



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

Another Look at Looking Time: Surprise as Rational Statistical Inference

Zi L. Sim, Fei Xu

Department of Psychology, University of California, Berkeley

Received 31 August 2016; received in revised form 4 October 2018; accepted 4 October 2018

Abstract

Decades of developmental research have capitalized on the fact that infants are surprised (i.e., look longer) at some events but not others. Differences in looking time have been considered to be a reflection of perceptual discrimination, or a reaction toward witnessing a violation of prior expectations. Here, we provide an overview of a new perspective on infant surprise that examines the underlying cognitive processes that drive this response. We suggest that looking time may reflect sophisticated statistical inference, and we review empirical evidence and computational modeling results from several recent studies to support this conjecture (Kidd, Piantadosi, & Aslin, 2012; Piantadosi, Kidd, & Aslin, 2014; Sim, Griffiths, & Xu, 2018; Sim & Xu, 2017; Téglás et al., 2011). We also discuss how our view relates to other new developmental research on surprise and learning (Stahl & Feigenson, 2015, 2017) and outline some suggestions for future research.

Keywords: Looking time; Surprise; Rational statistical inference; Violation-of-expectation method; Infant cognition; Cognitive development

“What percepts do you have? What knowledge are you born with? How do you learn about the world?” As students of cognitive development, we would give the world to be able to pose such intricate questions to young infants, but they simply do not have the verbal capacity, among many other necessary tools, to satisfy our curiosity about their

Correspondence should be sent to Zi L. Sim or Fei Xu, 2121 Berkeley Way, Department of Psychology, University of California, Berkeley, CA 94720. E-mails: zi@berkeley.edu or fei_xu@berkeley.edu

inner world. What infants do have, however, is the ability to explore their environment with their *eyes*, right from day one. And decades of developmental research have capitalized on such looking behavior to probe the inner workings of an infant's mind.

Initially, looking time methods were used to evaluate infants' perceptual capacities, such as whether infants can perceive forms or color (Berlyne, 1958; Fantz, 1961; Hershenson, 1964; Spears, 1966). Visual habituation methods, in which infants are repeatedly presented with a set of stimuli until their looking to the stimuli decreases to some criterion, also revealed that infants would dishabituate, that is, increase in looking time, upon being shown a novel stimulus (Fantz, 1964; see Colombo & Mitchell, 2009 and Aslin, 2007 for reviews).

These methods paved the way for the violation-of-expectation paradigm (VOE), which in the last three decades has become a staple in infant labs for characterizing the initial state of an infant. With the advent of this tool, researchers have demonstrated that the human infant is highly competent — we now know that infants possess knowledge in a variety of domains very early on, such as physical knowledge (e.g., Baillargeon, 2008; Baillargeon, Li, Gertner, & Wu, 2011; Baillargeon, Spelke, & Wasserman, 1985; Spelke, Breinlinger, Macomber, & Jacobson, 1992), numerical knowledge (e.g., Coubart, Izard, Spelke, Marie, & Streri, 2014; McCrink & Wynn, 2015; Xu & Spelke, 2000), statistical and probabilistic intuitions (e.g., Fiser & Aslin, 2002; Kirkham, Slemmer, & Johnson, 2002; Téglás, Girotto, Gonzalez, & Bonatti, 2007; Xu & Garcia, 2008), and theory of mind (e.g., Gergely, Nádasdy, Csibra, & Bíró, 1995; Onishi & Baillargeon, 2005).

The underlying logic of the VOE method is as such: By conducting looking time experiments together with adequate controls, infants' longer looking at an unexpected than an expected event indicates that, first, infants possess the expectation being examined; second, they have detected a violation of that expectation; and third, they are *surprised* by the violation, measured as increased attention or interest (Wang, Baillargeon, & Brueckner, 2004; but see Jackson & Sirois, 2009; Sirois & Jackson, 2007). This method has been traditionally used as a qualitative measure, in the sense that there is typically only a binary set of events/stimuli, and in the ideal situation, infants look longer at one but not the other, suggesting surprise.

In this paper, we provide an overview of a new way of thinking about looking time responses, focusing on examining the underlying cognitive processes. We argue that looking time may reflect sophisticated statistical inference, captured by Bayesian probabilistic models. We will provide preliminary empirical and computational modeling evidence in support of this conjecture. We will also discuss other recent developmental work on surprise and learning, and outline some suggestions for future research.

In the last few years, several studies have modeled patterns of looking times in a quantitative manner and argued that we can use a combination of empirical research and computational models to best understand the nature of looking time (Kidd et al., 2012; Piantadosi et al., 2014; Sim & Xu, 2017; Sim et al., 2018; Téglás et al., 2011). These studies follow from research in the last decade or so that has demonstrated that the precision afforded by computational models may give us a great additional tool for studying developmental phenomena (e.g., Gopnik & Wellman, 2012; Gopnik et al., 2004; Perfors,

Tenenbaum, Griffiths, & Xu, 2011; Xu, 2016; Xu & Kushnir, 2012, 2013; Xu & Tenenbaum, 2007).

In Téglás et al. (2011), researchers presented infants with dynamic displays of three objects of one type (e.g., blue objects) and one object of another type (e.g., yellow objects) bouncing inside a container. They recorded infant looking times for the outcomes of 12 different events that were generated using a combination of three factors: (a) how long the container was occluded for before an object exited, (b) whether a majority or minority object exited, and (c) the physical arrangement of the four objects just before occlusion. The results showed that the average looking times of 12-month-old infants were systematically related to the predictions of an ideal Bayesian learner computing $1 - P(\text{outcome})$, where $P(\text{outcome})$ was the probability of observing a particular object type (i.e., a blue object or a yellow object) exiting the container. This model explained 88% of the variance in infants' looking times to the 12 different events. Besides demonstrating that infants can integrate numerosity, spatial, and temporal cues to form rational expectations about novel events, the findings also indicated that patterns of looking time may be driven by a computation of event probabilities. Infants were surprised by low-probability events, looking longer at these events compared to high-probability events.

Kidd et al. (2012) likewise found this relationship between looking behavior and event probability for 8-month-old infants. Their results indicated that infants' probability of *looking away* during a visual sequence was predicted by the *complexity* of the sequence, formalized as the *negative log probability* of observing a specific event given the distribution of previously observed events. More specifically, infants tended to look away when an observed event had very low complexity (i.e., very high in probability given previously observed events) or very high complexity (i.e., very low in probability given previously observed events). Even more striking, the modeling results held up within individual infants and was not an artifact of averaging infants with different types of behavior (Piantadosi et al., 2014).

Our recent work builds upon these earlier results indicating that infants' looking times are well predicted by the probability of observed events (Kidd et al., 2012; Piantadosi et al., 2014; Téglás et al., 2011). We hypothesized that the cognitive processes that drive infants' looking times may be more sophisticated, in that infants may have the ability to go beyond considering the mere probability of observed events. Specifically, we hypothesized that the level of surprise shown by infants, as measured by their looking times, may be predicted by infants considering alternative hypotheses that could account for the observed data.

Under this account, observed events are considered surprising not simply because they are low in probability, but because their occurrence is more consistent with a set of alternative hypotheses being true than that for the original hypothesis. Within a context of die rolls, a sequence "1, 1, 1, 1" has the same probability as a sequence "2, 4, 3, 6," but we consider the former to be far more surprising because it provides more support for a set of alternative hypotheses, such as the die is loaded, or the die has the number "1" on all its faces, and so on, as compared to our original hypothesis, that the die is fair. This proposal follows a Bayesian account of the sense of coincidence in adults (Griffiths &

Tenenbaum, 2007) and is congruent with principles of Bayesian learning (Griffiths, Charter, Kemp, Perfors, & Tenenbaum, 2010; Tenenbaum & Griffiths, 2001; Tenenbaum, Kemp, Griffiths, & Goodman, 2011): *The ideal learner begins the learning process by assessing the fit between the observed evidence and her current highest-probability hypothesis and compares it to the fit provided by a set of lower-probability hypotheses.*

To test this, we used a combination of behavioral experiments and computational modeling. In a series of experiments, we familiarized 8-month-old infants to a population box containing six different-colored balls. An experimenter then tossed out different sequences of balls from the box (sampling with replacement). For example, infants observed sequences of balls such as “yellow, yellow, yellow, yellow” or “orange, green, blue, yellow.” Looking times for the various sequences were measured. Using this method, infants were presented with a total of 10 different sequences. These looking times were then compared to a Bayesian model that posits that infants may implicitly evaluate the support that an observed sequence provides for a set of alternative theories as compared to the currently favored theory, which is termed the *likelihood ratio*. An event with a likelihood ratio greater than 1 indicates that the event would be much better accounted for by alternative theories, and thus generates longer looking times in infants during VOE experiments. Therefore, under this account, surprise increases as likelihood ratio increases. We found that the Bayesian model provided a high quality of model fit, performing far better than several alternative models that were tested (see Supplementary Material). This finding suggests that when infants observe events in their environment, they reason in a manner that is consistent with evaluating those events according to how well they support different hypotheses. Observing a sequence of balls “yellow, yellow, yellow, yellow” prompts long looking times in infants because it better supports a set of alternative theories (e.g., the box has a hidden compartment containing many yellow balls; the yellow balls are heavier therefore more likely to fall out, etc.) as compared to the currently favored theory (i.e., the balls are being randomly sampled from the box). On the other hand, observing a sequence of balls “orange, green, blue, yellow” prompts short looking times in infants because the likelihood ratio is small—the observed sequence provides far more support for the currently favored theory than the set of alternative theories. This form of inferential reasoning is an important first step in enabling the infant to move toward theory building and theory revision, allowing her to form the best generative model for the data observed in the world (Sim & Xu, 2017; Sim, Griffiths, & Xu, unpublished data).

In another study, we replicated a part of these results with 13-month-old infants, namely that infants would look longer at a uniform sequence of balls (e.g., yellow, yellow, yellow) being tossed out of a box containing six different-colored balls, as compared to a variable sequence of balls being tossed out (e.g., orange, green, blue, yellow). We then used an exploration measure to investigate the downstream consequences of observing a surprising event. We presented 13-month-old infants with two boxes, each containing six different-colored balls. From one of the boxes, infants observed a uniform sequence of balls being tossed out, while from the other box, infants observed a variable sequence of balls being tossed out. When these boxes were later presented to the infants, we found that infants preferentially approached and explored the source of the surprising

event, that is, the box that had generated the uniform sequence of balls. These results indicate that infants spontaneously explore sources that violate their expectations, potentially providing themselves with new learning opportunities (Sim & Xu, 2017).

Overall, recent work in our laboratory demonstrates that infants' surprise response, as indicated by looking times, is driven by an evaluation of how well observed data support different hypotheses, and that such an evaluation has downstream consequences on infants' exploration, potentially influencing future learning.

So far we have provided a computational level of analysis for how an ideal learner should evaluate evidence against alternative hypotheses. By using a fairly abstract and domain-general task of random draws from an urn, we are able to provide good quantitative fits between the empirical results and a Bayesian probabilistic model. This line of reasoning opens up many new research questions, most important, how do learners generate alternative hypotheses and revise their beliefs?

As far as we know, no empirical studies with infants have tried to answer this question directly. We see a number of possibilities in how young learners may generate alternative hypotheses in order to revise their beliefs. First, infants may decide, especially with more data, that it was just a blip in the system — that is, extraneous factors produce the surprising data that are observed, and no belief revision is necessary. In our example, if an infant observes over and over again that a sequence of different-colored balls is randomly drawn from the box and occasionally a sequence of same-colored balls is drawn, then no revision is needed in their understanding of random events and probability. Second, as more evidence accumulates, infants may entertain the possibility that the urn is rigged, so an extra variable is needed in order to explain the observed data. This may be the beginning of minor belief revisions. Third, with more evidence, infants may decide to overhaul their whole theory of probabilistic reasoning, resulting in genuine conceptual change. It is unlikely that one piece of surprising evidence will change how infants reason about random events, but if more surprising data accumulate over time, infants may begin to entertain alternative hypotheses that could result in major belief revisions.

Alternatively, as Maguire, Moser, Maguire, and Keane (this volume) suggest, learners may use a heuristic such as *randomness deficiency* to detect anomalous events. That is, people are surprised if they find patterns when their current model predicts only random noise. Furthermore, this conception of surprise may be formalized using Algorithmic Information Theory, and there is empirical evidence that people's behavior accords well with this formal measure. The Maguire et al. account may provide a different interpretation of the infant looking time results we have discussed above (Sim et al.): Infants may be surprised at a sequence such as yellow, yellow, yellow, yellow compared with a sequence such as yellow, blue, green, purple because the former suggests non-randomness. Note that this view does not posit that the infant entertains any alternative hypotheses. The empirical evidence to date does not distinguish between these two accounts; these different theoretical perspectives offer a fertile ground for more fine-grained investigations of surprise in both adults and infants in future research. One interesting possibility is that a young learner may detect anomalies based on the randomness deficiency heuristic — this is the first step. Then the learner needs to go further in deciding whether the

anomaly warrants belief revision by considering alternative hypotheses and how much support is provided by this new piece of evidence.

The studies reviewed above all point in the same direction, that looking times reflect rational statistical inference (not merely a binary response that researchers can use as a reliable dependent measure), and they all hint at the possibility that surprise, construed as such, will provide opportunities for future learning. However, none of these studies have provided direct evidence for downstream learning after measuring surprise. Much to our delight (and surprise!), Stahl and Feigenson (2015, 2017) have developed new, ingenious methods to investigate what infants can learn after a surprise reaction, tackling these questions head on.

Stahl and Feigenson (2015) start with the observation that an infant's environment is highly complex. The surprise engendered by detecting an expectation violation helps infants to clarify their learning space, enabling them to know which aspects of their world to attend to and to learn from. In four experiments, Stahl and Feigenson (2015) demonstrated that after witnessing a violation of expectation, such as seeing a toy car pass through a solid wall, 11-month-old infants showed better learning of a hidden auditory property of the violation object. They also found better learning for the violation object but not for a novel object presented after the violation event, suggesting that the enhanced learning was not due to a general increase in attention or a preference for novelty. Finally, they also showed that these infants were more interested in playing with the violation object as compared to a novel object, and they engaged in behaviors that were directly related to the specific violation event that they observed. For example, the infants spent more time banging the violation object instead of dropping the violation object after observing a solidity violation, and this pattern was reversed after infants observed a support violation. These results suggest that infants recognize violation events to be a signal for special learning opportunities. The differential learning rate that infants showed after observing unexpected and expected events is very impressive. As the authors suggest, "expectancy violations offer a wedge into the problem of what to learn" (Stahl & Feigenson, 2015, p. 91).

Stahl and Feigenson (2017) found parallel results with 3- to 6-year-old children in a different domain, namely word learning. Using similar methods, their results showed that children learned the referents of novel nouns and verbs after witnessing a violation (e.g., spatiotemporal continuity), and they did not learn the new words after witnessing a novel event that did not violate any principles of intuitive physics. Furthermore, the enhanced learning effects were specific to the objects involved in the violation events, and not the result of general arousal.

These are truly innovative and elegant experiments that open up new ways of thinking about surprise and learning in cognitive development. One characterization of these findings is the following: Increased focused attention (i.e., longer looking time that reflects surprise) enhances infants' subsequent learning of a non-obvious property (e.g., internal sound, word) of the object that has violated a principle of physical reasoning. As Stahl and Feigenson pointed out, many questions remain about how domain-general these findings are, and whether less or more extreme violations would also result in enhanced learning.

In the last section of this review, we discuss three implications of these new lines of research on looking time, surprise, and learning in infants, and suggest new directions for future theoretical and empirical work.

One important question raised by all the recent research reviewed above is the type of learning that occurs after infants observe a surprising event. In Stahl and Feigenson (2015), infants showed enhanced learning for a hidden auditory property of an object that had participated in a violation event, and in Stahl and Feigenson (2017), young children showed enhanced learning of novel nouns and verbs. The learning that occurs centers on the entity that had violated their expectation; the infant is now attuned to any new, arbitrary piece of information that is available about the specific violation object. Within the rational statistical inference framework, the type of learning that occurs after infants witness a violation of prior expectations may center on belief revision and conceptual change. Since we assume that infants are building causal, generative models of the world of a particular domain, that is, intuitive theories (Carey, 2009; Gopnik & Meltzoff, 1997; Gopnik & Wellman, 2012; Tenenbaum et al., 2011; Xu & Kushnir, 2012, 2013), surprising data may serve as a driving force for theory change. These two approaches, however, may complement each other. The striking findings of Stahl and Feigenson (2015, 2017) — that infants and young children experience enhanced learning of a somewhat arbitrary piece of information associated with the violation object — may be the first step in belief revision and theory change. So far enhanced learning has only been demonstrated for a particular violation object, but infants may subsequently believe that all objects that make the same novel sounds will produce a solidity violation. By connecting different pieces of knowledge and inferences, infants may begin to revise their intuitive theory about the physical world.

A second important open question is the nature of the surprise response. Past studies have largely assumed that surprise (reflected in longer looking time) is binary — either infants are surprised or they are not. The quantitative computational modeling efforts suggest otherwise — surprise is graded, and the gradedness is indicative of the level of surprise, predictability, and the plausibility of alternative hypotheses (Kidd et al., 2012; Piantadosi et al., 2014; Téglás et al., 2011; Sim et al., unpublished data). In particular, Kidd et al. (2012) suggest that extremely surprising and unpredictable events result in the infants turning away, and giving up since they may view these as truly random and an unlikely source of learning. As Stahl and Feigenson (2017) pointed out, their focused-attention-enhances-learning findings may be extended and integrated into a rational statistical inference framework of Kidd et al. (2012). Suppose the object that violates solidity also floats in mid-air and undergoes unexpected featural change, would infants still show enhanced learning, or would they give up in the face of multiple violations of core beliefs?

The third important open question is how surprise may help infants gauge what to learn in a cluttered environment. Stahl and Feigenson (2015, 2017) have demonstrated elegantly that violations of core knowledge principles (e.g., solidity, continuity) enhanced learning in infants and young children. In real life, many surprising events happen and infants (indeed all learners) may be surprised to different degrees by different events, for example, being surprised that someone radically changed her hairstyle versus being surprised that the sun

did not rise on July 13, 2018. In other words, not all surprises are created equal; only some surprises are worthy of further investigation and provide potential learning opportunities. Our suggestion is that the rational statistical inference framework may help learners decide which surprising events may be more valuable for learning than others. In our work, surprising events are considered more valuable if they provide support for alternative hypotheses, relative to how much support they provide for the currently held hypothesis. In Kidd et al. (2012), surprising events are considered more valuable if they are at a mid-level of complexity and predictability. Both of these suggestions, however, have been stated at an abstract algorithmic level. The challenge for future research will be to generate empirical studies that investigate these possibilities more directly.

For the past few decades, looking time has been used to index and characterize infants' *pre-existing* knowledge. In contrast, the two new lines of work reviewed here examine the role surprise plays in *advancing* an infant's cognitive development, shedding much-needed light on a question that has remained largely ignored in the infant cognition literature. These new approaches provide empirical support for the view that surprise has consequences for early learning, enabling the infant learner to construct a more accurate model of the world. The empirical findings and computational modeling results complement work on surprise and learning in adults (e.g., Munnich & Ranney, this volume), and work on prediction error in non-human animals (e.g., Holland & Gallagher, 2006). The research reviewed here represents exciting new directions for the study of surprise and learning in cognitive development. We believe that further theoretical, computational, and empirical work will generate new insights and help us develop a full account of why infants are surprised and how they learn.

Acknowledgments

We thank Tom Griffiths and Noah Goodman for helpful discussion.

References

- Aslin, R. N. (2007). What's in a look? *Developmental Science*, *10*(1), 48–53. <https://doi.org/10.1111/j.1467-7687.2007.00563.x>
- Baillargeon, R. (2008). Innate ideas revisited: For a principle of persistence in infants' physical reasoning. *Perspectives on Psychological Science*, *3*(1)2–13. <https://doi.org/10.1111/j.1745-6916.2008.00056.x>
- Baillargeon, R., Li, J., Gertner, Y., & Wu, D. (2011). How do infants reason about physical events? *The Wiley-Blackwell Handbook of Childhood Cognitive Development*, *2*, 11–48.
- Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition*, *20*(3), 191–208.
- Berlyne, D. E. (1958). The influence of the albedo and complexity of stimuli on visual fixation in the human infant. *British Journal of Psychology*, *49*(4), 315–318. <https://doi.org/10.1111/j.2044-8295.1958.tb00669.x>
- Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.

- Colombo, J., & Mitchell, D. W. (2009). Infant visual habituation. *Neurobiology of Learning and Memory*, 92(2)225–234. <https://doi.org/10.1016/j.nlm.2008.06.002>
- Coubart, A., Izard, V., Spelke, E. S., Marie, J., & Streri, A. (2014). Dissociation between small and large numerosities in newborn infants. *Developmental Science*, 17(1), 11–22.
- Fantz, R. L. (1961). The origin of form perception. *Scientific American*, 204(5), 66–72. <https://doi.org/10.1038/scientificamerican0561-66>
- Fantz, R. L. (1964). Visual experience in infants: Decreased attention to familiar patterns relative to novel ones. *Science*, 146(3644), 668–670.
- Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences of the United States of America*, 99(24), 15822–15826. <https://doi.org/10.1073/pnas.232472899>
- Gergely, G., Nádasdy, Z., Csibra, G., & Bíró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, 56(2), 165–193.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, 111(1), 3–32.
- Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. Cambridge, MA: MIT Press.
- Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, 138(6), 1085–1108.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences*, 14(8), 357–364. <https://doi.org/10.1016/j.tics.2010.05.004>
- Griffiths, T. L., & Tenenbaum, J. B. (2007). From mere coincidences to meaningful discoveries. *Cognition*, 103(2), 180–226. <https://doi.org/10.1016/j.cognition.2006.03.004>
- Hershenson, M. (1964). Visual discrimination in the human newborn. *Journal of Comparative and Physiological Psychology*, 58(2), 270–276. <https://doi.org/10.1037/h0047340>
- Holland, P., & Gallagher, M. (2006). Different roles for amygdala central nucleus and substantia innominata in the surprise-induced enhancement of learning. *Journal of Neuroscience*, 26, 3791–3797.
- Jackson, I., & Sirois, S. (2009). Infant cognition: Going full factorial with pupil dilation. *Developmental Science*, 12(4), 670–679. <https://doi.org/10.1111/j.1467-7687.2008.00805.x>
- Kidd, C., Piantadosi, S. T., & 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(5), e36399. <https://doi.org/10.1371/journal.pone.0036399>
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, 83(2), B35–B42.
- McCrink, K., & Wynn, K. (2015). Large-number addition and subtraction by 9-month-old infants. *Psychological Science*, 15(11), 9–14.
- Onishi, K. H., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, 308(5719), 255–258. <https://doi.org/10.1126/science.1107621>
- Perfors, A., Tenenbaum, J. B., Griffiths, T., & Xu, F. (2011). A tutorial introduction to Bayesian models of cognitive development. *Cognition*, 120, 302–321.
- Piantadosi, S. T., Kidd, C., & Aslin, R. (2014). Rich analysis and rational models: Inferring individual behavior from infant looking data. *Developmental Science*, 17(3), 321–337. <https://doi.org/10.1111/desc.12083>
- Sim, Z., Griffiths, T. L., & Xu, F. (2018) Inferential reasoning in 8-month-old infants: An empirical and computational investigation. Manuscript under revision.
- Sim, Z., & Xu, F. (2017). Infants preferentially approach and explore the unexpected. *British Journal of Developmental Psychology*, 35, 596–608.
- Sirois, S., & Jackson, I. (2007). Social cognition in infancy: A critical review of research on higher order abilities. *European Journal of Developmental Psychology*, 4(1), 46–64. <https://doi.org/10.1080/17405620601047053>

- Spears, W. C. (1966). Visual preference in the four-month old infant. *Psychonomic Science*, 4(6), 237–238.
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, 99(4), 605–632.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348(91), 91–94.
- Stahl, A. E., & Feigenson, L. (2017). Expectancy violations promote learning in young children. *Cognition*, 163, 1–14.
- 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 of the United States of America*, 104(48), 19156–19159. <https://doi.org/10.1073/pnas.0700271104>.
- 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(6033), 1054–1059. <https://doi.org/10.1126/science.1196404>.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, 24, 629–640.
- 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. <https://doi.org/10.1126/science.1192788>.
- Wang, S.-H., Baillargeon, R., & Brueckner, L. (2004). Young infants' reasoning about hidden objects: Evidence from violation-of-expectation tasks with test trials only. *Cognition*, 93(3), 167–198. <https://doi.org/10.1016/j.cognition.2003.09.012>.
- Xu, F. (2016). Preliminary thoughts on a rational constructivist approach to cognitive development. In D. Barner & A. Baron (Eds.), *Core knowledge and conceptual change*. Oxford, UK: Oxford University Press.
- 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(13), 5012–5015. <https://doi.org/10.1073/pnas.0704450105>.
- Xu, F., & Kushnir, T. (2012). Rational constructivism in cognitive development. In F. Xu & T. Kushnir (Eds.), *Advances in child development and behavior*, Vol. 43. Waltham, MA: Academic Press.
- Xu, F., & Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science*, 22, 28–32.
- Xu, F., & Spelke, E. S. (2000). Large number discrimination in 6-month-old infants. *Cognition*, 74, B1–B11.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114, 245–272.

Supporting Information

Additional Supporting Information may be found online in the supporting information tab for this article:

Appendix S1. Summary of infant experiments and model results

Appendix

To characterize the underlying processes that give rise to different levels of looking, we formalize two accounts using probabilistic models: 1) a mere probability model and 2) a Bayesian inferential model. As the assumptions of each model are outlined explicitly, we can calculate the predictions made by these models and thus determine which model best accounts for infant looking times. For these models, h is a hypothesis, d is a specific sequence of balls observed, k is the number of uniquely colored balls in the large population box, and N is the number of independent draws from the box.

Mere Probability Model

First, infants could simply consider low-probability events to be surprising given their current expectations. This account is intuitive, as we tend to consider surprising events as having a low probability of occurrence. Furthermore, previous studies have demonstrated that infants as young as 6 months do look longer at low-probability events as compared to high-probability events (Denison, Reed, & Xu, 2013; Téglás et al., 2011; Téglás, Girotto, Gonzalez, & Bonatti, 2007; Xu & Garcia, 2008). This model thus predicts that infants' looking times will be closely related to the mere probability, $P(d|h_{random})$, of individual sequences of balls, d , given h_{random} , the hypothesis that the balls are being tossed out randomly from the box. We use the negative log probability of these events, as this measure quantifies how surprising it is to see a particular outcome (Kidd et al., 2012). The probability of the sequences is

$$P(d | h_{random}) = \left(\frac{1}{k}\right)^N,$$

(1)

as the probability of each color is inversely proportional to the number of colors k and the draws are assumed to be independent.

Bayesian Model

Infants may evaluate the evidence that a certain sequence, d , provides for the alternative theory, as compared to the currently favored theory, $P(d|h_{\text{alternative}})/P(d|h_{\text{current}})$ (the likelihood ratio; Griffiths & Tenenbaum, 2007), in a process that is consistent with principles of Bayesian learning (Griffiths et al., 2010; Tenenbaum & Griffiths, 2001; Tenenbaum et al., 2011). For our experiments, h_{current} refers to the currently favored hypothesis that sampling is random (i.e., h_{random}), as the experimenter appears to have no control over the outcome of the tosses from the box. The alternative hypothesis, $h_{\text{alternative}}$, is thus that sampling is not random (i.e., h_{biased}).

Assuming each biased distribution is equally likely, a derivation yields the likelihood ratio

$$\frac{P(d | h_{\text{alternative}})}{P(d | h_{\text{current}})} = \frac{\binom{k^N}{k-1} \left(\prod_{i=1}^k n_i! \right)}{(N+k-1)!} .$$

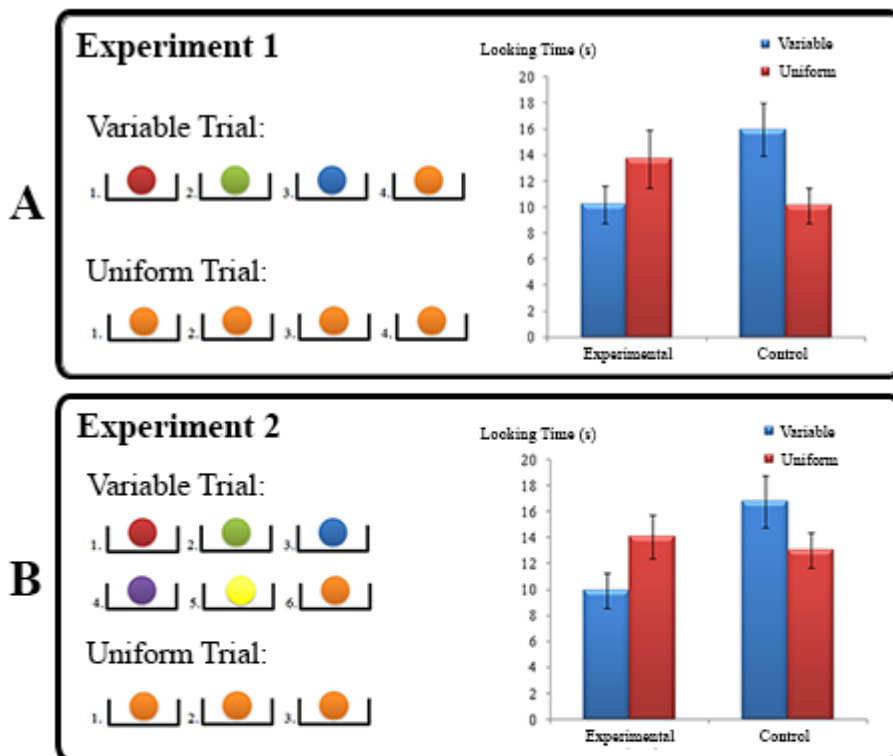
(5)

A likelihood ratio larger than 1 indicates that there is a greater probability of observing event d under the alternative theory than the currently favored theory. Under this model, an event becomes more surprising as the likelihood ratio increases. For example, the probability of observing a ball pass through a wood panel is virtually zero given our solidity expectations: objects move only on unobstructed paths (Spelke et al., 1992). When such an event is observed (typically only in an infant VOE experiment), the likelihood ratio is practically infinite since the event is much better accounted for by alternative hypotheses. As such, a solidity violation would produce longer looking times in infants.

Model Predictions

These models make different predictions, which we test in two looking time experiments. In Experiment 1 (Figure 1A), we considered two sequences of balls randomly

sampled from a population of 6 different colored balls: a uniform sequence (e.g. orange, orange, orange, orange), and a variable sequence (e.g. red, green, blue, orange). As the two sequences are equal in probability given random sampling, the mere probability model predicts that looking times should not be different for the two sequences. However, the Bayesian model predicts otherwise: infants should look longer at the uniform sequence because it has a higher likelihood ratio.



In Experiment 2 (Figure 1B), we considered sequences of different lengths, e.g. a uniform sequence of 3 balls, and a variable sequence of 6 balls. The mere probability model predicts that infants should look longer at the variable sequence as it is a lower probability event. In contrast, the Bayesian model predicts that infants should look longer at the uniform sequence, as it continues to have a higher likelihood ratio.

Results from Experiments

Experiment 1. Forty 8-month-old infants were tested, half in the Experimental condition and half in the Control condition. In the Experimental Condition, on each

Familiarization trial, the box containing 6 different-colored balls was shown to the infant. The Test phase consisted of two test trials, a Uniform trial and a Variable trial (Figure 1A). The Control condition was the same as the Experimental condition, except that on the test trials, the balls were pulled out of the experimenter's pocket instead of the box. Looking times for the test trials were analyzed using a 2 x 2 repeated-measures ANOVA with Condition (Experimental vs. Control) as the between-subjects factor and Trial Type (Uniform vs. Variable) as the within-subjects factor. There were no main effects. There was a significant interaction between Condition and Trial Type, $F(1, 38) = 11.58, p = .002, \eta_p^2 = .23$. In the Experimental condition, infants looked significantly longer in the Uniform trial ($M = 13.68s, SD = 9.87$) than the Variable trial ($M = 10.22s, SD = 6.35$); in the Control condition infants looked significantly longer in the Variable trial ($M = 15.96s, SD = 9.02$) than the Uniform trial ($M = 10.14s, SD = 6.01$).

Experiment 2. Forty 8-month-old infants were tested, half in the Experimental condition and half in the Control condition. Design and procedure were similar to those of Experiment 1, except that the Uniform sequence now consisted of 3 balls of the same color, and the Variable sequence consisted of 6 balls of different colors (Fig. 1B). A repeated-measures ANOVA with Condition (Experimental vs. Control) as the between-subjects factor and Trial Type (Uniform vs. Variable) as the within-subjects factor was performed on the obtained looking times. There were no main effects. There was a significant interaction between Condition and Trial Type, $F(1, 38) = 12.16, p = .001, \eta_p^2 = .24$. Infants in the Experimental condition looked significantly longer in the Uniform trial ($M = 14.07s, SD = 7.55$) than the Variable trial ($M = 9.91s, SD = 6.07$); infants in the Control condition looked significantly longer in the Variable trial ($M = 16.76s, SD = 8.88$) than the Uniform trial ($M = 13.06s, SD = 6.02$).

Quantitative Analysis

Next we correlated infants' looking times from Experiments 1 and 2 with the model predictions. The Bayesian model provided a high quality of model fit ($r = .96$, $df = 8$, $p = .0002$). The predictions correlated well with the looking times, unlike that of the mere probability model ($r = .53$, $df = 8$, $p = .32$).

References

- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: exploring representations and inductive biases. *Trends in Cognitive Sciences*, *14*(8), 357–364. <http://doi.org/10.1016/j.tics.2010.05.004>
- Griffiths, T. L., & Tenenbaum, J. B. (2007). From mere coincidences to meaningful discoveries. *Cognition*, *103*(2), 180–226. <http://doi.org/10.1016/j.cognition.2006.03.004>
- Kidd, C., Piantadosi, S. T., & 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*(5), e36399. <http://doi.org/10.1371/journal.pone.0036399>
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, *99*(4), 605–632.
- 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 of the United States of America*, *104*(48), 19156–9. <http://doi.org/10.1073/pnas.0700271104>
- 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 (New York, N.Y.)*, *332*(6033), 1054–9. <http://doi.org/10.1126/science.1196404>
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, *24*, 629–640.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science*, *331*, 1279–85.

<http://doi.org/10.1126/science.1192788>

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*(13), 5012–5.

<http://doi.org/10.1073/pnas.0704450105>