition, 69*, 1-34.

Wynn, K. (1992). Addition and subtraction by human infants. *Nature*, 358(6389), 749-750.

Xu, F., & Denison, S. (2009). Statistical inference and sensitivity to sampling in 11-month-old infants. *Cognition*, 112, 97–104.

Xu, F. & Garcia, V. (2008) Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences of the United States of America*, 105, 5012-5015.