Cognition 163 (2017) 1–14

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Original Articles

Expectancy violations promote learning in young children

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

[a] The College of New Jersey, 2000 Pennington Road, Ewing, NJ 08628, United States
[b] Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, United States

ARTICLE INFO

Article history:
Received 11 July 2016
Revised 8 February 2017
Accepted 15 February 2017
Available online 27 February 2017

Keywords:
Children
Word learning
Surprise
Expectations
Object knowledge

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.

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;

* Corresponding author at: Psychology Department, The College of New Jersey, 2000 Pennington Road, Ewing, NJ 08628, United States.
E-mail addresses: stahla@tcnj.edu (A.E. Stahl), feigenson@jhu.edu (L. Feigenson).

http://dx.doi.org/10.1016/j.cognition.2017.02.008
0010-0277/

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., Kundey, 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 et al. (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