RESEARCH ARTICLES

COGNITIVE DEVELOPMENT

Observing the unexpected enhances infants' learning and exploration

Aimee E. Stahl* and Lisa Feigenson

Given the overwhelming quantity of information available from the environment, how do young learners know what to learn about and what to ignore? We found that 11-month-old infants (N = 110) used violations of prior expectations as special opportunities for learning. The infants were shown events that violated expectations about object behavior or events that were nearly identical but did not violate expectations. The sight of an object that violated expectations enhanced learning and promoted information-seeking behaviors; specifically, infants learned more effectively about objects that committed violations, explored those objects more, and engaged in hypothesis-testing behaviors that reflected the particular kind of violation seen. Thus, early in life, expectancy violations offer a wedge into the problem of what to learn.

Humans are capable of remarkable achievements, from learning a language to designing skyscrapers and mastering calculus. These achievements would be impossible without learning. Yet, as many theorists have noted, the problems of when learning should occur, and what should be learned at all, are highly underdetermined (*1*, *2*). In an environment that is dynamic and complex, how can a learner know which aspects of the world to attend to and learn from, and which to ignore? Without a filter for determining when and what to learn, or a teacher to provide guidance (*3*), information overload can, in practice, make learning impossible.

At the same time, some aspects of the world appear to be represented even prior to learning. These cognitive primitives, sometimes collectively called "core knowledge," can be observed in newborn creatures (*4*, *5*) and emerge across diverse rearing conditions (*6*) and cultures (*7*). But far from obviating the need for learning, core knowledge may be a foundational understanding from which learning begins. One way this could be so is if core knowledge offers a wedge into the hard problems of knowing when and what to learn. If a learner has a basic repertoire of core expectations about the world, then detecting a violation of these expectations—a conflict between what was predicted and what is observed—might signal a special opportunity for learning.

A clue that core knowledge may in fact guide early learning comes from infants' behavior

Department of Psychological & Brain Sciences, Johns Hopkins University, Baltimore, MD 21218, USA.
*Corresponding author. E-mail: astahl4@jhu.edu

in tests of preverbal cognition. Across hundreds of studies, infants respond when basic expectations are violated, including expectations generated by core knowledge (*8*). For example, infants look longer when a ball appears to pass through a wall than when it is stopped by the wall, suggesting a core understanding of object solidity (*9*), and they look longer when an object hidden in one location is revealed in a different location, suggesting a core understanding of object continuity (*10*). Seeing surprising events like these can trigger increases in infants' looking, as well as alterations in facial expression (*11*), pupil dilation (*12*), and changes in cerebral blood flow or brain electrical activity (*13*, *14*). These various responses have been taken to indicate the detection of a discrepancy between what was expected and what is observed, and have been documented across many knowledge domains. Infants detect violations when, for example, a hidden object vanishes (*15*), when 5 + 5 = 5 (*16*), and when a social entity approaches someone mean rather than someone nice (*17*). Responses to such surprising physical, numerical, and social events have been invaluable in efforts to characterize the roots of human cognition. Yet it remains unknown what purpose these surprise responses serve and what the cognitive consequences of experiencing an expectancy violation might be.

Here we tested the hypothesis that, early in life, violations of core expectations signal a special opportunity for learning. First we asked whether infants more effectively learn new information about objects that violate expectations than about objects that accord with expectations (experiments 1 to 3). Then we asked whether in-

fants preferentially seek information from objects that violated expectations, and whether their exploratory actions test plausible explanations for an observed violation (experiment 4).

Infants' learning about objects that violated expectations

In experiments 1 to 3, we showed infants an event whose outcome either was expected because it accorded with core knowledge of object behavior or was surprising because it violated core knowledge, using events modeled on those in many previous studies. Then we taught infants something new about the object that had participated in the event, and finally we measured how well they learned this new information.

Three aspects of our design were crucial. First, we ensured that events that violated core knowledge differed minimally from events that accorded with core knowledge, by perceptually matching the events in all respects except for their outcomes. Second, we ensured that any observed learning enhancement was caused by experiencing a violation of core expectations, rather than by longer perceptual exposure to objects that violated expectations, by matching the duration of infants' looking across outcome types. Third, we ensured that infants were learning something genuinely new by teaching them information that could not have been known beforehand and that could not have been acquired just by seeing the objects themselves (i.e., we taught infants an object's hidden property).

In experiment 1, 11-month-old infants saw an event that either accorded with or violated object solidity or spatiotemporal continuity, two core physical principles to which young infants have consistently shown sensitivity (*18*–*20*) (N = 40; movies S1 to S4). In the solidity event (Fig. 1A), infants saw an object (a toy car for half the infants; a ball for the other half) roll down a ramp and pass behind a screen. A solid wall, partially visible above the screen, clearly blocked the object's path. Infants then saw the screen removed to reveal either that the object had been stopped by the wall, thereby according with expectations about solidity (Knowledge-Consistent outcome, n = 10), or that the object appeared to have passed through the wall, thereby violating expectations about solidity (Knowledge-Violation outcome, n = 10). In the spatiotemporal continuity event (Fig. 1B), a separate group of infants saw two screens placed on an empty stage. The experimenter hid an object (a ball for half the infants; a block for the other half) behind the left screen, then lifted both screens to reveal either that the object was still behind the left screen, thereby according with expectations about continuity (Knowledge-Consistent outcome, n = 10), or that the object was now behind the right screen, thereby violating expectations about continuity (Knowledge-Violation outcome, n = 10).

Unlike previous studies designed to measure differences in infants' looking to expected versus

violation events, here we gave all infants the same limited visual exposure to the Knowledge-Consistent and Knowledge-Violation outcomes; all infants had just 10 s to encode the event outcome. A univariate analysis of variance (ANOVA), with looking time to the event outcome as the dependent variable and event type (Solidity or Continuity) and outcome type (Knowledge-Consistent or Knowledge-Violation) as fixed factors, showed no main effect of outcome type [$F(1,36) = 0.002$, $P = 0.96$] (18) (table S1), which was as predicted given the short encoding window in our design. Thus, any subsequent differences in learning cannot be attributed to longer perceptual exposure to the object in the Knowledge-Violation events.

Immediately after this 10-s exposure to the outcome of the Knowledge-Consistent or the Knowledge-Violation event, we taught all infants new information about the object in the event. The experimenter demonstrated that the object had a hidden auditory property (e.g., it squeaked) by moving it up and down while the sound played synchronously from a hidden central location for 12 s. Our dependent measure was infants' learning of this object-sound mapping. In the test trial, infants saw the target object from the preceding event and a new distractor object resting silently on the stage (baseline; 5 s). For half the infants, the ball was the target and either the car or the block was the distractor; this was reversed for the other half. Then the experimenter moved both objects up and down simultaneously while the previously taught sound (e.g., squeaking) played from a hidden central location (mapping test; 10 s). For each infant we calculated a learning score by determining the proportion of time that infants looked at the target object (relative to the new distractor object) during the baseline, then subtracting this value from the proportion of time they looked at the target object during the mapping test, when the taught sound played (table S1). If infants had successfully learned the object-sound mapping, they should increase the proportion of time they looked at the target object when the sound played;

such auditory-visual "matching" is the pattern typically observed in studies of infants' mapping abilities (21).

We found that infants' learning of the object-sound mapping depended on whether they had just seen a Knowledge-Consistent or a Knowledge-Violation event. A univariate ANOVA, with learning score as the dependent variable and event type (Solidity or Continuity) and outcome type (Knowledge-Consistent or Knowledge-Violation) as fixed factors, yielded only a significant main effect of outcome type [$F(1,36) = 10.691$, $P = 0.002$, partial $\eta^2 = 0.229$]. Infants' learning scores were significantly greater after Knowledge-Violation events than after Knowledge-Consistent events (Fig. 2A). We then compared infants' learning scores to chance (zero). Infants showed no evidence of learning after events consistent with object solidity [$t(9) = -1.088$, $P = 0.31$] or continuity [$t(9) = 1.62$, $P = 0.14$] but showed significant learning after violations to object solidity [$t(9) = 3.092$, $P = 0.01$] and spatiotemporal continuity [$t(9) = 3.715$, $P = 0.005$] (18) (Fig. 2A and table S1).

In experiment 2, we asked whether this pattern reflected actual learning or simply indicated greater attention to objects that had violated expectations. As in experiment 1, infants saw an object violate the core principle of solidity ($n = 10$) or continuity ($n = 10$) and were then taught that the object had a hidden auditory property (e.g., it squeaked). However, during the mapping test, we played an entirely novel sound (e.g., rattling). This time, infants did not increase their proportion of looking to the target object when the novel sound played after violations of either solidity [$t(9) = 1.453$, $P = 0.18$] or continuity [$t(9) = 0.036$, $P = 0.97$] (table S1). A univariate ANOVA, with learning score as the dependent variable and event type (Solidity or Continuity) and sound type (taught sound from the Knowledge-Violation condition of experiment 1 or novel sound from experiment 2) as fixed factors, yielded only a significant main effect of sound type. Infants' learning scores were significantly greater when the taught sound played in the mapping test (experiment 1) than

when the novel sound played (experiment 2) [$F(1,36) = 5.349$, $P = 0.03$, partial $\eta^2 = 0.129$] (18). This confirms that infants' performance in experiment 1 reflected successful learning of an object property, rather than heightened visual preference for an object that had violated expectations.

In experiment 3, we asked whether violations of expectation enhance learning specifically about objects that violated expectations, rather than about anything that might follow a violation. We showed infants ($n = 10$) the continuity violation from experiment 1, with an object (i.e., ball) hidden behind the left screen but revealed behind the right. After the object was revealed in the surprising location, the experimenter reached in with a new object (i.e., a block) and demonstrated that it had a hidden auditory property (e.g., it squeaked). We then measured infants' learning about this new object. As in experiment 1, we calculated learning scores by determining the proportion of time that infants looked at this new object (relative to a distractor object) during the silent baseline, then subtracting this value from the proportion of time they looked at it during the mapping test, when the taught sound played. We found that infants did not map the sound to the new object in the mapping test; their learning scores did not differ from chance [$t(9) = 0.074$, $P = 0.94$] (table S1). An independent-samples t test confirmed that this pattern differed significantly from that of experiment 1, in which infants were taught about the very object that had violated continuity [$t(18) = 2.126$, $P = 0.048$] (18). Hence, violations of expectation enhanced learning only for the object involved in the violation event, not for unrelated objects. Further, infants' failure to learn about the new object shows that the enhanced learning in experiment 1 was not due to general arousal or novelty. When taught about an object that was completely perceptually novel (because it had never been seen before) but did not violate any expectations, infants showed no evidence of learning.


Fig. 1. Knowledge-Consistent and Knowledge-Violation outcomes in experiments 1 to 4. (**A**) Solidity events (movies S1 and S2). (**B**) Continuity events (movies S3 and S4). (**C**) Support events (movies S5 and S6).

Infants' exploration and hypothesis testing after violations of expectation

Our finding that violations shaped infants' learning in a targeted way, enhancing learning only about objects relevant to the observed violation, raises a further question about the nature of the new information learned. In experiments 1 to 3, the new information taught to infants was arbitrary, in the sense that it did not clearly causally relate to the surprising violations (because the sound made by an object does not offer a direct explanation for its behavior). Besides enhancing learning for such arbitrary mappings [like those acquired by nonhuman animals (22)], do violations of expectation privilege the learning of particular kinds of information that are relevant to the nature of the surprising event? When an observation conflicts with prior knowledge, an effective learning strategy would be to seek evidence that could explain the discrepancy between what was predicted and what is observed. Older children engage in this kind of hypothesis testing, performing targeted actions to support or rule out possible explanations for an event (23, 24). But it is unknown whether preverbal infants actively test hypotheses about events, especially events involving violations of core knowledge.

In experiment 4, we first asked whether infants (N = 40) preferentially seek information from an object that violated expectations over an object that did not. Infants saw an event that either accorded with or violated the principles of object solidity or (extending our inquiry to another principle) object support (18) (movies S5 and S6). The solidity events were identical to those in experiment 1 (Knowledge-Consistent outcome, n = 10; Knowledge-Violation outcome, n = 10) (Fig. 1A). In the support event (Fig. 1C), infants saw an object (e.g., car) either pushed along a surface while remaining completely supported, thereby according with expectations about support (Knowledge-Consistent outcome, n = 10), or pushed over the surface edge without falling, thereby violating expectations about support (Knowledge-Violation outcome, n = 10) (25). As before, we limited infants' visual exposure to the event outcomes; a univariate ANOVA, with looking time to the event outcome as the dependent variable and event type (Solidity or Support) and outcome type (Knowledge-Consistent or Knowledge-Violation) as fixed factors, showed no main effect of outcome type [$F(1,36)$ = 0.794, P = 0.379] (18) (table S2).

After infants saw the outcome of the solidity or support event, we gave them two objects to freely explore for 60 s: the target object from the preceding event (e.g., car) and a new distractor object (e.g., ball; for half the infants the car was the target and the ball was the distractor, and for the other half this was reversed). We calculated infants' exploration preference scores by subtracting the amount of time they explored the new distractor object from the amount of time they explored the target object (table S2). We predicted that infants who had seen a Knowledge-Consistent event would show no preference, whereas infants who had seen a Knowledge-Violation event would prefer to explore the object that had just violated their expectations. A univariate ANOVA, with infants' exploration preference score as the dependent variable and event type (Solidity or Support) and outcome type (Knowledge-Consistent or Knowledge-Violation) as fixed factors, yielded a significant main effect of outcome type [$F(1,36)$ = 5.933, P = 0.02, partial η^2 = 0.14]: Infants who had seen the Knowledge-Violation event explored the target object more than infants who had seen the Knowledge-Consistent event. We then compared infants'
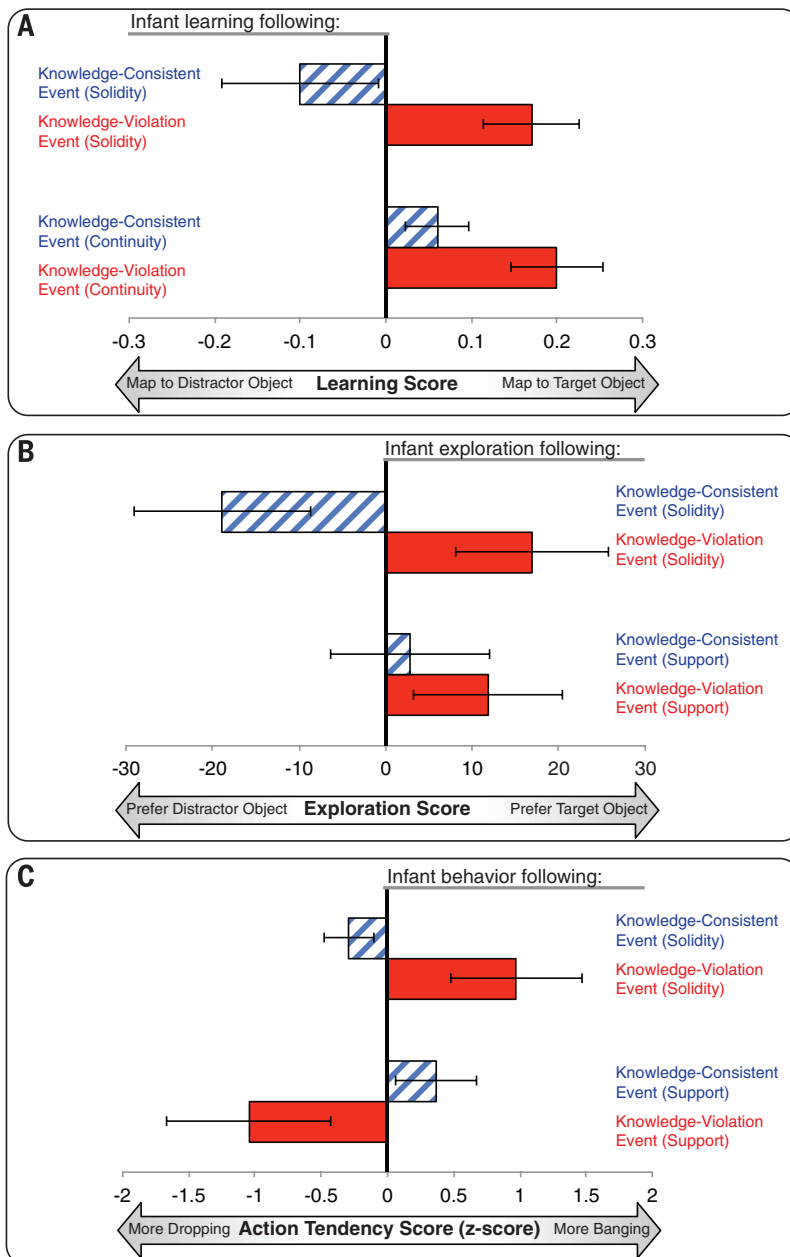

Fig. 2. Results from experiments 1 and 4. (**A**) Infants' learning after Knowledge-Consistent and Knowledge-Violation events in experiment 1. Bars represent average learning scores (proportion of looking at target object during mapping test minus proportion of looking at target object during baseline). (**B**) Infants' exploration after Knowledge-Consistent and Knowledge-Violation events in experiment 4. Bars represent looking at and/or touching the target object minus looking at and/or touching the new distractor object. (**C**) Infants' exploratory behaviors on the target object after Knowledge-Consistent and Knowledge-Violation events in experiment 4. Bars represent infants' z-scored object-banging behaviors minus z-scored object-dropping behaviors. All error bars represent SEM.

exploration preference scores to chance (zero). Collapsed across the solidity and support events, infants who had seen a Knowledge-Consistent event explored the target and distractor objects equally [$t(19) = -1.128$, $P = 0.27$], whereas infants who had seen a Knowledge-Violation event preferred to explore the target object [$t(19) = 2.395$, $P = 0.027$] (18) (Fig. 2B and table S2).

Infants who saw a violation event showed enhanced interest in the violation object, preferring to explore it over a new object. Because details of infants' exploratory behaviors might reveal an even richer interplay between knowledge and exploration, we next asked whether infants explored the target object qualitatively differently depending on which violation they had seen. We analyzed two common exploratory behaviors, each relevant to one of the presented events: banging an object (relevant to testing object solidity) and dropping an object onto the table or floor (relevant to testing object support). Because dropping an object takes longer than banging an object, we converted the frequency of these behaviors into z scores to enable direct comparison (table S2). To calculate infants' tendency to bang versus drop objects, we subtracted each infant's z-scored dropping frequency from their z-scored banging frequency. A univariate ANOVA, with action tendency score on the target object as the dependent variable and event type (Solidity or Support) and outcome type (Knowledge-Consistent or Knowledge-Violation) as fixed factors, yielded a significant interaction between event type and outcome type [$F(1,36) = 9.43$, $P = 0.004$, partial $\eta^2 = 0.208$] (Fig. 2C). An independent-samples t test revealed that infants who had seen an object appear to pass through a wall (Knowledge-Violation solidity event) banged that object more than they dropped it, relative to infants who had seen the same object stopped by the wall (Knowledge-Consistent solidity event) [$t(18) = 2.378$, $P = 0.029$]. By contrast, infants who had seen an object appear to hover in midair (Knowledge-Violation support event) did the reverse: They dropped the object more than they banged it, relative to infants who had seen the same object fully supported (Knowledge-Consistent support event) [$t(18) = -2.045$, $P = 0.056$] (18) (Fig. 2C and table S2). This double dissociation in infants' behavior—wherein infants who saw a solidity violation tended to actively bang the target object, whereas infants who saw a support violation tended to drop it—shows that infants tailored their exploratory actions to the type of violation seen.

The infants' banging and dropping of the new distractor object that had not participated in the solidity or support event did not differ across event types or outcomes. A univariate ANOVA, with action tendency score on the distractor object as the dependent variable and event type (Solidity or Support) and outcome type (Knowledge-Consistent or Knowledge-Violation) as fixed factors, yielded no significant interaction [$F(1,36) = 0.062$, $P = 0.80$]. Critically, a repeated-measures ANOVA that examined

action tendency scores across object type (target or distractor), event type (Solidity or Support), and outcome type (Knowledge-Consistent or Knowledge-Violation) yielded a significant interaction among these three factors [$F(1,36) = 4.95$, $P = 0.032$, partial $\eta^2 = 0.12$] (18); infants performed differential actions only after Knowledge-Violation events and only on the objects that had committed the violation. This dissociation in infants' actions on just the target object reveals two senses in which infants' behaviors were highly directed: They focused on the entity that had violated expectations, and they were relevant to the nature of the observed violation. Thus, infants' behaviors are not merely reflexive responses to the novelty of surprising outcomes but instead reflect deeper attempts to learn about aspects of the world that failed to accord with expectations.

Conclusions

Our findings show that infants' learning is changed when their expectations are violated. Much as scientists faced with unexpected patterns of data are propelled to think harder, run further experiments, or change their methods of inquiry, untutored preverbal minds are sensitive to conflict between the predicted and the observed, and use this conflict as a scaffold for new learning.

In our experiments, we tested learning after violations of expectations drawn from core knowledge of object behavior—knowledge that is available from early in life, is universal across human cultures, and is present in other species. The existence of these foundational expectations has been used to argue for the presence of rich innate knowledge in infants; given our finding that violations of these expectations lead to enhanced learning, early knowledge and early learning are mutually reinforcing. In addition, expectancy violations involving other types of knowledge are also likely to be important in learning. Children form new expectations by tracking experienced contingencies (26), by receiving others' testimony (27), and by using abstract knowledge to form probabilistic predictions about events they have never observed (28). Some of these sophisticated behaviors have been interpreted in terms of Bayesian inferences that generate knowledge by weighing new evidence against prior beliefs (29, 30). Our findings accord well with such a framework and suggest avenues to explore how violations detected in different domains of prior knowledge, or using different kinds of new evidence, shape exploration and learning throughout the life span and across species.

Together, our experiments reveal that when infants see an object defy their expectations, they learn about that object better, explore that object more, and test relevant hypotheses for that object's behavior. Seen through this lens, the decades of findings that infants look longer at surprising events suggest not only that infants are equipped with core knowledge about fundamental aspects of the world but also

that this knowledge is harnessed to empower new learning even in infancy. Thus, core knowledge is not an alternative to learning but is instead a key ingredient in driving learning forward.

REFERENCES AND NOTES

1. N. Chomsky, in *Language and Learning: The Debate Between Jean Piaget and Noam Chomsky*, M. Piatelli-Palmarini, Ed. (Routledge and Kegan Paul, London, 1980), pp. 393–396.
2. J. B. Tenenbaum, C. Kemp, T. L. Griffiths, N. D. Goodman, *Science* 331, 1279–1285 (2011).
3. G. Csibra, G. Gergely, *Trends Cognit. Sci.* 13, 148–153 (2009).
4. V. Izard, C. Sann, E. S. Spelke, A. Streri, *Proc. Natl. Acad. Sci. U.S.A.* 106, 10382–10385 (2009).
5. L. Regolin, G. Vallortigara, *Percept. Psychophys.* 57, 971–976 (1995).
6. C. Chiandetti, G. Vallortigara, *Anim. Cognit.* 13, 463–470 (2010).
7. S. Dehaene, V. Izard, P. Pica, E. Spelke, *Science* 311, 381–384 (2006).
8. E. S. Spelke, K. D. Kinzler, *Dev. Sci.* 10, 89–96 (2007).
9. E. S. Spelke, K. Breinlinger, J. Macomber, K. Jacobson, *Psychol. Rev.* 99, 605–632 (1992).
10. T. Wilcox, L. Nadel, R. Rosser, *Infant Behav. Dev.* 19, 309–323 (1996).
11. L. A. Camras *et al.*, *Emotion* 2, 179–193 (2002).
12. G. Gredebäck, A. Melinder, *Cognition* 114, 197–206 (2010).
13. T. Wilcox, H. Bortfeld, R. Woods, E. Wruck, D. A. Boas, *J. Biomed. Opt.* 10, 011010–011019 (2005).
14. A. Berger, G. Tzur, M. I. Posner, *Proc. Natl. Acad. Sci. U.S.A.* 103, 12649–12653 (2006).
15. R. Baillargeon, E. S. Spelke, S. Wasserman, *Cognition* 20, 191–208 (1985).
16. K. McCrink, K. Wynn, *Psychol. Sci.* 15, 776–781 (2004).
17. V. Kuhlmeier, K. Wynn, P. Bloom, *Psychol. Sci.* 14, 402–408 (2003).
18. See supplementary materials on *Science* Online.
19. R. Baillargeon, *Cognition* 23, 21–41 (1986).
20. E. S. Spelke, R. Kestenbaum, D. Simons, D. Wein, *Br. J. Dev. Psychol.* 13, 113–142 (1995).
21. P. K. Kuhl, A. N. Meltzoff, *Infant Behav. Dev.* 7, 361–381 (1984).
22. J. M. Pearce, G. Hall, *Psychol. Rev.* 87, 532–552 (1980).
23. E. B. Bonawitz, T. J. P. van Schijndel, D. Friel, L. Schulz, *Cognit. Psychol.* 64, 215–234 (2012).
24. C. H. Legare, *Child Dev.* 83, 173–185 (2012).
25. A. Needham, R. Baillargeon, *Cognition* 47, 121–148 (1993).
26. J. R. Saffran, R. N. Aslin, E. L. Newport, *Science* 274, 1926–1928 (1996).
27. V. K. Jaswal, *Cognit. Psychol.* 61, 248–272 (2010).
28. E. Téglás *et al.*, *Science* 332, 1054–1059 (2011).
29. J. B. Tenenbaum, T. L. Griffiths, C. Kemp, *Trends Cognit. Sci.* 10, 309–318 (2006).
30. L. Schulz, *Trends Cognit. Sci.* 16, 382–389 (2012).

SUPPLEMENTARY MATERIALS

www.sciencemag.org/content/348/6230/91/suppl/DC1
Materials and Methods
Supplementary Text
Tables S1 and S2
Movies S1 to S6

Observing the unexpected enhances infants' learning and exploration
Aimee E. Stahl and Lisa Feigenson
Science **348**, 91 (2015);
DOI: 10.1126/science.aaa3799

This copy is for your personal, non-commercial use only.