TOPICS
TOPICS IN COGNITIVE SCIENCE

Topics in Cognitive Science (2018) 1–18

ISSN: 1756-8765 online
DOI: 10.1111/tops.12389

This article is part of the topic "The Ubiquity of Surprise: Developments in Theory, Converging Evidence, and Implications for Cognition," Edward Munnich, Mark Keane and Meadhbh Foster (Topic Editors). For a full listing of topic papers, see http://onlinelibrary.wiley.com/journal/10.1111/(ISSN)1756-8765/earlyview.

Violations of Core Knowledge Shape Early Learning

Aimee E. Stahl,[a] Lisa Feigenson[b]

[a]*Department of Psychology, The College of New Jersey*
[b]*Department of Psychological & Brain Sciences, Johns Hopkins University*

Received 29 July 2016; received in revised form 17 September 2018; accepted 17 September 2018

Abstract

Research on cognitive development has revealed that even the youngest minds detect and respond to events that adults find surprising. These surprise responses suggest that infants have a basic set of "core" expectations about the world that are shared with adults and other species. However, little work has asked what purpose these surprise responses serve. Here we discuss recent evidence that violations of core knowledge offer special opportunities for learning. Infants and young children make predictions about the world on the basis of their core knowledge of objects, quantities, and social entities. We argue that when these predictions fail to match the observed data, infants and children experience an enhanced drive to seek and retain new information. This impact of surprise on learning is not equipotent. Instead, it is directed to entities that are relevant to the surprise itself; this drive propels children—even infants—to form and test new hypotheses about surprising aspects of the world. We briefly consider similarities and differences between these recent findings with infants and children, on the one hand, and findings on prediction errors in humans and non-human animals, on the other. These comparisons raise open questions that require continued inquiry, but suggest that considering phenomena across species, ages, kinds of surprise, and types of learning will ultimately help to clarify how surprise shapes thought.

Keywords: Infants; Children; Object knowledge; Surprise; Expectations; Learning

Correspondence should be sent to Aimee E. Stahl, Department of Psychology, The College of New Jersey, 2000 Pennington Road, Ewing, NJ 08628. E-mail: stahla@tcnj.edu

1. Introduction

Where does human knowledge come from? Which aspects of our thinking originate from observation and instruction, and which emerge independent of specific experience? The tension between empiricist and nativist accounts of human thoughts and beliefs has long animated research in psychology and the cognitive sciences (Chomsky, 1980; Elman et al., 1996; Samuels, 2002; Skinner, 1957; Spelke & Newport, 1998), generating both lively debate and empirical progress. Although most theorists agree that some combination of innate abilities and experience-driven learning is required to yield the impressive suite of knowledge in place by the first few years of life, this "middle ground" is vast and often lacks specificity. What could it mean to endorse a view on which both experience and innate knowledge have roles to play? One possibility is that some concepts or abilities are acquired through learning, whereas others emerge prior to or without experience. Another (not mutually exclusive) possibility is that innate knowledge constrains and guides what is acquired through learning. An example of this latter idea is the research program that considers young children as Bayesian learners, continually adjusting the likelihood of their hypotheses about the world using new observations as evidence (e.g., Perfors, Tenenbaum, Griffith, & Xu, 2011; Xu & Kushnir, 2013; Xu & Tenenbaum, 2007).

Here, we explore a related sense in which innate knowledge and new learning may interact—in this case, with *surprise* as a driving force for early learning. Over the past several decades, research in cognitive development finds that even preverbal infants respond to surprising events.[1] Characterizing the events they find surprising suggests that infants have rich, interconnected expectations about objects, people, and quantities; they respond when these entities act in ways that adults would judge to be impossible, often by the first weeks or months of life (Spelke & Kinzler, 2007). Infants' surprise reactions have been invaluable in the effort to document the existence of early knowledge. Yet little is known about the consequences of surprise for infants' and children's learning. In this paper we consider the idea that surprise not only marks the detection of an instance in which what was predicted failed to match what was observed, but also that surprise acts as a force shaping what and when children learn. We describe some of our recent work showing that violations of expectations generated from very early, "core" knowledge affect infants' and children's acquisition of new information, both inside and outside the laboratory. This research program is in its early days, and much remains unknown. For example, the idea that babies and preschoolers learn better from surprising events has parallels with findings on learning in non-human animals and in human adults, but it is not yet clear whether these are examples of a single phenomenon, or are merely analogous. Our hope is that by discussing some of our new findings, highlighting commonalities with results across different literatures, and, in places, taking license to speculate, we can contribute to progress on questions such as this. Therefore, what follows is less a review of a mature subfield of research than what we hope is a starting point for further inquiry.

2. Core knowledge in early cognition

Infants' responses to surprising events have been one of the most important tools in the effort to characterize early thought. One experimental method—violation of expectation—has been especially fruitful in revealing the nature of infants' understanding. In this method, infants watch events that either accord with or defy the expectations an adult would generate, and their looking times to these outcomes are compared. Across hundreds of studies, infants reliably look longer at the surprising than the expected outcomes, even when controlling for perceptual differences between the two. Such violation-of-expectation studies measuring looking time—or, alternatively, changes in facial expression (Camras et al., 2002), increases in social referencing (Walden, Kim, McCoy, & Karass, 2007), or changes in brain activity (Berger, Tzur, & Posner, 2006; Wilcox, Bortfeld, Woods, Wruck, & Boas, 2005)—have revealed that infants have basic expectations about aspects of the world nearly from the start of life. At least some of this knowledge appears to be organized into a handful of "core" domains that support thinking about objects, numerical quantities, spatial relationships, and social entities. This "core knowledge" is often conceived of as the product of evolutionary pressures for organisms to solve domain-specific problems, with the solutions conserved over ontogeny and phylogeny (Spelke & Kinzler, 2007 for review). Core knowledge has been documented in a variety of species (e.g., Agrillo, Dadda, Serena, & Bisazza, 2008; Cantlon & Brannon, 2006; Kundey, De Los Reyes, Taglang, Baruch, & German, 2010; Regolin & Vallortigara, 1995; Santos, 2004; Spelke & Lee, 2012) and across diverse human cultures (Dehaene, Izard, Pica, & Spelke, 2006; Gordon, 2004; Pica, Lemer, Izard, & Dehaene, 2004), and in some cases has been shown to be present, in both humans and non-human animals, at birth (Izard, Sann, Spelke, & Streri, 2009; Regolin & Vallortigara, 1995). Finally, because it can emerge under controlled rearing conditions in which animals' opportunities for learning are drastically reduced, core knowledge is often proposed to be independent of specific learning experience (e.g., Chiandetti & Vallortigara, 2010).

One of the earliest explored cases of core knowledge is that of the *object concept*: Infants in their first year of life are sensitive to the behavior of mid-sized physical objects. Spelke (1990, 1994) and Baillargeon (2004) have proposed that this core knowledge of objects is organized around a few central principles of object behavior, each of which has been implicated by numerous violation-of-expectation studies. These studies find that infants are surprised when inanimate objects move of their own accord (e.g., Leslie & Keeble, 1987; Woodward, Phillips, & Spelke, 1993), break into pieces rather than move cohesively (e.g., Kellman & Spelke, 1983; Spelke & Van de Walle, 1993), disappear into thin air, or simultaneously occupy the same space as another object (e.g., Baillargeon, Spelke, & Wasserman, 1985; Spelke, Breinlinger, Macomber, & Jacobson, 1992; Spelke, Kestenbaum, Simons, & Wein, 1995). Longer looking to surprising events also has revealed infants' core knowledge about other aspects of the world, including approximate number (e.g., Feigenson, 2011; McCrink & Wynn, 2004), social agents (e.g., Kuhlmeier, Wynn, & Bloom, 2003; Woodward, 1998), and emotions (Skerry & Spelke, 2014). Because infants have to be at least a few months old to be tested in some of these experimental designs, it is hard to determine whether their expectations are innate, or just acquired very early. However, research in non-human species has found evidence for many similar abilities in controlled rearing conditions that do not offer any relevant learning experience (e.g., object permanence in chicks reared without having seen a single instance of object occlusion). These studies lend credence to the conclusion that core knowledge in humans may also be, at least in part, innate (Vallortigara, 2012; Vallortigara, Regolin, Chiandetti, & Rugani, 2010).

While such evidence has been used to argue that very young infants are equipped with basic expectations about the world, a frequent criticism of the idea of core knowledge is that it denies or ignores the contributions of development and experience. This complaint sees core knowledge as static and unresponsive to the environment. For example, Spencer and colleagues suggest that "developmental psychologists should no longer embrace 'endowments', 'primitives', 'core knowledge', 'essences', or other *static concepts* that *devalue developmental process*" (emphasis added) (Spencer et al., 2009, p. 79). But does core knowledge devalue development or ignore learning? An alternative is that having some basic expectations in place from very early in development, even before associative expectations about the world have been accrued, may actually scaffold learning.

3. How might core knowledge violations influence cognition?

One way in which core knowledge might plausibly affect learning is by reducing the learning space confronting the observer. Faced with an infinite number of objects, object parts, events, relations, and features that could potentially be learned about, violations of expectation might focus the learner's cognitive resources on aspects of the world for which the learner had the wrong prior model—about which the learner made an incorrect prediction. Targeting resources to these, rather than entities about which the learner's predictions were already accurate, would be one way to constrain the problem of the infinite learning space—and thereby make learning more effective. Our own recent work has begun to explore this idea in a series of experiments that ask whether violations of core knowledge help infants learn. Such an effect is far from a foregone conclusion. Core knowledge violations may be detected but have no measurable effect on the learning that follows. Or their effect could be detrimental: Core knowledge violations might be so striking—because they violate deeply held beliefs about the workings of the world—that they instead disrupt further processing. Indeed, in at least some situations, children learn better from predictable events (visual sequences) than unpredicted events (Benitez & Saffran, 2018).

In our experiments we focused on core knowledge of objects, targeting principles that have been observed in very young infants and newborn animals (e.g., Spelke et al., 1992; Vallortigara, Regolin, Rigoni, & Zanforlin, 1998). We showed 11-month-old infants an event that either accorded with or violated core expectations about object behavior. One group saw an object (e.g., a ball) roll down a ramp and either appear to be stopped by a solid wall in its path (an outcome that adults find unsurprising) or appear to pass straight through the wall, thereby violating the principle of object solidity. A separate group saw an event that accorded with or defied a different principle of object behavior: support. These infants either saw an object (e.g., a car) pushed along the top of a box while remaining fully supported (unsurprising) or saw the object pushed over the box's edge but appear to hover in mid-air despite the lack of support (surprising). In all cases, infants experienced just a single trial that was either expected or surprising, so there was no opportunity to learn about the object's behavior during the experiment itself.

Previous studies using the violation-of-expectation method find that young infants look longer at the surprising outcomes of such solidity and support events (e.g., Baillargeon & Hanko-Summers, 1990; Spelke et al., 1992). Our question was different: We wanted to know whether seeing a violation of core knowledge would change infants' exploratory actions or their learning. To find out, immediately following the solidity or support event we gave infants two objects: the object from the preceding event and an entirely new distractor object. We found that infants who had just seen the expected outcome of either event (ball stopped by wall, or car fully supported by box) showed the standard preference for novelty; they chose to explore the new distractor object over the familiar object that had behaved according to core principles of objecthood. However, infants who had just seen the surprising outcome of either event (ball passed through wall, or car hovered in mid-air) reversed this preference, spending most of their time exploring the toy that had just violated their expectations (Stahl & Feigenson, 2015). This suggests that infants were motivated to learn about objects that behaved in surprising ways—that violations of core object knowledge propelled them to seek further information.

What information were infants seeking during these exploratory bouts? Analyses of infants' behavior revealed that infants explored the target object in qualitatively different ways depending on the type of core knowledge violation they had seen. Infants who saw a solidity violation, in which the object passed through the wall, tended to explore that object by banging it with their hands or against the high chair tray. Infants who saw a support violation, in which the object appeared to float in mid-air, repeatedly dropped it (Stahl & Feigenson, 2015). This suggests that infants not only notice events that violate their expectations, and not only choose to explore objects that behave in surprising ways, but also that they test specific hypotheses about the violations they see. Hence, violations of core knowledge expectations cause infants to seek information that is relevant to explaining the violation.

These findings show that predictions generated from core knowledge can change the opportunities infants have for learning. Do they also change learning itself? Beyond asking whether infants are motivated to explore surprising events, we asked whether infants *learn better* from surprising events. We again showed 11-month-olds an event that accorded with or violated core expectations of object behavior. Half of the infants saw a solidity event with an expected or surprising outcome (e.g., ball stopped by wall, or ball passed through wall), and half saw a spatiotemporal continuity event with either an expected or surprising outcome (e.g., ball hidden behind left screen and revealed behind that same screen, or ball hidden behind left screen but revealed behind right screen). We then taught infants new information about the ball: that it had a hidden auditory property (e.g., it squeaked). The subsequent test trial asked whether infants had mapped the sound to the object. We found that infants who saw the expected outcome of either event failed to learn the new information—as predicted, because we had designed the one-shot-learning task to be hard. In contrast, infants who saw the surprising outcome of either event successfully learned the object's hidden property (Stahl & Feigenson, 2015). Thus, infants' surprise at seeing violations of core knowledge expectations enhanced their learning.

Most investigations of core knowledge have focused on preverbal infants. But core knowledge, critically, is thought to be continuous across development and to guide intuitions throughout the lifespan (Spelke & Kinzler, 2007). If older learners also rely on core knowledge, they too might experience enhanced learning when core knowledge expectations are violated. We asked whether this was so with children well beyond infancy, using an explicit task rather than a looking-time task. We tested preschoolers and early school-aged children in the dynamic setting of a busy science museum (Stahl & Feigenson, 2017). We first showed 3- to 6-year-olds an event that accorded with or violated the core principle of spatiotemporal continuity: A toy was hidden under one of two blue cups and was either revealed under the original cup (as expected) or magically under the other cup. For all children, immediately after the object was revealed, the experimenter labeled the action just performed, pointing to the cups and saying, "These cups *blicked* the toy!" Children were then tested on their learning of this novel verb: Children saw a blue cup with two other distractor cups (that had been previously labeled with verbs to describe novel but possible actions) and were asked to point to the cup that could *blick* the toy.

We found that children who saw the expected outcome performed at chance, showing no evidence of having learned the novel word. Again, this was as predicted because we had designed the task to be challenging. But children who saw the surprising outcome robustly learned the novel verb, despite having had just one exposure, and despite verbs being notoriously difficult to learn (e.g., Golinkoff & Hirsh-Pasek, 2008). We also replicated these findings with a different core knowledge violation (featural continuity, in which an object appeared to magically change its features) and a different type of verbal label (nouns; e.g., the blue cups were labeled as *blickers*). Our experiments ruled out the possibility that children in the expected condition of the above experiment only failed to learn what *blicking* was because the event was not novel enough to merit a novel label, and they confirmed that children truly learned the labels following surprising events, rather than simply being attracted to objects that behaved surprisingly (Stahl & Feigenson, 2017). These results confirm that surprising events enhance older children's learning of new information, even outside of tightly controlled laboratory conditions.

4. What is the scope of the effect of core knowledge violations on cognition?

Violations of predictions generated from core knowledge appear to present special opportunities to learn. What is the scope of this learning enhancement, and how might seeing a violation of core knowledge affect learning? One way that surprise might enhance learning is by increasing overall arousal or attention, thereby benefitting learning in a diffuse way. Such generalized learning boosts have been seen in adults. For example, increasing attention to a target task can enhance memory for information presented in an unrelated concurrent task (Swallow & Jiang, 2010), and being more curious about one type of information can enhance memory for unrelated incidental information (Gruber, Gelman, & Ranganath, 2014). These effects point to mechanisms that affect learning in a non-targeted fashion. If violations of core knowledge also enhance learning in a general way, then infants and children should show better learning about any entity, as long as the surprising event directly precedes learning.

To begin to address this, we again presented infants and children with events that violated core object knowledge. This time instead of teaching them about the very object that had just violated their expectations, we immediately taught them about an entirely new object that had not participated in the violation. Infants and children were no better at learning about this "innocent bystander" object when taught following a surprising violation event than following an expected event (Stahl & Feigenson, 2015, 2017). These findings suggest that violations of core knowledge likely do not affect learning diffusely, for example, by increasing overall arousal. Rather, learning improves for the particular entities involved in failed predictions.

A second question about the scope of surprise's effect on early learning also concerns the specificity of learning—in this case, how far newly learned information is extended. Our first studies showed that infants who saw a violation of core object knowledge (e.g., ball passes through wall) learned better about that object, explored that object more, and tested hypotheses for the object's surprising behavior. These enhancements do not extend to objects that are entirely unrelated to the surprising event (Stahl & Feigenson, 2015). But are they focused on object tokens—in which case learning and exploration should only change for the very object that had violated expectations (e.g., the ball that had passed through a wall)? Or do the enhancements occur at a more abstract level—in which case learning and exploration might also change for objects perceived as similar to the object that violated expectations?

Previous research has shown that infants readily generalize object properties to new exemplars from the same category (e.g., Baldwin, Markman, & Melartin, 1993; Mandler & McDonough, 1996; Welder & Graham, 2001). In recent work we have begun to ask whether infants also generalize their surprise-induced responses beyond the particular objects that violated core knowledge expectations (Stahl & Feigenson, unpublished data). We showed 11-month-olds the same events as in our previous work. For example, infants saw a red- and blue-striped ball pass through a wall in its path. Then, instead of giving them the opportunity to explore that very ball, we gave infants an object from the same kind category as in the preceding event, but with novel features, and an object from a completely different category (e.g., a car). This allowed us to ask whether seeing one ball behave surprisingly would change the way infants interacted with (i.e., attempted to learn about) balls more generally. Our findings to date suggest that infants did not generalize the effects of seeing the core knowledge violation. Infants who saw the surprising outcome showed no significant preference for the new ball over the car, and their preferences did not differ from those of infants who saw the ball behave in an expected way. Although preliminary, these results suggest that infants reserve their surprise-induced information-seeking for the very entities that violated their expectations (Stahl & Feigenson, unpublished data).

In sum, this emerging body of work suggests that infants and children can harness their core knowledge to guide learning. When they see an event that violates a principle of core knowledge (e.g., solidity, support, spatiotemporal continuity, featural continuity), they increase their exploration, produce hypothesis-testing behaviors, and show enhancements in their ability to learn new object-sound mappings and novel words. These effects of surprise on learning are not equipotent—they are reserved for entities directly involved in the core knowledge violation, not unrelated or even similar entities.

5. Are violations of core knowledge privileged for learning?

Is it surprising that violations of core knowledge affect learning in infants and young children? One the one hand, we suggest that the answer is yes. Despite the decades of research showing that infants detect violations in the behavior of objects, quantities, and other people, it was not previously known whether these detection responses had any cognitive consequences at all. But in another sense, the finding that core knowledge violations change learning fits well with a number of existing ideas. First, the claim that expectancy violations might promote learning has a long history in a literature that has proceeded almost entirely separately from the literature on infant cognition: the study of incremental, associative learning in non-human animals and human adults. Classic theories of associative learning emphasize the notion of prediction error—the difference between the event an animal expected and the event it observed—as a powerful learning impetus. When the animal experiences an event that violates learned contingencies (i.e., is surprising), it modifies its behavior to better predict future events. The larger the prediction error, the faster the learning. In contrast, when a novel stimulus is paired with an event already fully predicted based on previously experienced contingencies, no error is generated and no learning about the new stimulus occurs (e.g., Kamin, 1969; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Schultz & Dickinson, 2000). Researchers have identified the neural underpinnings of such error-driven learning in both non-human animals (e.g., Friston, 2005; Hayden, Heilbronner, Pearson, & Platt, 2011; Holland & Gallagher, 2006; Hollerman & Schultz, 1998; Schultz, Dayan, & Montague, 1997) and human adults, who also generate prediction errors to violations of learned contingencies (e.g., Behrens, Woolrich, Walton, & Rushworth, 2007; Brown & Braver, 2005; Carter et al., 1998; Fletcher et al., 2001; McClure, Berns, & Montague, 2003; Mestres-Missé, Trampel, Turner, & Kotz, 2017; O'Reilly, 2013; den Ouden, Friston, Daw, McIntosh, & Stephan, 2009; Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006; Yu & Dayan, 2005). Evidence suggests that in both non-humans and humans, dopamine neurons play a key role in signaling prediction error and in driving neuronal and behavioral learning (e.g., Holroyd & Coles, 2002; Waelti, Dickinson, & Schultz, 2001).

The effects of expectancy violation on learning also have been considered in parallel in the literature on children's thinking. This tradition, too, has a long history—at least as far back as Piaget, who suggested that children experience some aspects of the world as fitting neatly within existing mental models (in which case new information is *assimilated* into children's representations), and other aspects of the world as conflicting with existing beliefs or expectations (in which case the new information requires *accommodation*) (Piaget, 1970; see also Kagan, 2002). Contemporary research in cognitive development finds that young children indeed modify their exploration and explanations following events that fail to accord with expectations (e.g., Bonawitz, van Schijndel, Friel, & Schulz, 2012; Legare, 2012; Legare, Gelman, & Wellman, 2010; Legare, Schult, Impola, & Souza, 2016; van Schijndel, Visser, van Bers, & Raijmakers, 2015; Schulz, 2012; Schulz & Bonawitz, 2007). For example, children prefer to explore a familiar toy over a novel one when the familiar toy produced an outcome that was unexpected (Bonawitz et al., 2012), and they show an increased tendency to produce verbal explanations of events that do not fit with newly acquired knowledge (Legare et al., 2010, 2016).

Even infants use violations of recently learned causal relationships to guide behavior. They more persistently attempt to produce behaviors that had previously been effective at activating toys that no longer work, relative to when the toy behaved as expected (Baldwin et al., 1993). Furthermore, infants show some ability to infer the reasons for objects' failure to produce a learned effect. When the data were consistent with a violation being caused by a problem with the toy (e.g., toy might have been broken), infants sought a different toy with which to play. When the data were consistent with the violation being caused by a problem with the infant's own actions (e.g., the infant might not have properly activated the toy), infants instead sought adults' help (Gweon & Schulz, 2011).

The findings reviewed above—from non-human animals, human adults, and children—show that there are cognitive and behavioral consequences when learned contingencies change unexpectedly. When expectations are violated by new evidence, learning can be enhanced (as in the work on prediction error) and information-seeking behaviors promoted (as in the work on children's exploration and explanation). An open question is whether the learning enhancements observed after core knowledge violations reflect the same phenomenon as the learning that follows changes to newly learned contingencies. Does surprise affect learning in the same way when a ball appears to pass through a solid wall, compared to when a previously rewarded tone is suddenly no longer followed by a reward? While our experiments cannot directly address this question, the findings that infants and children require only a single exposure to a core knowledge violation—and no training within the experimental context—to learn new information that typically requires extensive exposure (e.g., verb learning; Golinkoff & Hirsh-Pasek, 2008) suggests that core knowledge violations might be particularly powerful for propelling learning.

One framework that might capture the potency of core versus acquired knowledge on learning comes from the Bayesian approach. In Bayesian models, learners are hypothesized to have existing hypotheses about the world. These priors have varying strengths, and they might be innate (e.g., core knowledge) or learned. When learners observe an event, they adjust their priors to posterior probabilities—the greater the update (i.e., the larger the discrepancy between the prior and the posterior), the greater the surprise, and perhaps the greater the subsequent learning (e.g., Baldi & Itti, 2010; Courville, Daw, & Touretzky, 2006; Perfors et al., 2011; Schulz, 2012; Téglás et al., 2011; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). It may that priors set by core knowledge are stronger than priors set by acquired knowledge, and therefore that they lead to greater learning enhancement. This version of the difference between violations of core knowledge driving learning, on the one hand, and violations of acquired associations driving learning, on the other, is quantitative in nature. Whether this is the right characterization remains to be seen, and it will require extended efforts using multiple empirical approaches.

However, as a first step, we have recently begun to ask whether young learners experience enhanced learning from surprising events that are improbable, rather than impossible. Infants certainly detect violations to probabilistic associations, in addition to violations of core knowledge. For example, infants and young children look longer and show physiological responses when an established auditory or visual statistical pattern suddenly changes (e.g., Emberson, Richards, & Aslin, 2015; Kirkham, Slemmer, & Johnson, 2002; Kouider et al., 2015; Saffran, Aslin, & Newport, 1996; Stahl, Romberg, Roseberry, Golinkoff, & Hirsh-Pasek, 2014). Furthermore, infants can infer the likelihood of an event using the statistics of visual scenes. They are surprised, for example, when “random” draws from a container of mostly white balls repeatedly produce red balls, even though such an outcome is not impossible (Sim & Xu, 2017; Téglás, Girotto, Gonzalez, & Bonatti, 2007; Téglás et al., 2011; Xu & Denison, 2009; Xu & Garcia, 2008). Does the surprise generated by these types of improbable events also promote learning in infants and young children? Or are impossible events processed differently from events that are merely improbable (Téglás & Bonatti, 2016)?

As a first start at answering this, we showed 2- and 3-year-old children a population of objects inside a transparent gumball machine (Stahl & Feigenson, unpublished data). Children saw that putting a coin in the machine made one object from the population emerge through a chute (apparently at random). We first replicated our previous finding that impossible events enhance children’s learning relative to possible events. One group of children saw that the machine contained 10 identical novel purple toys and 10 identical novel pink toys, all intermixed. Children had time to view this population, and then they were prompted to insert the coin. This caused a pink toy to come out—an outcome that was not surprising, as each type of toy had a 50% likelihood of being drawn. Children were then taught a novel word for the pink toy. A separate group of children saw that the machine contained 20 identical novel purple toys. When they put the coin in, a pink toy also emerged (through laboratory trickery). This event was designed to appear impossible, because no pink toys were in the machine and therefore there was a 0% chance of a pink toy being drawn. These children were also taught the novel word for the pink toy. Finally, we tested children’s learning of the novel word by asking them to identify its referent from an array. As in our previous work (Stahl & Feigenson, 2015, 2017), children who saw the probable outcome (50% likelihood) failed to learn the novel word, whereas children who saw the impossible outcome (0% likelihood) succeeded (Stahl & Feigenson, unpublished data).

Critically, three additional groups of children saw events that were improbable rather than impossible. These children saw a pink toy emerge from a machine that contained either 2 pink and 18 purple toys (10% likelihood of drawing the pink), 1 pink and 19 purple toys (5% likelihood), or 1 pink and 39 purple toys (2.5% likelihood) (we ensured that children always detected the pink toy within the population prior to inserting the coin). In all these cases, children failed to learn the novel word used to label the pink toy. Interestingly, there was no evidence of graded performance; learning did not improve as the event became more surprising (i.e., as the likelihood of drawing the pink toy decreased). Instead, performance was step-like: equivalent failures to learn across all the improbable events, no matter how unlikely, and a sudden boost in learning for the impossible event (Stahl & Feigenson, unpublished data). The literature on prediction error finds that the size of prediction errors (and resultant learning) scales continuously with the extent to which a stimulus was statistically expected (e.g., Abler, Walter, Erk, Kammerer, & Spitzer, 2006; Schönberg, Daw, Joel, & O'Doherty, 2007). Our data from children, although preliminary, do not seem to follow this pattern. Rather, our findings suggest that impossible events might produce an immediate drive to learn, whereas possible events (no matter how unlikely) do not, at least under the conditions we tested.

One might question the utility of a learning mechanism that benefits from the impossible. Why learn about what cannot happen? Other labs' recent work comparing infants' and children's responses to events that were surprising to varying degrees indirectly bears on this question. One study found that infants tend to allocate their attention to events that are moderately unpredictable (in terms of their statistical likelihood within the experimental session), relative to events that are highly predictable or highly unpredictable (the "Goldilocks Effect"; Kidd, Piantadosi, & Aslin, 2012, 2014). Similarly, children show better memory for stories that contain neither too few nor too many surprising facts (Banerjee, Haque, & Spelke, 2013), indicating that engaging with moderately surprising concepts might support learning. Although it may seem counterintuitive, we tentatively suggest that the core knowledge violations we presented actually might have been optimally surprising in this same sense. Although physically impossible, the events were structured such that observers would be likely to generate a prediction about their outcome. Just one aspect of that prediction was then violated (e.g., a ball should have been stopped by the wall in its path, but instead it appeared to pass through). The surprising outcome—the ball unexpectedly resting on the far side of the wall—was only minimally distinct from the expected outcome—the ball resting on the near side of the wall. When one considers the vast range of possible surprising outcomes that did *not* occur (ball passes through wall and changes identity, ball and wall disappear entirely, ball disappears and five new objects appear, etc.), the impossible outcomes in our experiments seem perhaps less surprising and, critically, highly interpretable using the mental vocabulary with which infants likely represented the initial scene. In this way, our events may have been "moderately impossible," such that they stimulated more learning than merely improbable events, and more than outlandishly impossible events.

Still, although infants may in fact represent impossible events differently from improbable ones (Téglás & Bonatti, 2016), it remains unclear whether this distinction correctly captures the learning dissociation we observed in the gumball machine experiments. It could be, rather, that what promotes learning is the need for causal explanation (see Adler, 2008; Foster & Keane, 2015). Core knowledge expectations reflect non-arbitrary commitments about the workings of the world—commitments that, if overturned, would have far-reaching implications for the way we reason about many aspects of our environment. If it were actually true that, for example, an object that was absent a moment ago could suddenly appear (like the pink object that was never in the gumball machine but then magically emerged), this would radically change our beliefs about how objects behave in the material world. Violations of this type of core knowledge expectation may trigger learners to search for the underlying cause of the surprising event. For example, the learner might revise his or her view of how the machine works (there was a secret compartment) or of the toys themselves (the toys might look different from different visual perspectives). In contrast, arbitrary contingencies or statistical likelihoods may require less belief revision to explain; hence, violations of these types may prompt a shallower search for information. The pink toy was rarer than the purples, but seeing it emerge from the population requires no reformulation of prior concepts about the machine or the toys. On this view, improbable events should also be able to enhance learning, to the extent that they propel learners to seek causes. For example, seeing a person repeatedly drawing a rare pink toy from a population of mostly purples is certainly not impossible, but its statistical unlikelihood, coupled with the opportunity to construe the event as goal-directed, might cause learners to make inferences about the underlying cause for the improbable event (e.g., "She must be finding that pink one because she prefers it"). We fully acknowledge that these ideas are speculative at present—and what this speculation underscores, again, is that future work is needed. In this case, it will be useful to find a way to quantify the degree of surprise learners experience and then to more carefully map out the extent to which this predicts subsequent learning.

6. Conclusions

Here we reviewed findings converging on the idea that failed predictions enhance early learning. Our own recent work suggests that this phenomenon guides learning from the start of human life—it is seen in preverbal infants, and it relies on predictions that, coming from core knowledge, require little or no prior experience. In addition, this learning phenomenon extends beyond infancy, as surprise influences learning in older children as well, and does so outside of the controlled laboratory setting. However, we see some of the broader questions arising from this body of research as being wide open. A key issue is whether there is a unitary phenomenon of surprise that shapes learning across the wide range of results reviewed here. Alternatively, does "surprise" mean something different in the context of core knowledge versus associative learning, in humans versus non-human animals, in human learning versus machine learning? Although we see no definitive answers at present, the diversity of contributions to this volume represents an important step. The time is right for researchers across disciplines to seriously consider what is meant by "surprise" and to think about how broadly they expect the signature effects observed in their own domain to apply across other domains and preparations. Researchers in cognitive development have much to gain from the formal approach taken in classic work on animal learning, and perhaps from the neuroscience approach taken in the literature on reward prediction in humans and non-human animals. The contributions of cognitive development to this dialogue are newer and sparser. However, in our view, data from infants and young children offer unique benefits. These data encourage us to consider learning over a longer timescale (changes over years, not just over experimental trials) and highlight the importance of thinking about how different types of knowledge may engender different types of predictions. Characterizing the effects of violating these predictions, in learners and problems of all types, should be a major goal in the effort to understand human thought.

Acknowledgments

This work was supported by National Science Foundation Graduate Research Fellowship DGE-1232825 to A.E.S., National Institutes of Health grant HD007407 to L.F., and by a grant to the Museum of Science, Boston from the National Science Foundation under Award Number 1113648.

Note

1. Here we are neutral about whether infants experience surprise in the same affective sense as adults do. We use *surprise* to refer to the attentional state of heightened interest following an unpredicted event.

References

Abler, B., Walter, H., Erk, S., Kammerer, H., & Spitzer, M. (2006). Prediction error as a linear function of reward probability is coded in human nucleus accumbens. *NeuroImage*, *31*, 790–795.

Adler, J. E. (2008). Surprise. *Educational Theory*, *58*, 149–173.

Agrillo, C., Dadda, M., Serena, G., & Bisazza, A. (2008). Do fish count? Spontaneous discrimination of quantity in female mosquitofish. *Animal Cognition*, *11*, 495–503.

Baillargeon, R. (2004). Infants' physical world. *Current Directions in Psychological Science*, *13*, 89–94.

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.

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*, 1214–1221.

Benitez, V. L., & Saffran, J. R. (2018). Predictable events enhance word learning in toddlers. *Current Biology*, *28*, 2787–2793.

Berger, A., Tzur, G., & Posner, M. I. (2006). Infant brains detect arithmetic errors. *Proceedings of the National Academy of Sciences*, *103*, 12649–12653.

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.

Brown, J. W., & Braver, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science*, *307*, 1118–1121.

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.

Cantlon, J. F., & Brannon, E. M. (2006). Shared system for ordering small and large numbers in monkeys and humans. *Psychological Science*, *7*, 401–406.

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, *280*, 747–749.

Chiandetti, C., & Vallortigara, G. (2010). Experience and geometry: Controlled-rearing studies with chicks. *Animal Cognition*, *13*, 463–470.

Chomsky, N. (1980). Rules and representations. *Behavioral and Brain Sciences*, *3*, 1–15.

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.

den Ouden, H. E. M., Friston, K. J., Daw, N. D., McIntosh, A. R., & Stephan, K. E. (2009). A dual role for prediction error in associative learning. *Cerebral Cortex*, *19*, 1175–1185.

Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness*. Cambridge, MA: MIT Press.

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.

Feigenson, L. (2011). Predicting sights from sounds: 6-month-olds' intermodal numerical abilities. *Journal of Experimental Child Psychology*, *110*, 347–361.

Fletcher, P. C., Anderson, J. M., Shanks, D. R., Honey, R., Carpenter, R. A., Donovan, T., Papadakis, N., & Bullmore, E. T. (2001). Responses of human frontal cortex to surprising events are predicted by formal associative learning theory. *Nature Neuroscience*, *4*, 1043–1048.

Foster, M. I., & Keane, M. T. (2015). Why some surprises are more surprising than others: Surprise as a metacognitive sense of explanatory difficulty. *Cognitive Psychology*, *81*, 74–116.

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*, 815–836.

Golinkoff, R. M., & Hirsh-Pasek, K. (2008). How toddlers begin to learn verbs. *Trends in Cognitive Sciences*, *12*, 397–403.

Gordon, P. (2004). Numerical cognition without words: Evidence from Amazonia. *Science*, *306*, 496–499.

Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, *84*, 486–496.

Gweon, H., & Schulz, L. E. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, *332*, 1524.

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.

Holland, P. C., & Gallagher, M. (2006). Different roles for amygdala central nucleus and substantia innominata in the surprise-induced enhancement of learning. *The Journal of Neuroscience*, *26*, 3791–3797.

Hollerman, J. R., & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, *1*, 304–309.

Holroyd, C. B., & Coles, M. G. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, *109*, 679.

Izard, V., Sann, C., Spelke, E. S., & Streri, A. (2009). Newborn infants perceive abstract numbers. *Proceedings of the National Academy of Sciences*, *106*, 10382–10385.

Kagan, J. (2002). *Surprise, uncertainty, and mental structures*. Cambridge, MA: Harvard University Press.

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.

Kellman, P. J., & Spelke, E. S. (1983). Perception of partly occluded objects in infancy. *Cognitive Psychology*, *15*, 483–524.

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.

Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, *83*, B35–B42.

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., 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.

Mandler, J. M., & McDonough, L. (1996). Drinking and driving don't mix: Inductive generalization in infancy. *Cognition*, *59*, 307–335.

McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, *38*, 339–346.

McCrink, K., & Wynn, K. (2004). Large-number addition and subtraction by 9-month-old infants. *Psychological Science*, *15*, 776–781.

Mestres-Missé, A., Trampel, R., Turner, R., & Kotz, S. A. (2017). Uncertainty and expectancy deviations require cortico-subcortical cooperation. *NeuroImage*, *144*, 23–34.

O'Reilly, J. X. (2013). Making predictions in a changing world — Inference, uncertainty, and learning. *Frontiers in Neuroscience*, *7*, 1–10.

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.

Pessiglione, M., Seymour, B., Flandin, G., Dolan, R., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behavior in humans. *Nature*, *442*, 1042–1045.

Piaget, J. (1970). Piaget's theory. In P. H. Mussen (Ed.), *Carmichael's manual of child psychology* (vol. 1, pp. 703–732). New York: Wiley.

Pica, P., Lemer, C., Izard, V., & Dehaene, S. (2004). Exact and approximate arithmetic in an Amazonian indigene group. *Science*, *306*, 499–503.

Regolin, L., & Vallortigara, G. (1995). Perception of partly occluded objects by young chicks. *Perception & Psychophysics*, *57*, 971–976.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*, 1926–1928.

Samuels, R. (2002). Nativism in cognitive science. *Mind & Language*, *17*(3), 233–265.

Santos, L. R. (2004). 'Core knowledges': A dissociation between spatiotemporal knowledge and contact-mechanics in a non-human primate? *Developmental Science*, *7*, 167–174.

Schönberg, T., Daw, N. D., Joel, D., & O'Doherty, J. P. (2007). Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. *The Journal of Neuroscience*, *27*, 12860–12867.

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.

Sim, Z., & Xu, F. (2017). Infants preferentially approach and explore the unexpected. *British Journal of Developmental Psychology*, *35*, 596–608.

Skerry, A. E., & Spelke, E. S. (2014). Preverbal infants identify emotional reactions that are incongruent with goal outcomes. *Cognition*, *130*, 204–216.

Skinner, B. F. (1957). *Verbal behavior*. New York: Appleton-Century-Crofts.

Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, *14*, 29–56.

Spelke, E. S. (1994). Initial knowledge: Six suggestions. *Cognition*, *50*, 431–445.

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.

Spelke, E. S., & Lee, S. A. (2012). Core systems of geometry in animal minds. *Philosophical Transactions of the Royal Society, B.*, *367*, 2784–2793.

Spelke, E. S., & Newport, E. L. (1998). Nativism, empiricism, and the development of knowledge. *Infant Behavior and Development*, *21*, 181–200.

Spelke, E. S., & Van de Walle, G. (1993). Perceiving and reasoning about objects: Insights from infants. In N. Eilan, R. McCarthy, & W. Brewer (Eds.), *Spatial representation* (pp. 132–161). Oxford, UK: Basil Blackwell.

Spencer, J. P., Blumberg, M. S., McMurray, B., Robinson, S. R., Samuelson, L. K., & Tomblin, J. B. (2009). Short arms and talking eggs: Why we should no longer abide the nativist-empiricist debate. *Child Development Perspectives*, *3*, 79–87.

Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*, 91–94.

Stahl, A. E., & Feigenson, L. (2017). Expectancy violations promote learning in young children. *Cognition*, *163*, 1–14.

Stahl, A. E., & Feigenson, L. (in preparation a). Surprise–induced exploratory play in infants: Does surprise generalize across exemplars?

Stahl, A. E., & Feigenson, L. (in preparation b). Impossibly special: Impossible but not improbable events boost children's learning.

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.

Swallow, K. M., & Jiang, Y. V. (2010). The Attentional Boost Effect: Transient increases in attention to one task enhance performance in a second task. *Cognition*, *115*, 118–132.

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*, 19156019159.

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.

Vallortigara, G. (2012). Core knowledge of object, number, and geometry: A comparative and neural approach. *Cognitive Neuropsychology*, *29*(1–2), 213–236.

Vallortigara, G., Regolin, L., Chiandetti, C., & Rugani, R. (2010). Rudiments of mind: Insights through the chick model on number and space cognition in animals. *Comparative Cognition and Behavior Reviews*, *5*, 78–99.

Vallortigara, G., Regolin, L., Rigoni, M., & Zanforlin, M. (1998). Delayed search for a concealed imprinted object in the domestic chick. *Animal Cognition*, *1*, 17–24.

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.

Waelti, P., Dickinson, A., & Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, *412*, 6842.

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.

Welder, A. N., & Graham, S. A. (2001). The influence of shape similarity and shared labels on infants' inductive inferences about nonobvious object properties. *Child Development*, *72*, 1653–1673.

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.

Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, *69*, 1–34.

Woodward, A. L., Phillips, A., & Spelke, E. S. (1993). Infants' expectations about the motions of inanimate vs. animate objects. In *Proceedings of the Cognitive Science Society* (pp. 1087–1091). Hillsdale, NJ: Erlbaum.

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.

Xu, F., & Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science*, *22*, 28–32.

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*, 245.

Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*, 681–692.