Rational Approaches to Learning and Development

by

Celeste Kidd

Submitted in Partial Fulfillment of the Requirements

for the Degree Doctor of Philosophy

Supervised by Professor Richard N. Aslin

Department of Brain & Cognitive Sciences

Arts, Sciences and Engineering

School of Arts and Sciences

University of Rochester

Rochester, New York

May 2013

Dedication.

This thesis is dedicated to the children of the Orange County Catholic Worker

(Isaiah House) in Santa Ana, California.

Biographical Sketch.

The author was born in New Orleans, Louisiana. She completed coursework in computer science and investigative journalism at the University of California, Santa Cruz before attending the University of Southern California. While at USC, she worked as a research assistant to Dani Byrd, Shrikanth Narayanan, Rachel Walker, and Toben Mintz; she also conducted independent research on speech processing and language acquisition under the supervision of Toben Mintz, Dani Byrd, Rachel Walker, and Elsi Kaiser. She graduated from USC with Bachelor of Arts degrees in Linguistics and Print Journalism. In 2008, the National Science Foundation awarded her a Graduate Research Fellowship in support of her research on early human development and attention under the direction of Richard N. Aslin in Brain and Cognitive Sciences at the University of Rochester. In 2010, she was awarded a Computational Modeling Prize in Perception/Action from the Cognitive Science Society for her computational work on infant attention. She received the Master of Arts degree in Brain and Cognitive Sciences at the University of Rochester in 2011. The following publications were the result of work conducted during her doctoral study:

Peer-reviewed research

(1) Piantadosi, S.T., Kidd, C., & Aslin, R.N. (In press.) Rich Analysis and Rational Models: Inferring individual behavior from infant looking data. *Developmental Science.*

(2) Kidd, C., Palmeri, H., Aslin, R.N. (2013). Rational Snacking: Young children's decision-making on the marshmallow task is moderated by beliefs about environmental reliability. *Cognition, 126*(1), 109-114.

(3) 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.

(4) Kidd, C., White, K.S., & Aslin, R.N. (2011). Toddlers use speech disfluencies to predict speakers' referential intentions. *Developmental Science, 14*(4): 925–934.

Under review and in progress

(5) Kidd, C., Piantadosi, S.T., & Aslin, R.N. (Under review.) The Goldilocks effect in infant auditory cognition.

(6) Mintz, T.H., Walker, R.L., Welday, A., & Kidd, C. (Under review). Infants' Universal Sensitivity to Vowel Harmony and Its Role in Segmenting Speech.

(7) Kidd, C., Blanchard, T., Aslin, R.N., & Hayden, B.Y. (In prep). Goldilocks Effect in Human Infants and Juvenile Non-Human Primates: Cross-species learners preferentially attend to moderately unexpected visual events over those that are too predictable or too improbable.

(8) Kidd, C. & Aslin, R.N. (In prep). Infants' rational use of social cues during word learning.

Peer-reviewed proceedings papers

(9) Kidd, C., White, K.S., & Aslin, R.N. (2009). Children's Use of Disfluencies for Pragmatic Inference in Lexical Development. *Proceedings of the 31st Annual Meeting of the Cognitive Science Society (pp. 1494-1499).*

(10) Kidd, C., Piantadosi, S.T., & Aslin, R.N. (2010). The Goldilocks Effect: Infants' preference for visual stimuli that are neither too predictable nor too surprising. *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society (pp. 2476-2481).*

Book chapter

(11) Kidd, C., White, K.S., & Aslin, R.N. (2011). Learning the Meaning of "Um": Toddlers' developing use of speech disfluencies as cues to speakers' referential intentions. In I. Arnon & E.V. Clark (Eds.), *Experience, Variation, and Generalization: Learning a First Language* (Trends in Language Acquisition Research). Amsterdam: John Benjamins. (pp. 91–106).

Review article

(12) McGovern, G.C., Kidd, C., & Piantadosi, S.T. (2010). Psychology of authority in quid pro quo sexual harassment cases: 50 years of psychological science demonstrate the influence of authority and social norms. *Louisiana Advocates* (pp. 14-17). Baton Rouge: Louisiana Association for Justice.

Acknowledgements.

Abstract.

Rational cognitive theories posit that organisms act to optimize utility. This capacity depends on generating accurate predictions about the future—which, in turn, requires accurate mental models of the world. Adults' decisions are guided by their substantial experience in the world's chain of causality. Very young children, however, possess comparatively less data. In this dissertation, I study decision-making mechanisms in young children and non-human-primates across multiple domains—including visual attention and overt choice—in order to discover the efficacy and limitations of rational cognitive theories. I present empirical evidence that learners rely on utility maximization both to build complex models of the world starting from very little knowledge and, more generally, to guide their decisions and behavior. Four experiments were conducted on 4-year-olds, 7-month-olds, and monkeys using visual and auditory stimuli presented in sequences of events. These experiments show that children are capable of rational decisions to optimize future utility and exhibit a U-shaped relationship between stimulus complexity and attention. Similarly, monkeys' attentional patterns are guided by stimulus complexity.

Contributors and Funding Sources.

Table of Contents.

List of Tables.

List of Figures.

I. Introduction.

The utility of applying rational approaches to the study of learning and development.

Rationalist cognitive theories posit that organisms should choose actions that optimize utility (e.g., Anderson, 1991; Oaksford & Chater, 1994). Effective utility-maximization requires an agent to generate accurate expectations about what will happen in the future—which, in turn, requires the agent to mentally model the world. Adult agents have a substantial amount of experience in the world upon which they can base such models; thus, their previous experience plays an important role in guiding their decisions. Very young children, however, have far less world experience, and thus—at least initially—possess no such advantage. They must then sample observations from their environments in order to overcome their naïveté concerning the structure of those environments. Sample by sample, they gradually infer complex representations and form abstract theories about the world. Sample by sample, they continuously build upon, update, and revise those theories.

In this dissertation, I investigate both implicit and overt measures of young children's choice behavior throughout development in order to understand the

decision-making mechanisms that guide the acquisition of knowledge. This research aims to test the efficacy and limitations of rational cognitive by better understanding the decision-making mechanisms that guide learning and behavior in young children and non-human primates. The studies span across multiple domains—including visual attention and overt choice. In this thesis, I present empirical evidence that suggests that young learners rely on rational utility maximization both to build complex models of the world starting from very little knowledge and, more generally, to guide their decisions and behavior.

Bounded Rationality.

In a classical view of rationality—the one espoused by Tversky and Kahneman (1973)—rationality is defined as strict adherence to the laws of probability theory. More contemporary views (e.g., Chase, Hertwig, & Gigerenzer, 1998; Oaksford & Chater, 2001), however, suggest rational agents should instead adhere to principles of so-called "bounded rationality." Under these views, rationality is defined not only in terms of the organisms' goals, but also in terms of the context, which includes competing goals, environmental restrictions, and cognitive limitations. Thus, human reasoning may be considered rational if people are as accurate as possible given the appropriate constraints imposed by time, available information, and limited cognitive resources.

One cannot determine the degree to which an actor is rational without a set of assumptions about to that actor's context. Yet determining the correct set of background assumptions is somewhat problematic given everything that we do not yet understand about human cognition, particularly in the domains of development and learning. To begin to make progress in understanding human cognition with this difficulty in mind, rational cognitive theorists (e.g., Marr, 1982; Anderson,

1991; Oaksford & Chater, 1994) have proposed defining a small set of reasonable assumptions—the smaller and more empirically justified, the better— and then determining what behavior would maximize utility in light of the defined context and constraints (through logic, probability theory, or statistical models). Once the "best" behavior is well defined, it can be compared to empirical measurements of actual human behavior. When human behavior falls short of the theorized "best" behavior, the severity and nature of the failures can be informative for theorizing about the cognitive context under which human actors may be operating. These theories about cognitive limitations and constraints can then be tested empirically—and if justified—eventually used as assumptions in later rounds of rational cognitive theories to generate new predictions and draw new comparisons to actual human data. Thus, in many ways, the rational analysis approach is both a theory, which posits that actors optimize utility in context—but also a framework, which may be used to guide empirical investigations and further theorizing about cognition.

To paraphrase Thomas Nagle, any current scientific theory is almost certainly false. Our role as scientists is to build off of the empirical work of our predecessors, tweaking their theories to better explain everything we've learned since then, then formally testing those theories by collecting new sets of empirical data. What we derive are better—but still imperfect—formalized theories for the next generation of research to tweak, test, and tweak again. Rational analysis offers us a formal system for running through this process. It suggests a method for defining the next reasonable theory to test, and some guiding principles to use in investigating the utility and shortfalls of that (hopefully improved, but almost certainly imperfect) new theory.

The earliest studies of infant attention focused predominately on explaining infant behavior in terms of low-level stimulus properties (e.g., color, contrast,

luminance). The next generation of scientists cleverly used these data to develop *cognitive processing*-based theories of stimulus preference, which they painstakingly tested with the most reasonable assumptions about infant cognition available at the time (e.g., Haith, 1980). These scientists, however, lacked the ability to formally test many aspects of these processing-based theories as such theories necessarily required relating infants' attentional behavior to infants' existing knowledge. With no direct way of observing infants' existing knowledge, and very limited reasonable assumptions about what each infant might know before or over the course of an experiment, progress on testing these big-picture attentional theories slowed substantially for more than three decades. Since then, computers have become more powerful, enabling us to use more sophisticated empirical and computational methods to test and represent infants' knowledge states in ways that weren't previously possible.

Model-Based Behavioral Experimentation.

A key feature of the research approach I present is the combination of behavioral methods and computational modeling. This model-driven behavioral experimentation enables me to rigorously test competing theories of decision-making and learning by quantifying otherwise unobservable cognitive processes or variables. For instance, in the "Goldilocks" work, our model is primarily used as a measure of an otherwise unobservable feature of the world, the information conveyed by a stimulus. By relating the model-based measure of information to infants' behavior, we are able to formalize and test a hypothesis about infant attention that had previously only been studied qualitatively. This is a powerful approach because traditional infant methods typically only compare the preferences of groups of infants. My work builds upon these traditional methods,

but attempts to formalize and test detailed predictions about behavioral patterns, allowing me to examine a wide range of formal theories. As an added bonus, this approach offers the potential to generate specific predictions about the learning outcomes of individual children on the basis of their particular behavioral patterns.

Children as Rational Agents.

Parents may find the description of infants and young children as rational agents humorous—and with good reason. Young children are infamous for their apparently *irrational* behavior—especially toddlers (the so-called "terrible twos"). If you have never observed this sort of behavior in a child firsthand, you need only consult the internet. In the popular Tumblr blog "Reasons My Son Is Crying," Greg Pembroke, father of two in Rochester, NY, catalogs world events that drive his young sons William, 3, and Charlie, 21 months, to tears. Here are just a few of the reasons:

- "I washed the dirt and sand off his pear."
- "The neighbor's dog isn't outside."
- "I wouldn't let him eat Buzz Lightyear's head."
- "The milk isn't juice."

And temper tantrums are just the beginning. Children's capacity for questionable decision-making is at times seemingly boundless. They drink their own bathwater, stash Legos in their noses, and spontaneously discard their own socks and shoes in inconvenient locations. When they play hide-and-seek, they often fail to fully obscure themselves, pick predictable places (like the same place over and over again), or reveal themselves before you've even started looking for them by giggling. While these behaviors may not seem like sensible strategies on the surface, these behaviors must be interpreted in the context of childrens' current

beliefs and knowledge—not a full, adult-like conceptual system. For instance, children may be choosing the best hiding spot possible on the basis of their limited experiences and degraded models of how things operate in the world. Without experience, children may not have a great model of their own size, or what sorts of hiding places their playmates are likely to check. If a child *believes* they are fully obscured from view, the decision to hide under a bed sheet is actually completely reasonable. In other words, children may actually be adhering to the principles of bounded rationality—but *highly* bounded.

Throughout the work in this thesis, I will take an approach to understanding behavior that is based on rational analysis. The primary projects contributing to this work are described in the next sections.

Rational Decision-Making in Young Children.

I begin by applying this approach to perhaps the most classic example of seemingly irrational behavior in young children—their poor track record in delay-of-gratification tasks, such as in the Stanford Marshmallow experiments (e.g., Mischel & Ebbesen, 1970). Though children apparently fail to maximize utility in such delay-of-gratification tasks, the cause of these apparent failures was not well understood (Chapter 2, Kidd, Palmeri, & Aslin, 2012). For example, most 3- to 5-year-olds choose an immediately available low-value reward (e.g., one marshmallow) over one of high-value (e.g., two marshmallows) after a temporal delay (Mischel & Ebbesen, 1970). One possible explanation of this choice is a deficiency in self-control: young children may be incapable of inhibiting their immediate-response tendencies to seek gratification (e.g., Marcovitch & Zelazo, 1999; Piaget, 1954). However, following the framework of testing for rational behavior under different assumptions, another possibility is that children's

performance may result from their expectations and beliefs, which are likely different from adults' and vary across children. Under this second theory, children engage in rational decision-making about whether waiting for the high-value reward yields an expected gain in utility. Waiting is only the rational choice if a child believes that the high-value reward will arrive as promised. To compare these two theories, we tested 3- to 5-year-old children using a classic delay-of-gratification paradigm—the marshmallow task (Mischel, 1974). We preceded marshmallow-task testing with evidence that the experimenter running the study was either reliable or unreliable as a means of manipulating children's beliefs across conditions. Children who believed the experimenter was reliable waited about four times longer before eating the marshmallow than those who thought she was unreliable (12 min vs. 3 min, $p < 0.0005$). These results suggest that children's wait-times are modulated by a rational decision-making process that considers environmental reliability. They may also provide an alternative explanation for why marshmallow wait-times correlate with later life success (e.g., Mischel, Shoda, & Rodriguez, 1989)—successful people grow up in reliable situations. Broadly, this illustrates that children build a model of the reliability of others' behavior, and use this model to inform their decisions.

Rational Models of Attention and Learning.

The marshmallow study relied on the fact that learners have aggregated information about the reliability of adult reward-promises prior to being tested. However, learners do not enter the world with access to most of this information—how do infants begin to make sense of the world with little or no knowledge on which to base their inferences? In Chapters 3 and 4, I will apply the rational approach embodied by the Rational Snacking project to my primary line

of research, infant attention. These results suggest that key attentional mechanisms filter environmental stimuli in a particularly useful way, thereby providing infants with data that are "just right" for learning (which we referred to as a "Goldilocks" effect). This work (Kidd, Piantadosi, & Aslin, 2010, 2012, under review; Piantadosi, Kidd, & Aslin, in press) explored attentional behavior in 7- and 8-month-old infants. We showed infants visual event sequences of varying complexity, as measured by an idealized learning model, and measured when in each sequence infants decided to terminate their attention by looking away from the display. We found that infants' probability of looking away was greatest to events of either very low information content (highly predictable) or very high information content (highly surprising). This attentional strategy holds in multiple types of visual displays (Chapter 3, Kidd, Piantadosi, & Aslin, 2010, 2012), for auditory stimuli (Chapter 4, Kidd, Piantadosi, & Aslin, under review), and even within individual infants (Chapter 5, Piantadosi, Kidd, & Aslin, in press). These results suggest a broadly applicable principle of infant attention. They suggest that infants implicitly decide to direct their attention in order to maintain intermediate rates of information absorption. Thus, infants likely avoiding wasting cognitive resources on overly predictable or overly complex events.

In Chapter 3, I will explain how these findings likely represent a resolution to a classic methodological problem in developmental psychology, in which preferential looking and listening studies typically find either novelty or familiarity effects. In our work, we have an explicit metric of information complexity, in contrast to previous theories of the novelty/familiarity conundrum that offered no quantitative predictions for which direction would be expected in any given experiment. According to our model, novelty preferences arise when the complexity of the stimulus-set is relatively low, while familiarity preferences occur when it is relatively high (Kidd, Piantadosi, & Aslin, 2012, 2013). In both cases,

the preferred stimulus is closer to the infants' preferred medial-value of complexity. These findings also have implications for interpreting null results, which can occur when the stimulus-set comprises one item from each extreme of the complexity continuum (one overly simple and one overly complex). Though infants may show no preference for either item, their disinterest may stem from the competition between two fundamentally different processes (disinterest in simplicity and disinterest in randomness).

Monkeys as Rational Agents.

The final chapter reflect the result of a collaboration with Tommy Blanchard, Richard N. Aslin, and Benjamin Hayden to extend the foregoing work on infant attention to ongoing work with awake behaving rhesus macaques (Kidd, Blanchard, Aslin, & Hayden, in prep). We presented monkeys with sequential visual displays featuring visual events that varied in terms of their informational complexity (surprisal), as in Kidd, Piantadosi, & Aslin (2012). We analyzed these data similarly to the infant data, and the results of these analyses suggest that visual attention in macaques is also sensitive to the information-theoretic properties of the stimuli. Further, the data suggest that both infant and monkey learners employ rational strategies that favor more informative visual events. These results provide evidence that rational behavior during the implicit decision-making processes of deciding where to attend in the world may be a feature common to all learners, not just humans.

The monkeys, however, do not exhibit all of the same attentional strategies as human infants or human adults. I highlight a few of the ways in which monkeys are different from humans. In particular, I discuss empirical findings suggesting that humans devote far more of their attentional resources to tracking transitional

statistics in the world than do the monkeys, whose attentional behavior is far better predicted by simple zero-order frequency statistics about visual events in the world.

In Summation.

In general, my research illustrates cases where children's behavior is best understood as a form of utility maximization—either in terms of deciding how to optimize their attention to incoming stimuli (Kidd, Piantadosi, & Aslin, 2012) or deciding to delay gratification in order to achieve a greater reward when that reward is statistically likely (Kidd, Palmeri, & Aslin, 2013). The ability to make good decisions about stimuli that support learning may lead to a powerful feedback loop. The better the actor's model of the world, the better the actor can maximize their utility through good decision-making; the better their decision-making, the more efficiently they can choose the best input to attend to in order to improve their model of the world. Some of my work in progress explores the relationship between implicit attentional decisions and later learning outcomes. My future work will further examine these interactions by building and testing rational models that can capture the general principles shared across explicit and implicit decision-making mechanisms.

References.

Allopenna, P.D., Magnuson, J.S. & Tanenhaus, M.K. (1998). Tracking the time course of spoken word recognition: evidence for continuous mapping models. *Journal of Memory and Language, 38*, 419-439.

Anderson, J. (1991). Cognitive architectures in a rational analysis. In K. VanLehn (Ed.), *Architectures for Intelligence* (pp. 1–24). Hillsdale, NJ: Lawrence Erlbaum Associates.

Baldwin, D.A., Markman, E.M., Bill, B., Desjardins, R.N., & Irwin, J.M. (1996). Infants' reliance on a social cri-terion for establishing word-object relations. *Child Development, 67*(6), 3135-3153.

Chase, V.M., Hertwig, R., & Gigerenzer, G. (1998). Visions of rationality. *Trends in Cognitive Sciences, 2*(6), 206-213.

Clark, H.H. & Fox Tree, J.E. (2002). Using uh and um in spontaneous speech. *Cognition, 84*, 73-111.

Kidd, C., & Aslin, R.N., (2010). Infants' looking behavior demonstrates use of rational statistical inference. *Talk presented at the International Conference on Infant Studies,* 13 March 2010, Baltimore, MD.

Kidd, C., & Aslin, R.N., (In prep). Rational use of social cues in word learning.

Kidd, C., Blanchard, T., Aslin, R.N., & Hayden, B. (In prep). Information-seeking behavior in human-infant and macaque learners.

Kidd, C., Palmeri, H., & Aslin, R.N. (2013). Rational Snacking: Young children's decision-making on the marshmallow task is moderated by beliefs about environmental reliability. *Cognition.*

Kidd, C., Piantadosi, S.T., & Aslin, R.N. (2010). The Goldilocks Effect: Infants' preference for visual stimuli that are neither too predictable nor too surprising. *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society* (pp. 2476-2481).

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.

Kidd, C., Piantadosi, S.T., & Aslin, R.N. (Under review). The Goldilocks effect in infant auditory cognition. Manuscript submitted for publication.

Kidd, C., White, K.S., & Aslin, R.N. (2009). Children's use of disfluencies for pragmatic inference in lexical de-velopment. *Proceedings of the 31st Annual Meeting of the Cognitive Science Society* (pp. 1494-1499).

Kidd, C., White, K.S., & Aslin, R.N. (2011a). Toddlers use speech disfluencies to predict speakers' referential intentions. *Developmental Science, 14*(4), 925–934.

Kidd, C., White, K.S., & Aslin, R.N. (2011b). Learning the Meaning of "Um": Toddlers' developing use of speech disfluencies as cues to speakers' referential intentions. In I. Arnon & E.V. Clark (Eds.), *Experience, Variation, and Generalization: Learning a First Language* (Trends in Language Acquisition Research). Amsterdam: John Benjamins. (pp. 91–106).

Marcovitch, S., & Zelazo, P.D. (1999). The A-not-B error: Results from a logistic meta-analysis. *Child Development, 70*, 1297–1313.

Mischel, W. (1974). Processes in delay of gratification. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 7, pp. 249–292). New York: Academic Press.

Mischel, W., & Ebbesen, E.B. (1970). Attention in delay of gratification. *Journal of Personality and Social Psychology, 21*, 204–218.

Mischel, W., Shoda, Y., & Rodriguez. (1989). Delay of gratification in children. *Science, 244*, 933–938.

Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review, 101*(4), 608-631.

Piaget, J. (1954). *The construction of reality in the child.* Oxford: Basic Books.

Piantadosi, S.P., Kidd, C., & Aslin, R.N. (In press). Rich Analysis and Rational models: Inferring individual behavior from infant looking data. *Developmental Science.*

II. Rational Snacking.

Young children's decision-making on the marshmallow task is moderated by beliefs about environmental reliability.

Celeste Kidd, Holly Palmeri, & Richard N. Aslin

Children are notoriously bad at delaying gratification to achieve later, greater rewards (e.g., Piaget, 1970)—and some are worse at waiting than others. Individual differences in the ability-to-wait have been attributed to self-control, in part because of evidence that long-delayers are more successful in later life (e.g., Shoda, Mischel, & Peake, 1990). Here we provide evidence that, in addition to self-control, children's wait-times are modulated by an implicit, rational decision-making process that considers environmental reliability. We tested children (M = 4;6, N = 28) using a classic paradigm—the marshmallow task (Mischel, 1974)—in an environment demonstrated to be either unreliable or reliable. Children in the reliable condition waited significantly longer than those in the unreliable condition ($p < 0.0005$), suggesting that children's wait-times reflected reasoned beliefs about whether waiting would ultimately pay off. Thus, wait-times on sustained delay-of-gratification tasks (e.g., the marshmallow task) may not only reflect differences in self-control abilities, but also beliefs about the stability of the world.

Introduction.

When children draw on walls, reject daily baths, or leave the house wearing no pants and a tutu, caretakers may reasonably doubt their capacity for rational decision-making. However, recent evidence suggests that even very young children possess sophisticated decision-making capabilities for reasoning about physical causality (e.g., Gopnik et al., 2004; Gweon & Schulz, 2011), social behavior (e.g., Gergely, Bekkering, & Király, 2002), future events (e.g., Denison & Xu, 2010; Kidd, Piantadosi, & Aslin, 2012; Téglás et al., 2011), concepts and categories (e.g., Piantadosi, Tenenbaum, & Goodman, 2012; Xu, Dewar, & Perfors, 2009), and word meanings (e.g., Xu & Tenenbaum, 2007). Here we demonstrate that young children also use their rational decision-making abilities in a domain of behavioral inhibition: a sustained delay-of-gratification task.

Decision-making is said to be rational if it maximizes benefit or utility (Anderson, 1991; Anderson & Milson, 1989; Marr, 1982), yet young children's decisions during delay-of-gratification tasks often appear to do just the opposite (e.g., Mischel & Ebbesen, 1970). When asked to resist the temptation of an immediately available low-value reward to obtain one of high-value after a temporal delay, 75% of children failed to do so, succumbing to their desire after an average of 5.72 min. The cause of these apparent failures of rationality, however, is not fully understood. While children's failures to wait are likely the result of a combination of many genetic and environmental variables, two potentially important factors are *self-control capacity* and *established beliefs*.

DEFICIENT CAPACITY HYPOTHESIS.

One possible explanation for failing to wait for a larger reward is a deficiency in *self-control*; some children are simply incapable of inhibiting their

immediate-response tendency to seek gratification. Young infants, for example, have not yet developed the executive functions necessary for inhibitory control (e.g., Piaget, 1970), as evidenced by the perseveration errors made by up to 2-year-old children in A-Not-B tasks (e.g., Marcovitch & Zelazo, 1999; Piaget, 1954). As predicted by this theory, children's ability to delay gratification improves with maturation (e.g., Mischel & Metzner, 1962). Maturational changes, however, are insufficient to account for all of the variance in task performance (e.g., Romer, Duckworth, Sznitzman, & Park, 2010). *Individual differences* in children's capacities for self-control may account for the remaining variance.

Self-control has been implicated as a major causal factor in a child's later life successes (or failures). Mischel, Shoda, and Peake (1988) analyzed data from adolescents who, many years earlier, had been presented with a laboratory choice-task: eat a single marshmallow immediately, or resist the temptation during a sustained delay to receive two marshmallows. With no means of distracting themselves from a treat left in view, the majority of children failed to wait for the maximum delay (15 or 20 min) before eating the marshmallow, with a mean wait-time of 6 min and 5 s. Importantly, longer wait-times among children were correlated with greater self-confidence and better interpersonal skills, according to parental report. Longer wait-times also correlated with higher SAT scores (Shoda et al., 1990), less likelihood of substance abuse (Ayduk et al., 2000), and many other positive life outcomes (e.g., Mischel, Shoda, & Rodriguez, 1989). Based on these findings, the marshmallow task was argued to be a powerful diagnostic tool for predicting personal well-being and later-life achievement—"an early indicator of an apparently long-term personal quality" (Mischel et al., 1988). The logic of the claim is that a child who possesses more self-control can resist fleeting temptations to pursue difficult goals; in contrast, children with less self-control fail to persist toward these goals and thus achieve less. To be clear, the evidence for

poor self-control in young children (e.g., Baumeister, Heatherton, & Tice, 1994; Goleman, 1995), in a wide variety of tasks and contexts, is undeniable. At issue is the *origin of failure* of delay-of-gratification in laboratory tests like the marshmallow task, which has remained largely speculative (Mischel et al., 1989, p. 936).

RATIONAL DECISION-MAKING HYPOTHESIS.

Another possibility is that the variance in children's performance may be due to differences in children's expectations and beliefs (Mahrer, 1956; Mischel, 1961; Mischel & Staub, 1965). Under this theory, children engage in rational decision-making about whether to wait for the second marshmallow. This implicit process of making rational decisions is based upon beliefs that the child acquired before entering the testing room. The basis for this theory centers on what it means to be rational in the context of the marshmallow task. Waiting is only the rational choice if you believe that a second marshmallow is likely to actually appear after a reasonably short delay—and that the marshmallow currently in your possession is not at risk of being taken away. This presumption may not apply equally to all children. Consider the mindset of a 4-year-old living in a crowded shelter, surrounded by older children with little adult supervision. For a child accustomed to stolen possessions and broken promises, the only guaranteed treats are the ones you have already swallowed. At the other extreme, consider the mindset of an only-child in a stable home whose parents reliably promise and deliver small motivational treats for good behavior. From this child's perspective, the rare injustice of a stolen object or broken promise may be so startlingly unfamiliar that it prompts an outburst of tears. The critical point of the foregoing vignette is that rational behavior is inferred by understanding the goals and expectations of the agent (Anderson, 1991; Anderson & Milson, 1989; Marr, 1982). Relevant to this

hypothesis is the fact that children with absent fathers more often prefer immediate, lesser rewards over delayed, more valuable ones (Mischel, 1961). Also, children's willingness to wait is negatively impacted by uncertainty about the likelihood, value, or temporal availability of the future reward (Fawcett, McNamara, & Houston, 2012; Mahrer, 1956; McGuire & Kable, 2012; Mischel, 1974; Lowenstein, Read, & Baumeister, 2003). These effects are consistent with the idea that children may be capable of engaging in a rational process when deciding whether or not to wait.

In support of this second hypothesis, we present evidence that the reliability of the experimenter in the testing environment influences children's wait-times during the marshmallow task. Half of the children observed evidence that the researcher was *reliable* in advance of the marshmallow task, while half observed evidence that she was *unreliable*. If children employ a rational process in deciding whether or not to eat the first marshmallow, we expect children in the reliable condition to be significantly more likely to wait than those in the unreliable condition. Our experiment provides a fundamental test of this perspective on children's rational behavior and provides compelling evidence that young children are indeed capable of delaying gratification in the face of temptation when provided with evidence that waiting will pay off.

Materials and Methods.

PARTICIPANTS.

Twenty-eight caretakers volunteered their children (ages 3;6 – 5;10) for the study. The children were all healthy, had not recently visited the lab (within 2 months), and had not interacted with researchers running the study since infancy.

These precautions ensured children had minimal prior expectations specific to the lab or researcher's reliability before this study. Children were randomly assigned to one of two experimental conditions—*unreliable* and *reliable*—such that each group was gender and age balanced (nine males, five females, and M = 4;6). Participants received a small treat bag and $10 as compensation.

PROCEDURE.

Art project task.

Before the marshmallow task, children were first provided with evidence about the reliability of the researcher through the completion of a two-part art project involving a Create-Your-Own-Cup kit (with which children could decorate a blank paper slip to be inserted into a special cup). Each of the project's two parts involved a crucial choice. In Choice 1, the child could either use well-used crayons or wait for a new set of art supplies. In Choice 2, the child could either use one small sticker or wait for a new set of better stickers. Upon arrival, children were escorted to the "art project room" that was not part of the normal lab space and where parents could covertly observe them from the main lab space.

For Choice 1, the researcher presented the child with a small set of well-used crayons in a tightly sealed widemouth jar. The researcher explained that the child could use the crayons now, or wait until the researcher returned with a brand-new set of exciting art supplies to use instead. The researcher then placed the tightly sealed crayon jar in the center of the table and left the child alone in the room to wait for 2.5 min. Though we wanted children to ostensibly have a choice, we wanted them to choose to wait. Thus, the chosen container was intentionally difficult to open. This manipulation was successful, and all children waited the full 2.5 min without using the well-used crayons. In the *unreliable* condition, the

researcher returned without the promised art set and provided the following explanation:

> "I'm sorry, but I made a mistake. We don't have any other art supplies after all. But why don't you just use these instead?"

The researcher then helped the child open the jar of well-used crayons. In the *reliable* condition, the researcher returned with a rotating tray featuring a large assortment of exciting art supplies. (See Appendix II-1 for full scripted dialog.) In both conditions, the researcher encouraged the child to draw for 2 min.

For Choice 2, the researcher produced a round 1/4-in. reward-style sticker from her pocket sealed inside of a plastic envelope. The researcher explained that they could use the small sticker now, or wait until the researcher returned with a larger number of better stickers to use instead. The researcher then placed the small sealed sticker in the center of the table and left the child alone in the room to wait for 2.5 min. As in Choice 1, the sticker packaging was also difficult-to-open by design: the sticker was glued down and covertly sealed inside the plastic envelope with superglue. This preparation was ultimately unnecessary, however, as children were so occupied with drawing during this delay that they did not examine the sticker. This manipulation was also successful, and all children waited the full 2.5 min. without using the 1/4-in. reward-style sticker. In the unreliable condition, the researcher returned without the promised stickers and provided the following explanation:

> "I'm sorry, but I made a mistake. We don't have any other stickers after all. But why don't you just use this one instead?"

The researcher then offered assistance to the child in opening the sealed sticker package, and then covertly swapped it out for an identical usable version. In the reliable condition, the researcher returned with 5–7 large die-cut stickers

featuring a desirable theme (e.g., Toy Story, Disney Princesses). Unbeknownst to the child, the caretaker selected that set of stickers to be especially desirable in advance of the study. In both conditions, the researcher then encouraged the child to work on their drawing for 2 min.

Thus, children were provided with two sources of evidence that the experimenter—and more generally the testing situation—was either unreliable or reliable.

Marshmallow task.

The marshmallow task immediately followed the two-part art task. Once the table was cleared, the researcher revealed a single marshmallow to the child and provided the following explanation:

> "You finished just in time, because now it's snack time! You have a choice for your snack. You can eat this one marshmallow right now. Or—if you can wait for me to go get more marshmallows from the other room—you can have two marshmallows to eat instead. How does that sound? [Response.] Okay, I'm going to go get more marshmallows from the other room and turn your picture into a cup! You should stay right here in that chair. Can you do that? [Response.] I'll leave this [marsh-mallow] here, and if you haven't eaten it when I come back, you can have two marshmallows instead!"

The researcher placed the marshmallow directly in front of the child, 4 in. from the table's edge. The researcher then quickly collected the art materials and drawing and exited the room. The child was left alone in the room, while under covert observation via webcam, until either they consumed the marshmallow or until 15 min had elapsed. Regardless of whether they waited, each child was ultimately given three additional marshmallows at the conclusion of the study. We note that this final portion of the experimental procedure is slightly different from those used by the studies analyzed in Shoda et al. (1990). Major features of

the delay situation are identical; however we did not require children to explicitly signal their desire to stop waiting before eating the lesser treat. The original paradigms involved training children to expect that the experimenter would return upon use of an explicit signal (e.g., ringing a bell). Since this would necessarily provide children with additional information about the experimenter's reliability (as well as add time and complication to our already lengthy experimental procedure), we omitted it. As an additional benefit, this simplified procedure ensures that even very young children could quickly and easily understand the task.

Coding.

Two naïve coders (who were unaware of the experimental conditions) reviewed blinded videos of children in the marshmallow task and recorded when each child's first taste—a lick or bite—occurred. Judgments were checked against one another to ensure reliability: 78.57% matched exactly, 14.29% differed by 1 s, and 7.14% differed by 2 s. When judgments differed, the later time was used. Coders also quantified excitement by measuring smiling time (s) and assigning a subjective rating of apparent contentedness (on 1–9 scale) at the onset of the waiting period (first 30 s). Additionally, the degree of physical movement (fidgetiness) was measured via a computer script that quantified the mean number of pixel changes across frames during the same 30-s time interval.

Results.

Mean wait-times are shown in Fig. II-1. Because the task was terminated at 15 min, children who had not eaten the marshmallow may have waited longer if the experimental design had permitted. Thus, this analysis is a conservative estimate of the true difference between the two conditions.
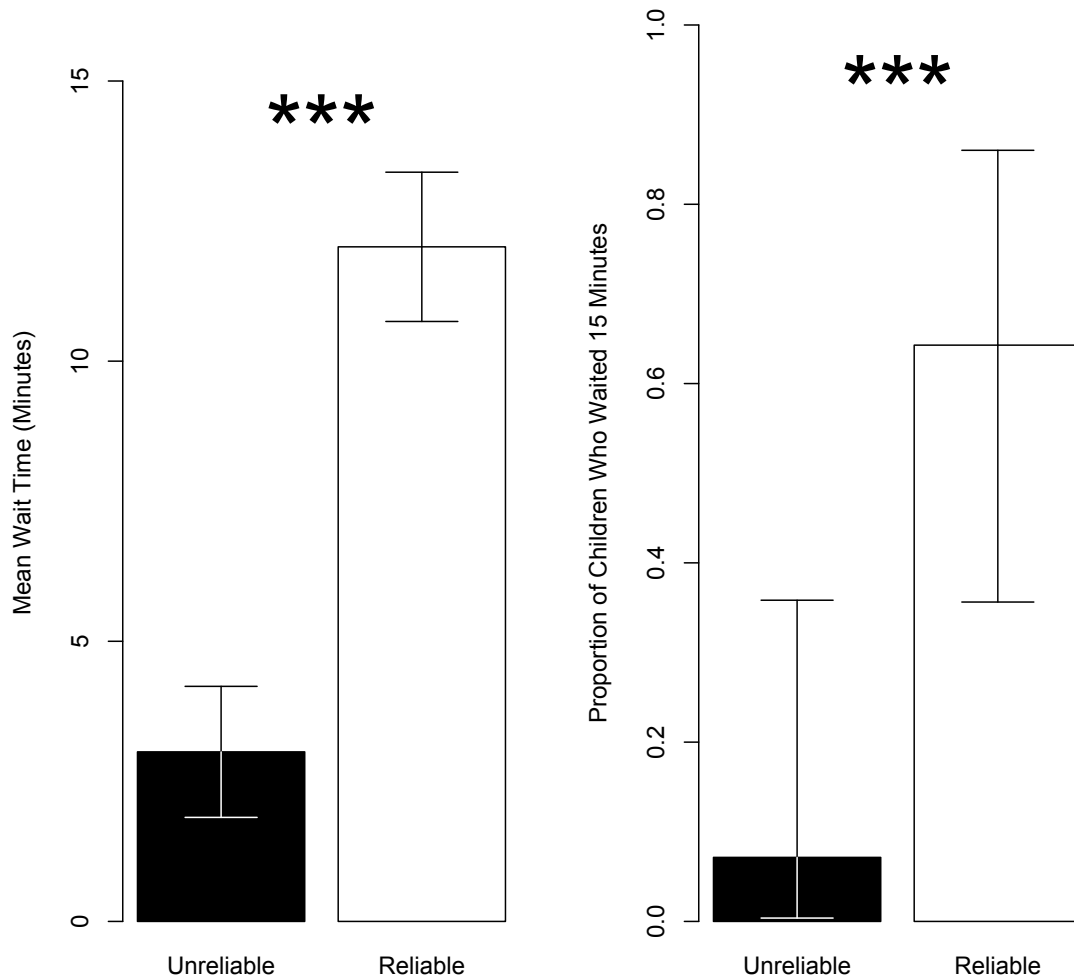
Fig. II-1. *(left)* **Mean Wait-Time of Children in Each Condition.** Error bars show 95% confidence intervals. Children in the unreliable condition waited without eating the marshmallow for a mean duration of 3 minutes and 2 seconds (M = 181.57 s). In contrast, those in the reliable condition waited 12 minutes and 2 seconds (M = 722.43 s). A Wilcoxon signed-rank test found this difference to be highly significant (W = 22.5, $p < 0.0005$). Here, 15 minutes was used as the wait-time for children who did not eat the marshmallow until the research returned, though these children may have actually waited for longer had the experimental design permitted

Fig. II-2. *(right)* **Proportion of Children Who Waited the Full 15 Minutes Without Eating the Marshmallow by Condition.** Error bars show 95% confidence intervals. In the unreliable condition, only 1 of the 14 children (7.1%) waited the full 15 min.; in the reliable condition, 9 out of the 14 children (64.3%) waited. We tested the difference using a two-sample test for equality of proportions with continuity correction at α_{2-tail} = 0.05. The test found it to be highly significant (X^2= 7.6222, df=1, $p < 0.006$).

.

Children in the *unreliable* condition waited without eating the marshmallow for a mean duration of 3 min and 2 s ($M = 181.57$ s). In contrast, children in the *reliable* condition waited 12 min and 2 s ($M = 722.43$ s). A Wilcoxon rank-sum test (also known as a *Mann–Whitney Wilcoxon* or a *Mann–Whitney U*) confirmed that this difference was highly significant ($W = 22.5$, $p < 0.0005$). Thus, children in the *unreliable* condition waited significantly less than those in the *reliable* condition.

We also conducted a binary analysis of whether children waited the entire 15 min without tasting the marshmallow (Fig. II-2). In the *unreliable* condition, only 1 out of the 14 children (7.1%) waited the full 15 min; in the *reliable* condition, however, 9 out of the 14 children (64.3%) waited. A two-sample test for equality of proportions with continuity correction at $\alpha_{2\text{-tail}} = 0.05$ (Newcombe, 1998) was highly significant ($X^2 = 7.62$, $df = 1$, $p < 0.006$). Thus, children in the *unreliable* condition were significantly less likely to wait the full 15 min than those in the *reliable* condition.

Additionally, we performed a linear regression with age and gender as predictors, controlling for condition. Neither factor—age ($\beta = 8.57$, t = 1.29, p > 0.20) nor gender (($\beta = -11.63$, t = -0.10, p > 0.92)—was significant in our sample. Detailed subject data appear in Appendix II-2.

Since these results might alternatively be explained by a difference in mood across the two groups (e.g., by differently induced levels of either frustration or excitement), an analysis of the three relevant measures—apparent contentedness, smiling, and fidgetiness—was conducted (see Appendix II-3). Results suggested that these variables did not vary systematically across the two conditions.

Discussion.

The results of our study indicate that young children's performance on sustained delay-of-gratification tasks can be strongly influenced by rational decision-making processes. If self-control capacity differences were the primary causal mechanism implicated in children's wait-times, then information about the reliability of the environment should not have affected them. If deficiencies in self-control caused children to eat treats early, then one would expect such deficiencies to be present in the *reliable* condition as well as in the *unreliable* condition. The effect we observed is consistent with converging evidence that young children are sensitive to uncertainty about future rewards (Fawcett et al., 2012; Mahrer, 1956; McGuire & Kable, 2012).

The resulting effect of our experimental manipulation was quite robust ($\Delta M_{delay} = 9$ min, $p < 0.0005$). Importantly, while there were small procedural differences between our study and past studies, children—age and gender-matched to the current study—who faced similar choices without prior explicit evidence of experimenter reliability waited for around 6 min (e.g., 6.08 min in Shoda et al. (1990)[1] and 5.71 min in Mischel & Ebbesen (1970)[2]. When we manipulated experimenter reliability, children waited twice that long in the reliable condition (12.03 min), and half as long in the unreliable condition (3.02 min). While further work will be required to explicitly test the relative contributions of different factors, preliminary comparisons suggest that the influence of a children's beliefs about the reliability of the world is at least comparable to their capacity for self-control[3].

---

[1] Condition: exposed reward, no ideation instructions.

[2] Condition: immediate reward.

[3] Two additional manipulation results from Shoda et al. (1990) that may inform relative effect-size estimates: (1) obscuring visual contact with the rewards during the wait (attention manipulation)

To be clear, our data do not demonstrate that self-control is irrelevant in explaining the variance in children's wait-times on the original marshmallow task studies. They do, however, strongly indicate that it is premature to conclude that most of the observed variance—and the longitudinal correlation between wait-times and later life outcomes—is due to differences in individuals' self-control capacities. Rather, an unreliable worldview, in addition to self-control, may be causally related to later life outcomes, as already suggested by an existing body of evidence (e.g., Barnes & Farrell, 1992; Smyke, Dumitrescu, & Zeanah, 2002).

Conclusions.

We demonstrated that children's sustained decisions to wait for a greater reward rather than quickly taking a lesser reward are strongly influenced by the reliability of the environment (in this case, the reliability of the researcher's verbal assurances). More broadly, we have shown that young children's performance on delay-of-gratification tasks can be strongly influenced by an implicit rational decision-making process.

Acknowledgements.

---

increased mean wait-times by 3.75 min and (2) suggesting that children think about the larger reward (ideation strategy) increased them by 2.53 min.

Appendix II. Supplementary material.

Supplementary material associated with this article includes additional scripted dialogue (Appendix II-1), detailed subject data (Appendix II-2), and the analysis of mood variables (Appendix II-3). This material is part of the online version at http://dx.doi.org/10.1016/j.cognition.2012.08.004.

References.

Anderson, J. R., & Milson, R. (1989). Human memory: an adaptive perspective. Psychological Review, 96, 703–719.

Anderson, J. (1991). Cognitive architectures in a rational analysis. In K. VanLehn (Ed.), Architectures for Intelligence (pp. 1–24). Hillsdale, NJ: Lawrence Erlbaum Associates.

Ayduk, O., Mendoza-Denton, R., Mischel, W., Downey, G., Peake, P. K., & Rodriguez, M. (2000). Regulating the interpersonal self: Strategic self-regulation for coping with rejection sensitivity. Journal of Personality and Social Psychology, 79, 776–792.

Barnes, G. M., & Farrell, M. P. (1992). Parental support and control as predictors of adolescent drinking, delinquency, and related problem behaviors. Journal of Marriage and Family, 54(4), 763–776.

Baumeister, R. F., Heatherton, T. F., & Tice, D. M. (1994). Losing control: How and why people fail at self-regulation. San Diego, CA: Academic Press.

Denison, S., & Xu, F. (2010). Twelve- to 14-month-old infants can predict single-event probability with large set sizes. Developmental Science, 13, 798–803.

Fawcett, T. W., McNamara, J. M., & Houston, A. I. (2012). When is it adaptive to be patient? A general framework for evaluating delayed rewards. Behavioural Processes, 89, 128–136.

Gergely, G., Bekkering, H., & Király, I. (2002). Rational imitation in preverbal infants. Nature, 415, 755.

Goleman, D. (1995). Emotional intelligence. New York: Bantam Books.

Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. Psychological Review, 111(1), 1–31.

Gweon, H., & Schulz, L. E. (2011). 16-Month-olds rationally infer causes of failed actions. Science, 332(6037), 1524.

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(5), e36399.

Loewenstein, G. F., Read, D., & Baumeister, R. (2003). Time and decision: Economic and psychological perspectives on intertemporal choice. New York: Russell Sage Foundation.

Mahrer, A. H. (1956). The role of expectancy in delayed reinforcement. Journal of Experimental Psychology, 52, 101–106.

Marcovitch, S., & Zelazo, P. D. (1999). The A-not-B error: Results from a logistic meta-analysis. Child Development, 70, 1297–1313.

Marr, D. (1982). Vision. San Francisco: W.H. Freeman.

McGuire, J. T., & Kable, J. W. (2012). Decision makers calibrate behavioral persistence on the basis of time-interval experience. Cognition, 124(2), 216–226.

Mischel, W. (1961). Father-absence and delay of gratification: Cross-cultural comparisons. Journal of Abnormal and Social Psychology, 63, 116–124.

Mischel, W. (1974). Processes in delay of gratification. In L. Berkowitz (Ed.), Advances in experimental social psychology (Vol. 7, pp. 249–292). New York: Academic Press.

Mischel, W., & Ebbesen, E. B. (1970). Attention in delay of gratification. Journal of Personality and Social Psychology, 21, 204–218.

Mischel, W., Shoda, Y., & Rodriguez. (1989). Delay of gratification in children. Science, 244, 933–938 [Article later reprinted as a chapter in G. Loewenstein & J. Elster (Eds.). Choice over time. New York: Russel-Sage, 1992.].

Mischel, W., & Metzner, R. (1962). Preference for delayed reward as a function of age, intelligence, and length of delay interval. Journal of Abnormal and Social Psychology, 64(6), 425–431.

Mischel, W., Shoda, Y., & Peake, P. K. (1988). The nature of adolescent competencies predicted by preschool delay of gratification. Journal of Personality and Social Psychology, 54, 687–696.

Mischel, W., & Staub, E. (1965). Effects of expectancy on working and waiting for larger rewards. Journal of Personality and Social Psychology, 2, 625–633.

Newcombe, R. G. (1998). Interval estimation for the difference between independent proportions: Comparison of eleven methods. Statistics in Medicine, 17, 873–890.

Piaget, J. (1954). The construction of reality in the child. Oxford: Basic Books.

Piaget, J. (1970). Piaget's theory. In P. Mussen (Ed.). Handbook of child psychology (Vol. 1). New York: Wiley.

Piantadosi, S. T., Tenenbaum, J. B., & Goodman, N. D. (2012). Bootstrapping in a language of thought: A formal model of numerical concept learning. Cognition, 123(2), 199–217.

Romer, D., Duckworth, A. L., Sznitzman, S., & Park, S. (2010). Can adolescents learn self-control? Delay of gratification in the development of control over risk taking. Prevention Science, 11(3), 319–330.

Shoda, Y., Mischel, W., & Peake, P. K. (1990). Predicting adolescent cognitive and social competence from preschool delay of gratification: Identifying diagnostic conditions. Developmental Psychology, 26, 978–986.

Smyke, A. T., Dumitrescu, A., & Zeanah, C. H. (2002). Attachment disturbances in young children. I: The continuum of caretaking casualty. Journal of the American Academy of Child & Adolescent Psychiatry, 41(8), 972–982.

Téglás, E., Vul, E., Vittorio, G., Gonzalez, M., Tenenbaum, J. B., & Bonatti, L.L. (2011). Pure reasoning in 12-month-olds as probabilistic inference. Science, 332(6033), 1054–1059.

Xu, F., Dewar, K., & Perfors, A. (2009). Induction, overhypotheses, and the shape bias: Some arguments and evidence for rational constructivism. In B. M. Hood & L. Santos (Eds.), The origins of object knowledge (pp. 263–284). Oxford University Press.

Xu, F., & Tenenbaum, J. B. (2007). Sensitivity to sampling in Bayesian word learning. Developmental Science, 10, 288–297.

III. The Goldilocks Effect.

Human infants allocate attention to visual sequences that are neither too simple nor too complex.

Celeste Kidd, Steve Piantadosi, & Richard N. Aslin

Human infants, like immature members of any species, must be highly selective in sampling information from their environment to learn efficiently. Failure to be selective would waste precious computational resources on material that is already known (too simple) or unknowable (too complex). In two experiments with 7- and 8-month-olds, we measure infants' visual attention to sequences of events varying in complexity, as determined by an ideal learner model. Infants' probability of looking away was greatest on stimulus items whose complexity (negative log probability) according to the model was either very low or very high. These results suggest a principle of infant attention that may have broad applicability: infants implicitly seek to maintain intermediate rates of information absorption and avoid wasting cognitive resources on overly simple or overly complex events.

Introduction.

Human infants face two daunting problems as they begin to learn about their surroundings. First, they enter the postnatal world with only rudimentary mechanisms–provided by their evolutionary heritage–for interpreting environmental information. Second, the potential information available in the environment is both voluminous and complex. These two problems led William James to coin his famous phrase about "the blooming, buzzing confusion" that confronts the newborn (James, 1980). Nonetheless, infants show remarkable feats of learning, beginning in the last trimester of fetal life, continuing through the perinatal period, and accelerating through infancy and early childhood (DeCasper & Fifer, 1980; Rovee-Collier et al., 1980; Siqueland & De Lucia, 1969; Stevenson 1972). Infants are able to extract the statistical properties of their environment in a diverse array of learning tasks and domains, including sounds, words, people, shapes, and objects (Fiser & Aslin, 2002; Kirkham, Slemmer, & Johnson, 2002; Maye, Werker, Gerken, 2002; Saffran, Aslin, & Newport, 1996; Saffran, Johnson, Aslin, & Newport, 1999; Saylor, Baldwin, Baird, & LaBounty, 2007). But how is it that infants are able to learn efficiently in such a complex environment? One solution is to have a small set of innate biases; for example, seeking to look at and listen to biologically significant stimuli such as faces and speech. However, innate biases alone cannot be the solution for the vast majority of stimuli from which infants must learn. Given the slow time-course of evolution, we also need general purpose learning mechanisms to deal with a changing environment and with classes of stimuli that could not plausibly be processed by a small set of specialized mechanisms.

Here, we focus on this general-purpose learning mechanism by avoiding the use of special stimuli and asking whether infants deploy a sensible (and likely implicit) strategy for allocating attention to arbitrary, neutral stimuli. Our goal is to determine whether infants are biased to gather information from the environment in a principled way that serves as a key component of an efficient learning mechanism (Berlyne, 1960; Piaget, 1970). Specifically, we provide evidence that infants avoid spending time examining stimuli that are either too simple (highly predictable) or too complex (highly unexpected) according to their implicit beliefs about the probabilistic structure of events in the world. Rather, infants allocate their greatest amount of attention to events of intermediate surprisingness–events that are likely to have just enough complexity so that they are interesting, but not so much that they cannot be understood. This approach builds on a longstanding tradition in developmental psychology, as exemplified by Piaget (*e.g.,* Piaget, 1970). He argued that when children are confronted with a new piece of information, they initially attempt to incorporate it within their existing knowledge structures through a process of assimilation. When this is not possible, children either fail to learn new structures (and move on to sample other information) or they adapt by creating new knowledge structures, a process he called accommodation.

Piaget had no objective measure of assimilation or accommodation; they remained hypothetical constructs. However, in subsequent research, a proxy for these theoretical constructs centered on the relative duration of visual attention to objects or events varying in complexity or familiarity. Many researchers have speculated about what underlying mental operations are indexed by infants' looking times or attentional patterns (Fantz, 1964; for review: Aslin, 2007). The generally accepted view is that looking times reflect some combination of (a) stimulus-driven attention, (b) memory of past stimuli, and (c) comparison between

the current and the past stimuli. If infants are presented with an already familiar stimulus, they prefer it over a novel stimulus, but quickly tire of it after a brief period of re-familiarization (habituation), and subsequently show preferences for novel stimuli. Similarly, if repeatedly exposed to an initially novel stimulus, infant looking times decline and then recover to the presentation of another novel (i.e., completely unfamiliar) stimulus. Theoretical accounts for these familiarity and novelty preferences all share a common theme: As infants attempt to encode various features of a visual stimulus, the efficiency or depth of this encoding process determines their subsequent preferences. Familiarity preferences arise when infants have not yet completed encoding the familiar stimulus into memory, or when the novel stimulus is too dissimilar from the infants' existing mental representations to be readily encoded (Dember & Earl, 1957; Hunter & Ames, 1988; Kinney & Kagan, 1976; Roder, Bushnell, & Sasseville, 2000; Rose, Gottfried, Melloy-Carminar, & Bridger, 1982; Sokolov, 1963; Wagner & Sakovits, 1986).

However, these theories lacked an objective measure of the relevant independent variable–an event's complexity or relationship to existing representations. Instead, researchers overwhelmingly relied on qualitative judgments of stimulus complexity to select materials to test infants' visual preferences. These qualitative judgments relied on inferences about infants' existing mental representations, to which researchers had no direct access. With no reasonable way of modeling infants' existing representations, it was impossible to quantitatively measure the complexity of the information conveyed by a particular stimulus. Thus, researchers had only post hoc estimates of stimulus complexity– those obtained by measuring the very patterns of visual preferences that the theories were designed to predict. Two exceptions are Civan, Teller & Palmer (2005) and Kaldy, Blaser, & Leslie (2006) in that both papers quantified the

perceptual salience of visual stimuli in order to effectively demonstrate its importance in eliciting infants' preferences for novel versus familiar stimuli.

We overcome these problems by formalizing a notion of stimulus complexity and behaviorally testing the relationship between complexity and infants' probability of looking away at each successive point in a sequence of events. We assume that at each point in the experiment–and in everyday life–infants have used observed data to form probabilistic expectations about what events are likely and unlikely to be observed next (Téglás et al., 2011; Xu & Garcia, 2008). We model these expectations using an idealized observer model of our experimental stimuli. We then measure complexity as the negative log probability of an event according to this idealized model. This measure quantifies each event's information content (Shannon, 1948). (This measure has also been called surprisal (Tribus, 1961), since it may also be interpreted as representing the "surprise" of seeing the outcome.) We show that infants preferentially look away at events that are either very simple (high probability) or very complex (low probability), according to the idealized model. Intuitively, high probability events convey little information–infants' attentional resources are best spent elsewhere. Low probability events may indicate that the observed stimuli are unlearnable, unstructured, or difficult to use predictively in the future. Negative log probability also quantifies the number of bits of information an ideal observer would require to encode that sequence of events in memory. Thus, infants may avoid stimuli that require encoding too much information or information that could only be extracted by prolonged attention to rare events, thereby incurring a higher processing cost than shifting attention to less complex events.

Experiment and Modeling Approach.

The behavioral experiment measured the point, in a sequence of events, when an infant looked away from a visual display. The displayed stimuli were easily captured by a simple statistical model. In Experiment 1, we presented each infant with 42 unique animated displays, each featuring one of 42 uniquely colored and patterned boxes occluding one of 42 unique familiar objects (e.g., a ball). Each scene display began with the occluder rising and falling, thus appearing to reveal and then re-obscure the object hidden behind it (Fig. III-1*a* and **Video S1**). To maintain infants' attention early in the experiment, the first reveal always showed an object. For example, a blue polka-dotted occluder might rise to reveal a toy fire truck. On subsequent reveals, the same object appeared in the box according to some probability randomly assigned to that trial. For example, if a trial were associated with the probability of 0.3, then 30% of the time an object would be present behind the box. Probabilities ranged from 0 to 1 in increments of 0.05 (i.e., 0.0, 0.05, 0.1, 0.15, etc.), such that there were 21 possible probabilities-of-appearance that could be associated with an object on a particular trial. The sequences of object reveals thus varied in terms of their information-theoretic properties: some events in a sequence were highly predictable (e.g., a ball appears still in the box after having appeared on each of ten previous reveals), and others were less predictable (e.g., a rattle appears to have disappeared from within the box after having appeared on each of the ten previous reveals). The objects, boxes, and order in which the probabilities-of-appearance were presented were randomized across infants, and each of the 21 probabilities-of-appearance occurred twice (for a total of 42 trials). Each animated sequence of events continued until the infant met the look-away criterion, which was defined as gaze directed off-screen for greater than 1 consecutive second (see **Video S3** for look-away example). To address uncertainty about infants' mental representations and

their age-related or uniquely individual processing speeds and biases for stimulus salience, we exhaustively randomized and counterbalanced all of these extraneous variables (e.g., sequence order, object identity, object familiarity, spatial location).
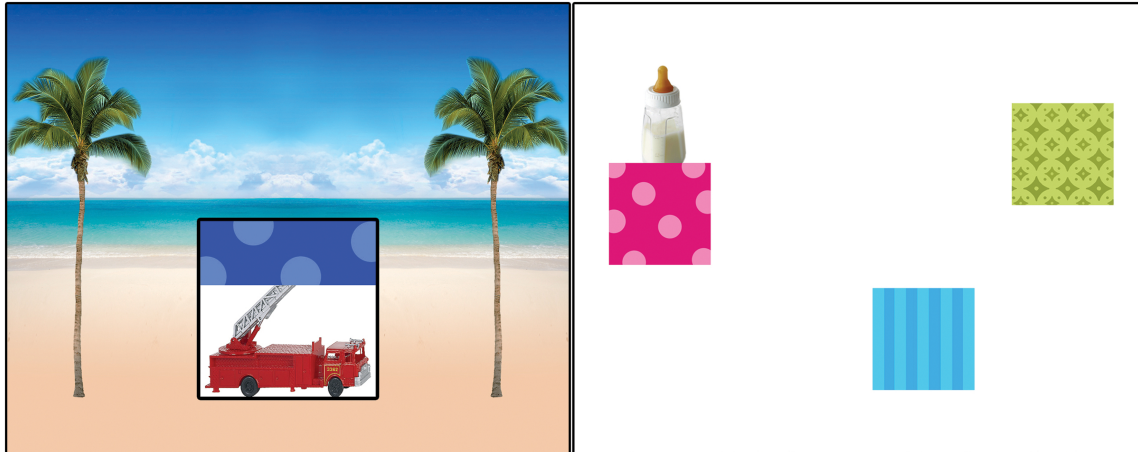
Fig. III-1. Examples of Visual Displays Used in Experiments 1 and 2. *(a)* The object (e.g., a toy fire truck) in the box for Experiment 1 was revealed (or not) by up-down animation of an occluder (e.g., a blue polka-dotted box). *(b)* In Experiment 2, one of three unique objects (e.g., a baby bottle) popped up from behind one of three highly distinctive boxes. Also, see Videos S1 and S2 for examples of animated displays used in these studies.

We modeled the sequences of reveals using a Markov Dirichlet-multinomial model (MDM). The Dirichlet-multinomial is a general-purpose statistical model that uses observed event counts to compute a posterior distribution for an underlying multinomial distribution on events. The Dirichlet-multinomial makes parametric assumptions about the form of the prior probability and the likelihood of an event and is often used in Bayesian statistics because of its computational simplicity (see Materials and Methods). We apply this to a time-series of events by making a Markov assumption that each event is statistically independent (i.e., not dependent on the ordering of the preceding events). Thus, the model can take some

previously observed sequence of events–corresponding to an individual infant's observations before they have looked away–and compute the probability of every possible next event. We hypothesize that infants' probability of looking away at the next event in a sequence is at least partially determined by the information-theoretic properties of that event, according to the model. Specifically, at each point in a sequence of events, the model assigns each event a probability, and the negative log of this probability provides a natural information-theoretic measure of the complexity of the next event according to the model's current expectations about which events are likely.

Fig. III-2 illustrates the logic of the experiment and analysis. In the first example, the observer sees a sequence of four *A* events in a row. In this case, the observed data consist of entirely *A*'s. These data are combined with the prior– essentially a smoothing term to avoid zero probabilities–to form an updated posterior belief with high probability of A but non-zero probability of B ("Updated belief" column). The complexity (negative log probability) of the next event is determined using this posterior, which represents the model's updated belief about the true distribution of events. Thus, if the next event is an A–an event that is highly likely according to the model's posterior–the complexity of that event would be low (i.e., the event would be highly predictable according to the model). We hypothesize that infants would be more likely to look away at this event. Conversely, if the previous observations assign A very low probability (second example), A will have very high complexity (i.e., the event would be highly surprising according to the model) and infants should terminate the sequence of events by looking away. If the previous observations make A moderately likely (third example), the occurrence of an A event will convey a "Goldilocks" amount

of information, leading infants to be less likely to look away. If infants do not look away, then the modeling step is repeated for the next item in the sequence. This means that infants may look away at different points in different sequences, but we predict systematicity in these look-aways: regardless of how far into a sequence an infant has made it without looking away, their probability of looking away on the next object will depend on its complexity, conditioning on all previous observations.
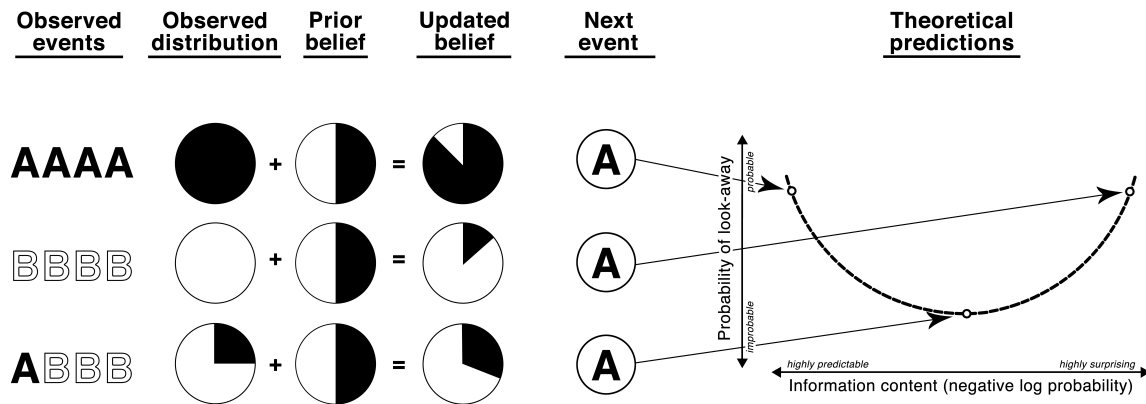
Fig. III-2. Ideal Observer Model Schematic. Schematic showing several example event sequences and how the Ideal Observer Model combines observed events with a simple prior to form expectations about upcoming events. The next event then conveys some amount of information according to these probabilistic expectations, which is related to infants' probability of look-away at a specific next event by a U-shaped function.

We note that this type of modeling and analysis contrasts with most previous infant studies, which typically tested for differences in overall mean looking times. Here, we are predicting a binary outcome (whether an infant looks away) at each individual event in the sequence. This is a more precise prediction based on probabilities computed on-line.

Results.

EXPERIMENT 1.

Fig. III-3 shows infants' probability of looking away, as a function of that event's negative log probability according to the model, and collapsing across infants, sequences, and sequence positions. The diamonds show raw probability of look-away, binning complexity into 5 discrete bins. The curve represents the fit of a Generalized Additive Model (Hastie & Tibshirani, 1990), which attempts to find a smooth relationship between complexity and look-away probability. This figure shows a U-shaped relationship between infant look-away probability and the on-line model-based estimate of complexity, with infants looking away from events that are especially predictable or especially surprising. There is a "Goldilocks" value of complexity around 1.25 bits, corresponding to infants' preferred information rate in this task.
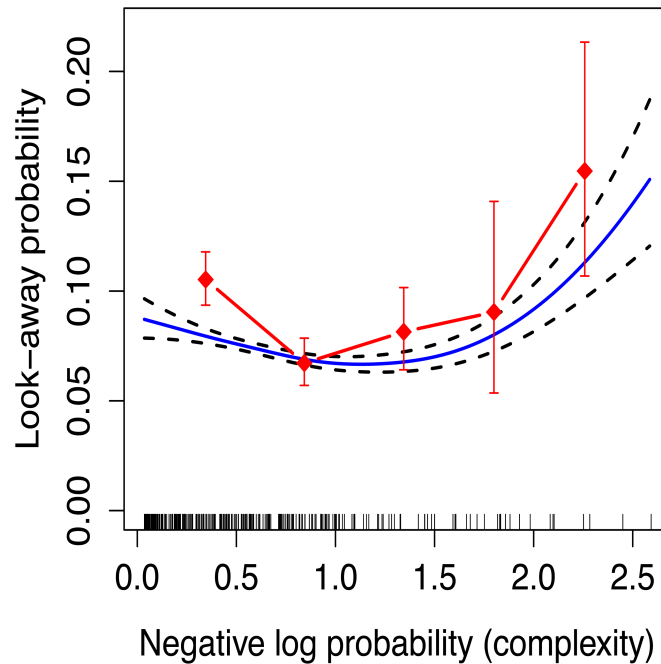
Fig. III-3. U-Shaped Curve for Single-Box Display Used in Experiment 1. The solid curve represents the fit of a Generalized Additive Model (GAM) (Hastie & Tibshirani, 1990) with binomial link function, relating complexity according to the MDM model (*x*-axis) to infants' look-away probability (*y*-axis). The dashed curves show standard errors according to the GAM. The GAM fits include the effect of complexity (negative log probability) and the effect of position in the sequence. Note, the error bars and GAM errors do not take into account subject effects. Vertical spikes on the *x*-axis represent data points collected at each complexity value. The red diamonds represent the raw look-away probabilities binned along the *x*-axis.

Although the plot in Fig. III-3 provides a revealing picture of the relationship between indexes of complexity and looking durations, there are likely other factors that influence when infants will look away from the displays. For instance, low-information and high-information events may tend to occur later in a sequence, after learners have developed expectations about the distributional properties of the events. If infants tend not to look away early, perhaps because they are initially captured by the salience of the display independent of its complexity, they would appear to disprefer low and high complexity. To address this potential confound, we performed a regression analysis that controls for the influence of temporal and other factors on look-away probability. When infants

look away in a trial, they provide no more data for the remainder of the trial. Because of this, such data violate the independence assumptions of standard logistic (or linear) regression. An appropriate model for this kind of data–used primarily in biostatistics to study, for example, predictors of mortality–is known as a survival analysis (Hosmer, Lemeshow, & May, 2008; Klein & Moeschberger, 2003). We used a type of survival analysis, known as a Cox regression, that measures the log linear influence of predictors on look-away probability, while respecting the fact that once infants look away they provide no additional data on the same trial. Importantly, this regression also controls for a baseline look-away distribution, which is fit non-parametrically to the data, thereby removing the influence of an average distribution of looking times before testing the significance of the other predictors. We note that this regression does not include subject effects, but we develop more sophisticated analysis methods that include a range of subject effects in forthcoming work (Piantadosi, Kidd, & Aslin, in press).

We included a number of control covariates that could plausibly influence infant look-aways using a stepwise procedure that only added variables that improved model fit. These variables included whether an object was present, whether the presence of the object was the same as the previous reveal, how many sequences the infant had already observed, and the uncertainty in the model about the correct distribution of events. This was measured by the differential entropy of the multinomial parameters in the MDM model. We also included linear and quadratic complexity terms. To aid in interpretation of the regression coefficients, complexity was standardized before being squared (i.e., it was shifted and scaled to have mean 0 and standard deviation 1 to test for a significant quadratic trend of complexity on look-aways. This stepwise procedure revealed a significant effect only for squared complexity (β=0.052, z=1.969, p<0.05), and no other variables (see Table III-1). This indicates that the U-shape observed in Fig. III-3 is

statistically significant, even after controlling for an overall baseline look-away distribution and the other potentially confounding variables (see *Materials and Methods*). The magnitude of this effect can be understood by considering $e^\beta = 1.05$, which is the *factor* that the baseline look-away probability is multiplied by for each increase in squared surprisal of one standard deviation from the overall mean in the experiment. This effect is relatively small, though statistically reliable.

Table III-1. Cox Regression Coefficients.

Covariate	Coefficient	exp(coefficient)	Standard error	Z-statistic	P-value
Experiment 1					
Squared complexity	0.052	1.05	0.026	1.969	0.049
Experiment 2 - Non-transitional model					
Complexity	−0.216	0.805	0.094	−2.29	0.022
Squared complexity	0.269	1.308	0.109	2.47	0.013
Trial number	0.029	1.030	0.007	3.99	$6.5.10^{-5}$
Model uncertainty	0.261	1.298	0.174	1.50	0.13
Experiment 2 - Transitional model					
Squared complexity	0.356	1.43	0.084	4.27	$1.9.10^{-5}$
Trial number	0.027	1.03	0.007	3.65	$2.7.10^{-4}$
First appearance	0.500	1.64	0.272	1.82	0.069

All variables found in Experiments 1 and 2 that were added by the stepwise procedure. Note that some non-significant variables are added because the stepwise comparison is based on the Akaike information criterion [40]. These results reveal significant quadratic effects of complexity in both experiments. Complexity and squared complexity were shifted and scaled to have mean of 0 and standard deviation of 1 before being entered into the regression.

EXPERIMENT 2.

In Experiment 1, objects were either present or absent from behind a single occluder. Perhaps a more typical context in real life, though, is for different events to occur in a multi-object scene, thereby allowing infants' attention to be attracted to both individual events and transitions between events. In Experiment 2, we presented each infant with 32 unique sequential-event displays (Fig. III-1*b* and **Video S2**). Each display presented an animated scene consisting of three uniquely patterned boxes, each concealing a unique familiar object (e.g., a cookie). The locations of the three boxes for a given sequence were chosen randomly but remained static throughout a scene. The box locations were randomly shuffled between event sequences, but no more than two boxes appeared on either half of the screen. Neither the patterns on the boxes nor the objects were repeated across event sequences so that each object-box pair was independent and unique. Each event in a sequence consisted of an object that popped out of a box, and then back into the box. Each event lasted 2 seconds in total duration (1-second "pop-up", 1-second "pop-down"). Events were presented sequentially with no overlap or delay. The same 32 event sequences were presented to every infant. However, the objects, boxes, and order in which the 32 event sequences were presented were randomized across infants. This design ensured that differences in looking times across event sequences were not driven by differences in scene items or presentation order. Each animated sequence of events continued until the infant met the look-away criterion, which was defined as gaze directed off-screen for greater than 1 consecutive second.

Results from Experiment 2 are shown in Fig. III-4. As in Experiment 1, there is a U-shaped relationship between look-away probability and complexity, as measured by the same MDM model (assuming event independence) used in Experiment 1. The Cox regression for Experiment 2 included all of the covariates

used in Experiment 1, except whether an object was present, since there was always an object popping up from behind one of the three boxes. However, because there are three different box-object pairs in each scene, we also included covariates measuring whether the current event is the first time an object has appeared from behind a box, and a factor measuring how many objects have not yet popped up. Results of this analysis are shown in Table III-1. As in Experiment 1, this analysis revealed significant effects of squared complexity (β=0.269, z=2.47, p<0.013). Here, e^β=1.308, meaning that each increase of squared complexity 1 standard deviation from the mean resulted in a look-away probability that was a factor of 1.31 times greater. This is a much larger effect than that found in Experiment 1. There was also a significant linear effect of complexity, indicating that the U is not symmetric about the mean (β=−0.216, z= −2.291, p< 0.05), and an effect of trial number, likely representing effects of fatigue (β=0.029, z= 3.994, p< 0.001), although this is small compared to the complexity effects (e^β=1.03).
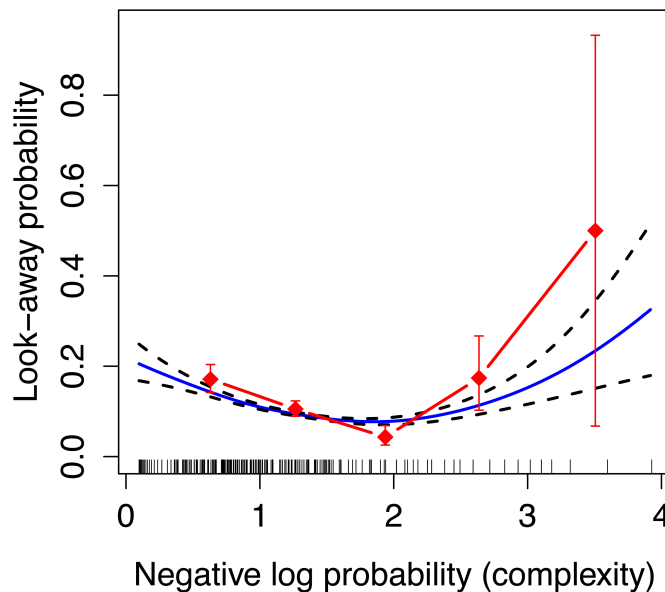
Fig. III-4. U-Shaped Curve for Three-Box Display Used in Experiment 2 (Non-Transitional MDM). The solid curve represents the fit of a GAM, relating complexity as measured by the non-transitional MDM (assuming event independence) to look-away probability. Dashed curves show GAM standard errors. The GAM fits include the effect of complexity (negative log probability) and the effect of position in the sequence. Note, the error bars and GAM errors do not take into account subject effects. Vertical spikes on the *x*-axis

represent data points collected at each complexity value. The red diamonds represent the raw look-away probabilities binned along the *x*-axis.

We also applied the MDM model to the data from Experiment 2 under an assumption of event-order dependence. That is, instead of treating every event as independent, we examined whether look-aways were predicted by the immediately preceding event (i.e., a transitional model). Fig. III-5 shows that a U-shaped function also describes this transitional model, and the Cox regression confirms that this effect is highly significant ($\beta=0.356$, $z=4.27$, $p<0.001$). This analysis also revealed an effect of trial-number ($\beta=0.027$, $z=3.645$, $p<0.001$).
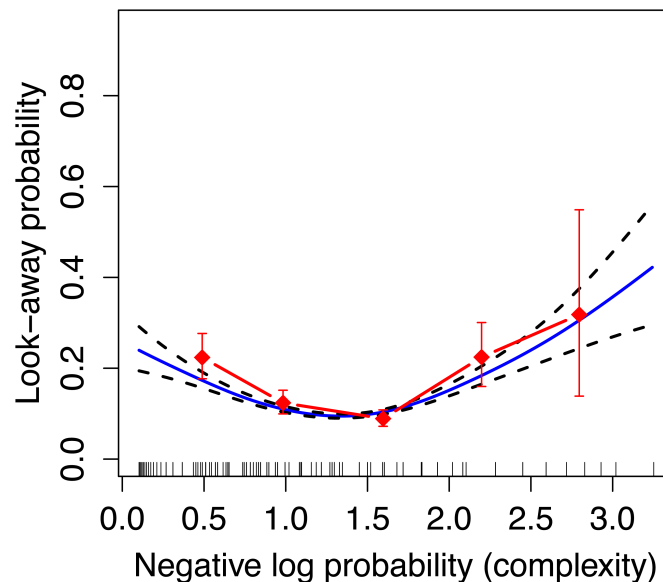
Fig. III-5. U-Shaped Curve for Three-Box Display Used in Experiment 2 (Transitional MDM). The solid curve represents the fit of a GAM, relating complexity as measured by the transitional MDM to look-away probability. Dashed curves show GAM standard errors. The GAM fits include the effect of complexity (negative log probability) and the effect of position in the sequence. Note, the error bars and the GAM errors do not take into account subject effects. Vertical spikes on the *x*-axis represent data points collected at each complexity value. The red diamonds represent the raw look-away probabilities binned along the *x*-axis.

Finally, one can ask which of the two models better accounts for infants' behavior on the task in Experiment 2. The predictions of the transitional and non-

transitional models are difficult to distinguish because they are closely related: Complexity of both models is correlated at $R = 0.62$ ($p < 0.001$). However, if both are entered into a Cox regression along with all variables found to be significant, the transitional complexity is significant ($\beta = 0.289$, $p < 0.01$), but the non-transitional complexity is not ($\beta = 0.015$, $p > 0.84$). This provides strong evidence that infants track transitional probabilities, but the null result for the non-transitional model is difficult to interpret due to its correlation with the transitional model and the noise inherent in infant data.

Discussion.

The results of the experiments reported here have important implications for two interrelated hypotheses concerning infants' attention. First, infants behave as if they are employing a principled inferential process for learning about events in the world. The particular MDM model used in our analyses took as inputs a series of observed events or transitions between events to form probabilistic expectations about what events are most likely to occur in the future. The model was necessary to determine what complexity a set of stimulus events conveys to an ideal observer. A failure of either of these components–the probabilistic model or the linking assumption that maps level of complexity onto looking times–would have yielded null results.

Second, infants appear to allocate their attention in order to maintain an intermediate level of complexity. A powerful feature of our analyses was the ability, via the Cox regression, to control for potential confounds such as the number of items that have not appeared yet, item repeats, and an arbitrary baseline distribution of look-away probabilities. To our knowledge, the hypothesis that infants prefer a particular level of information has not been tested while

controlling for these other variables, and our analyses therefore provide several methodological advances. Rather than predicting infants' average looking time to a stimulus, our analyses predicted the precise event in a sequence when an infant would terminate (i.e., look away from) the display. Although others have observed U-shaped behavior in infants under some circumstances, our results provide the first evidence that the information-theoretic properties of a formal model provide a significant predictor of infant look-aways, over and above the effects of other variables, for a large set of arbitrary, neutral visual stimuli. Interestingly, this U-shaped pattern is similar to those obtained with many earlier models of visual attention based on depth of processing or difficulty of encoding the stimulus (e.g., Hunter & Ames, 1988; Kinney & Kagan; 1976; Roder, Bushnell, & Sasseville, 2000). This could indicate that while earlier models did not computationally define the stimulus properties they hypothesized as the mediators of infant looking times (i.e., *complexity*), the properties they explored are nevertheless relevant in guiding infants' visual attention. Our results also provide a formal account for why infants show novelty preferences (when two test stimuli fall on the left half of the U-shaped function, the stimulus with greater complexity elicits more attention) or familiarity preferences (when two test stimuli fall on the right half of the U-shaped function, the stimulus with lesser complexity elicits more attention).

Similar hypotheses about how adults allocate their limited resources in the language domain—for example, those supporting a uniform information principle (Aylett & Turk, 2004; Genzel & Charniak, 2002; Piantadosi, Tily, & Gibson, 2011; Jaeger, 2010; Levy & Jaeger, 2007)—may suggest that what we have observed in infants reflects a ubiquitous constraint across domains and developmental levels. In addition, other theories propose that learners allocate attention to stimuli containing just the right level of complexity because optimal complexity triggers just the right amount of "arousal" in the learner (Yerkes &

Dodson, 1908). The U-shaped function may result from the basic response properties of neural systems (Turk-Browne, Scholl, & Chun, 2008), although determining the precise mechanism will require further research.

In summary, our findings are consistent with theories that suggest infants actively seek to maintain an intermediate level of information absorption, avoiding allocating cognitive resources to either overly predictable or overly surprising events. It is important to note that we are not claiming that this Goldilocks effect is the only factor in infants' allocation of attention. Certainly, there are species-typical preferences and effects of learning that can dominate infants' attentional behavior. We argue that when these other factors are controlled for, there remains a significant U-shaped effect of complexity that is well accounted for by our model. Further investigation is required to determine how infants' preference for intermediate levels of information affects the outcome of learning, either by enhancing the rate of learning or its asymptotic level.

Materials and Methods.

ETHICS STATEMENT.

All research was approved by the Research Subjects Review Board at the University of Rochester (protocol RSRB00024570). Parents volunteering their infants for the study were fully informed of the study procedures and completed written informed consent and permission forms in advance of the study.

VISUAL STIMULI.

In Experiments 1 and 2, we presented infants with animated displays depicting event sequences varying in their predictability. All displays featured uniquely colored and patterned boxes (e.g., pink polka dots) that were animated to reveal unique familiar objects (e.g., a ball; see **Videos S1** and **S2** for examples). A Matlab script was used to generate each of the animated displays. Neither the boxes nor the objects were repeated across event sequences so that each object-box pair was independent and unique. The objects, boxes, and order in which the event sequences were presented were also randomized across infants. This design ensured that differences in looking time across event sequences were not driven by differences in scene items or presentation order.

In Experiment 1, each animated sequence featured one unique object occluded by one box. The box opened (1 second) and closed (1 second) repeatedly, each time revealing the contents of the box. The object always appeared in the box on the first reveal event. On subsequent reveal events, the object was either present or absent depending on the predictability of the event sequence selected for that trial (a value between 0 and 1). So, for example, a single trial might feature a purple striped box occluding a small toy train with a probability-of-appearance of 0.5. The sequence of events (object appears = 1, empty box = 0) might be: 1, 1, 0, 1, 0, 1, 0, 1, 0, 0. The reveals were presented sequentially with no overlap or delay. There were 21 unique probabilities-of-appearance (increments of 0.05 between 0 and 1, e.g., 0, 0.05, 0.1, 0.15…) and all were presented to each infant twice (42 trials in total) in a random order.

In Experiment 2, each animated display featured three boxes of three unique colors and patterns (e.g., yellow stripes, blue polka dots, green stars), each concealing a unique object (e.g., a cookie, a spoon, a car). The locations of the three boxes for a given sequence were chosen randomly but remained static

throughout a scene. The box locations were randomly shuffled on the screen between event sequences, with the constraint that no more than two boxes appeared on either half of the screen. Each event in a sequence consisted of one of the three unique objects popping out from behind one of the three boxes (1 second), and then back into the box (1 second). Thus, the total duration of each event was 2 seconds, and events were presented sequentially with no overlap or delay. There were 32 unique event sequences that varied in the probability that each of the three objects appeared from behind their respective occluding boxes. Some sequences were simple (e.g., *A, A, A, A, A, A, …*), while others were more complex (e.g., *A, B, A, B, A, C, …*). All event sequences were presented to each infant (32 trials in total).

METHODS.

The procedures for Experiments 1 and 2 were identical, with two exceptions: the type of displays used (single-box in Experiment 1 and three-box in Experiment 2) and the number of trials presented to each infant (42 in Experiment 1 and 32 in Experiment 2). Each infant was seated on his or her parent's lap in front of a table-mounted Tobii 1750 eye-tracker. The infant was positioned such that his or her eyes were approximately 23 inches from the monitor, the recommended distance for accurate eye-tracking. At this viewing distance, the 17-inch LCD screen subtended 24×32 degrees of visual angle. Each of the three boxes was 5×5 degrees. To prevent parental influence on the infant's behavior, the parent holding the infant was asked to wear headphones playing music, wear a visor, lower their eyes, and abstain from interacting with their infant throughout the experiment.

Each trial was preceded by an animation designed to attract the infant's attention to the center of the screen–a laughing and cooing baby. Once the infant

looked at the attention-getter, an experimenter who was observing remotely pushed a button to start the trial. For each trial, an animated scene—featuring a single box in Experiment 1, or three boxes in Experiment 2—was played. The animated sequences of reveal events continued until the infant looked away continuously for 1 second, or until the sequence timed out at 60 seconds. The 1-second look-away criterion for trial termination was automatically determined by the Tobii eye-tracking software. If the infant looked continuously for the entire 60-second sequence, the trial was automatically labeled as a "time out" and discarded before the analysis (2.4% of trials in Experiment 1, 5.4% of trials in Experiment 2). If the trial was terminated before the infant actually looked away, the trial was labeled by an experimenter as a "false stop" and also discarded. False stops, as determined by a separate video recording of the infant's face, occurred as a result of the Tobii software being unable to detect the infant's eyes continuously for 1 second, usually due to the infant inadvertently blocking the eye-tracker camera's view of his or her own eyes with head or arm movements (22.1% of trials in Experiment 1, and 20.7% of trials in Experiment 2). Trials in which the infant looked for fewer than four events were also discarded, since it is presumed that too few observations are insufficient for establishing expectations about the distribution of events.

Subjects.

In Experiment 1, 42 infants (*mean* = 7.9 months, *range* = 7.0 - 8.9) were included in the analysis. Forty-four infants were tested; one infant was excluded due to excessive tiredness (he fell asleep within the first few trials and could not be awakened), and one was excluded due to fussiness. In Experiment 2, 30 infants (*mean* = 7.6 months, *range* = 7.0 - 8.8) were tested, and all participating infants completed the study. In both studies, all infants were born full-term and had no

known health conditions, hearing loss, or visual deficits according to parental report.

Ideal Learner Model.

Intuitively, infants observe how many times each event occurs in the world, and then use these event counts to infer an underlying probability model of their observations. In Experiment 1, the two possible events are that the screen lifts to reveal that an object is either present or absent. In Experiment 2, there are three possible events corresponding to which of three objects appears from behind its box.

An observer who sees only a single event happen would not likely infer that the single observed event is the only one possible (i.e., has probability of 1); instead, observers likely bring expectations to this learning task. In the MDM model used here, this prior expectation is parameterized by a single free parameter, α, which controls the strength of the learner's prior belief that the distribution of events is uniform. As α gets large, the model has strong prior beliefs that the distribution of events in the world is uniform; as α approaches zero, the model believes more strongly that the true distribution closely resembles that of the empirically observed event counts. In modeling, we chose a value of $\alpha = 1$, corresponding to a uniform prior expectation about the distribution of events (with expected values 50-50 in Experiment 1 and 33-33-33 in Experiment 2). However, the qualitative results–in particular, the U-shaped relationship between complexity and look-away probability–do not depend strongly on the choice of α.

Formally, suppose there are N events, x_1, x_2, ... , x_N and the ith event has been observed c_i times. We are interested in estimating (or *scoring*) a multinomial distribution parameterized by $\theta = (\theta_1, \theta_2, ... , \theta_N)$ where θ_i is the true (unobserved) probability of event x_i. Under a Dirichlet-Multinomial model,

$$(1) \qquad P(\theta | c_1, \ldots c_N, \alpha) = \frac{1}{B} \prod_{i=1}^{N} \theta_i^{\alpha + c_i - 1} \quad ,$$

where B is a normalizing constant that depends on the c_i and α. That is, after observing each event type occur some number of times, the infant may form a representation, θ, of their guess at the true distribution of events. Every distribution can be scored according to Equation 1, allowing one to compute how strongly a learner should believe that any particular θ is the correct one. We predict that infants' likelihood of looking away at a current event will depend upon the complexity of that current event, which is determined by both the previously observed events and the identity of the current event. We predict that events of either very low complexity (highly predictable) or very high complexity (highly surprising) will be more likely to trigger a look-away than events with moderate complexity.

When the ith event occurs, the main variable of interest here is its negative log probability according to the model. We compute this by integrating over the above posterior distribution on θ. This corresponds to a measure of the information conveyed by observing event i according to an ideal Bayesian learner who had seen all previous events. We predicted that infants would be more likely to look away during events that contained either too little or too much information, giving a U-shaped (quadratic) relationship between this negative log probability measure and the actual observed look-away probability.

Supporting Information.

The material is part of the online *PLoS ONE* version at

www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0036399

Video S1. An example of an animated single-box display used in Experiment 1.

Video S2. An example of an animated three-box display used in Experiment 2.

Video S3. A 7-month-old subject attending to a three-box display in Experiment 2, and then looking away (and subsequently terminating the trial).

Acknowledgments.

Author Contributions.

Conceived and designed the experiments: CK STP RNA. Performed the experiments: CK. Analyzed the data: CK STP. Contributed reagents / materials / analysis tools: CK STP RN. Wrote the paper: CK STP RN.

References.

Akaike, H. (1974). A New Look at the Statistical Model Identification. IEEE Transactions on Automatic Control, 19(6), 716–723.

Aslin, R.N. (2007). What's in a Look? Developmental Science, 10, 48–53.

Aylett M., & Turk A. (2004). The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. Language and Speech 47, 31.

Berlyne, D.E. (1960). Conflict, Arousal, and Curiosity. (Volume 331). New York: McGraw-Hill.

Civan, A., Teller, D.Y., & Palmer, J. (2005). Relations Among Spontaneous Preferences, Familiarized Preferences, and Novelty Effects: Measurements with forced-choice techniques. Infancy, 7, 111–142.

DeCasper, A.J., & Fifer, W.P. (1980). Of Human Bonding: Newborns prefer their mothers' voices. Science, 208(4448), 1174-1176.

Dember, W.N., & Earl, R.W. (1957). Analysis of Exploratory, Manipulatory, and Curiosity Behaviors. Psychological Review, 64, 91–96.

Fantz, R.L. (1964). Visual Experience in Infants: Decreased attention to familiar patterns relative to novel ones. Science, 146, 668.

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, 15822.

Genzel D, & Charniak E (2002). Entropy Rate Constancy in Text. Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics. Stroudsburg, PA: Association for Computational Linguistics. pp. 199–206

Hastie T., & Tibshirani R. (1990). Generalized Additive Models. Boca Raton: Chapman and Hall/CRC

Hosmer, D.W., Lemeshow, S., & May, S. (2008). Applied Survival Analysis: Regression Modeling of Time-to-Event Data, Second Edition. Hoboken, NJ: John Wiley and Sons.

Hunter, M.A., & Ames, E.W. (1988). A Multifactor Model of Infant Preferences for Novel and Familiar Stimuli. Advances in Infancy Research, 5, 69–95.

James, W. (1890). The Principles of Psychology. Cambridge, MA: Harvard University Press.

Kaldy, Z., Blaser, E.A., & Leslie, A.M. (2006). A New Method for Calibrating Perceptual Salience Across Dimensions in Infants: The case of color vs. luminance. Developmental Science, 9, 482–489.

Kinney, D.K. & Kagan, J. (1976). Infant Attention to Auditory Discrepancy. Child Development, 47(1), 155–164.

Klein J., & Moeschberger, M. (2003). Survival Analysis: Techniques for Censored and Truncated Data, Second Edition. New York: Springer-Verlag

Levy R., & Jaeger T.F. (2007). Speakers Optimize Information Density Through Syntactic Reduction. Advances in Neural Information Processing Systems, 19, 849.

Maye, J., Werker, J.F., & Gerken, L.A. (2002). Infant Sensitivity to Distributional Information Can Affect Phonetic Discrimination. Cognition, 82, B101–B111.

Piaget, J. (1970). Structuralism. New York: Harper & Row.

Piantadosi, S.P., Kidd, C., & Aslin, R.N. (In press). Rich Analysis and Rational models: Inferring individual behavior from infant looking data. Developmental Science.

Piantadosi, S.T., Tily, H., & Gibson, E. (2011). Word Lengths Are Optimized for Efficient Communication. Proceedings of the National Academy of Sciences of the United States of America 108(9), 3526–3529.

Roder, B.J., Bushnell, E.W., & Sasseville, A.M. (2000). Infants' Preferences for Familiarity and Novelty During the Course of Visual Processing. Infancy, 1, 491–507.

Rose, S.A., Gottfried, A.W., Melloy-Carminar. P., & Bridger, W.H. (1982). Familiarity and Novelty Preferences in Infant Recognition Memory: Implications for information processing. Developmental Psychology, 18, 704–713.

Rovee-Collier, C.K., Sullivan, M.W., Enright, M., Lucas, D., & Fagan, J.W. (1980). Reactivism of Infant Memory. Science, 208, 1159–61.

Saffran J.R., Aslin R.N., & Newport E.L. (1996). Statistical Learning by 8-Month-Old Infants. Science, 274, 1926.

Saffran, J.R., Johnson, E.K., Aslin, R.N., & Newport E.L. (1999). Statistical Learning of Tone Sequences by Human Infants and Adults. Cognition, 70, 27–52.

Saylor, M.M., Baldwin, D.A., Baird, J.A., & LaBounty, J. (2007). Infants' On-Line Segmentation of Dynamic Human Action. Journal of Cognition and Development, 8, 113–128.

Shannon, C.E. (1948). A Mathematical Theory of Communication. Bell Systems Technical Journal, 27, 379-423.

Siqueland, E.R., & De Lucia C.A. (1969). Visual Reinforcement of Non-Nutritive Sucking in Human Infants. Science, 165, 1144–1146.

Sokolov, E. (1963). Perception and the Conditioned Reflex. Oxford, England: Pergamon.

Stevenson, H.W. (1972). Children's Learning. New York: Appleton-Century-Crofts.

Téglás, E., Vul, E., Girotto, V., Gonzalez, M., Tenenbaum, J.B., et al. (2011). Pure Reasoning in 12-Month-Old Infants as Probabilistic Inference. Science, 332, 1054.

Tribus, M. (1961). Thermostatics and Thermodynamics: an introduction to energy, information and states of matter, with engineering applications. New York: D. Van Nostrand.

Turk-Browne, N.B., Scholl, B.J., & Chun, M.M. (2008). Habituation in Infant Cognition and Functional Neuroimaging. Frontiers in Human Neuroscience, 2, 1–11.

Wagner, S.H., & Sakovits L.J. (1986). A Process Analysis of Infant Visual and Cross-Modal Recognition Memory: Implications for an amodal code. Advances in Infancy Research, 4, 195–217.

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.

Yerkes, R.M., & Dodson, J.D. (1908). The Relation of Strength of Stimulus to Rapidity of Habit-Formation. Journal of Comparative Neurology and Psychology, 18, 459–482.

IV. The Goldilocks Effect in Infant Auditory Attention.

Celeste Kidd, Steven T. Piantadosi, & Richard N. Aslin

Infants must learn about many cognitive domains (e.g., language, music) from auditory statistics, yet limited cognitive resources restrict the quantity that they can encode. While we know infants can attend to only a subset of available acoustic input, few previous studies have directly examined infant auditory attention—and none have directly tested theorized mechanisms of attentional selection based on stimulus complexity. Using model-based behavioral experimentation methods we first developed to examine visual attention in infants (e.g., Kidd, Piantadosi, & Aslin, 2012), we demonstrate that infants' selectively attend to auditory stimuli that are intermediately predictable/complex with respect to their current beliefs/knowledge. Our results provide evidence of a broad principle of infant attention across modalities and suggest that auditory attention relies heavily on sound-to-sound transitional statistics.

Introduction.

Infants are remarkably sensitive to their auditory environment, showing the ability to learn from their mother's speech even before birth (DeCasper & Fifer, 1980). This process of learning from the auditory environment continues during the first postnatal year, as infants discover phonetic categories (Kuhl, 2004) and learn the sequences of speech that will form the words of their native language (Saffran, Aslin & Newport, 1996). These auditory milestones must be based on gathering input from the natural environment, where a myriad of novel sounds and sound-sequences (e.g., speech syllables, musical notes) unfold rapidly over time. A learner with an infinite information processing capacity could theoretically encode all available auditory input as it arrives at the ear. A human infant, however, possesses only limited cognitive resources (e.g., attention, memory, processing capacity). These cognitive constraints impose severe limits on the kind and quantity of auditory input an infant can encode in real time. Infants' learning is thus limited by constraints such as the temporal rate at which they can access sequential inputs (e.g., Conway & Christiansen, 2009), the number of elements they can hold in working memory (e.g., Ross-Sheehy, Oakes & Luck, 2003), and the depth to which they can ultimately encode the novel stimulus (e.g., Sokolov, 1969).

Even a single auditory stream (e.g., a mother speaking to her child in an otherwise silent room) expresses a complex composition and arrangement of acoustic variables (e.g., intensity, pitch, timbre) that additionally encode hierarchical levels of structure (e.g., sounds, syllables, phrases) and semantic meaning (e.g., salience, emotion, category, identity). Additionally, previous work with adults suggests that human auditory processing is likely inferior to visual processing in terms of resolution and capacity (e.g., Cohen, Horowitz, & Wolfe, 2009). Thus, the infant must pick and choose both *which auditory inputs* to attend

and *on which aspects of a single auditory stream* to focus. Locating and tracking the relevant statistics from within the continuous surge of incoming auditory data is then crucial for infants to solve the many auditory learning tasks they face. One reasonable strategy infants might employ in the natural environment is to allocate attention on an "as available" basis; that is, they might attempt to encode all auditory inputs, and effectively ignore stimuli that exceed their information processing capacity. However, such an undirected learning strategy would be inefficient at best, and futile at worst. Imagine, for example, attempting to complete an open-book exam on an unfamiliar subject in a vast library by drawing books from the shelves at random. An alternative strategy would be to make attention dependent upon relevant properties of the stimulus itself, perhaps actively allocating attention to auditory material that is most useful for learning. This latter strategy might be particularly advantageous for language learning, where the inventory of inputs is quite large (e.g., 40 phonemes, 1,000 syllables, 50,000 words) and combined in a huge variety of sequences.

A substantial amount of previous work on infant attention theorized that such a strategy might help infants focus on learning material that is sufficiently novel from—but also sufficiently related to—the infants' existing knowledge (e.g., Kinney & Kagan, 1976; Jeffrey & Cohen, 1971, Friedlander, 1970; Horowitz, 1972; Melson & McCall, 1970), Zelazo & Komer, 1971). Kinney and Kagan (1976) suggested that preferring stimuli that are moderately novel would prevent infants from wasting time on material that is already known. They further suggested that preferring stimuli that are somewhat related to existing knowledge might help infants focus on completing partially built cognitive representations. These partial representations could then facilitate more efficient construction of newer, bigger or more elaborate cognitive constructs later on in learning. This formulation of the *"discrepancy hypothesis"* thus suggests that the complexity of a

stimulus can be conceptualized as relating to the infant's current knowledge state. A "simple" stimulus would be one with little or no new information for the infant to learn. A "complex" stimulus would be one that contains almost entirely new information, distinct from nearly everything in the infant's current conceptual inventory. Further, these theories hold that infants should exhibit a U-shaped attentional pattern with respect to stimulus complexity: infants should more readily terminate attention to events that are either too simple (predictable) or too complex (surprising).

Our previous work (Kidd, Piantadosi, & Aslin, 2010; 2012) demonstrated that infants' visual attention was influenced by the complexity (or information content) of the visual stimulus. We used an idealized learning model in order to quantify the complexity of particular visual events in a sequence. We then measured at what point in a visual sequence an infant terminated their attention to the sequence. In these studies, infants looked away at visual events of either very low complexity (very predictable) or very high complexity (very surprising). Additional work demonstrated that this U-shaped pattern of preference for visual events of intermediate complexity occurred not only across a population of infants, but also within individual infants (Piantadosi, Kidd, & Aslin, in press). In the present study, we asked whether such an active strategy of attentional allocation extends from the visual modality to the auditory modality.

As suggested by the discrepancy hypothesis of infant auditory attention discussed earlier, the potential utility of such a strategy is substantial. In contrast to the large quantity of work examining auditory learning in infants (e.g., the literature on language learning and music cognition), few previous studies have directly examined infant auditory attention—and none to our knowledge have employed computationally well-defined stimuli varying in complexity. Previous work on infant auditory processing, however, provides some evidence suggesting

that the processes that govern selective auditory attention are not fully mature in infants. At 6 months of age, infants have difficulty detecting auditory signals in noise, detecting changes in intensity, and discriminating between low-frequency tones (e.g., Werner, 2002). Thus, infants may not have the capabilities required to identify and focus their auditory attention on the most informative parts of the auditory signal.

It is important to note that the general idea of a U-shaped function along a dimension of stimulus complexity is not new. In fact, several recent studies of infants (Gerken, Balcomb & Minton, 2011; Spence, 1996) have reported similar effects. What is new about our approach is to make a specific prediction about the U-shaped function based on a quantitative metric of complexity. Previous studies have either defined complexity after obtaining a U-shaped function or have contrasted learnable versus unlearnable information rather than exploring the space of complexity in a continuous manner. Moreover, it is important to determine whether the same general principles of attention allocation apply in the auditory modality as well as in the visual modality, especially given modality differences in the temporal and spatial statistics typically used to process natural stimuli in each domain.

In the present experiment with 7- and 8-month-olds, we measured infants' visual attention to sequential sounds that varied in complexity, as determined by an idealized learning model. Both the experiment and modeling approach were based on our earlier studies on visual attention (Kidd, Piantadosi, & Aslin, 2010; Kidd, Piantadosi, & Aslin, 2012; Piantadosi, Kidd, & Aslin, in press).

Fig. IV-1 illustrates the logic of the experiment and our analysis approach in which one of three possible sounds is presented in a sequence that varies in complexity across a series of trials.
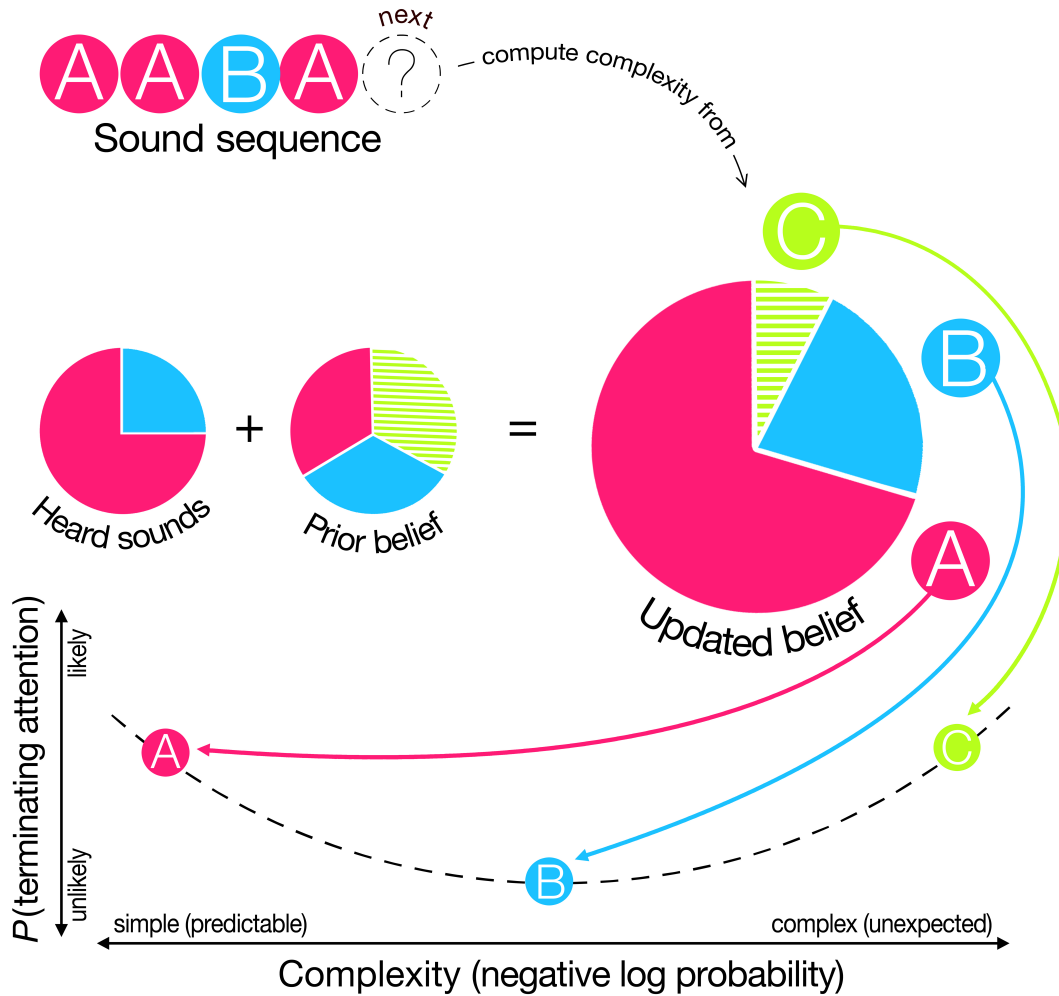
Fig. IV-1. Schematic of Idealized Learning Model. Schematic showing an example sound sequence and how the idealized learning model combines heard sounds with a simple prior to form expectations about upcoming sound events (the "updated belief" above). The next sound then conveys some amount of complexity according to these probabilistic expectations of the updated belief. The "Goldilocks" hypothesis holds that infants will be most likely to terminate their attention to the sequence at sounds that are either overly simple (predictable) or that are overly complex (unexpected), according to the model. Thus, sounds to which the updated belief assigns either a very high probability (e.g., sound A) or a very low probability (e.g., sound C) would be expected to be more likely to generate attentional termination (look-aways) than those to which it assigns an intermediate probability (e.g., sound B).

In this example trial, the observer has heard a sequence composed of three *A* sounds and one *B* sound, and the key question is whether the infant terminates the trial upon hearing the next sound in the sequence. The heard sounds (*AAAB*) comprise the observed data, which are combined with the prior—essentially a smoothing term to avoid zero probabilities—to form an updated (posterior) belief. In this example, the updated belief leads to an expectation that the next event has a high probability of being sound *A*, a moderate probability of it being sound *B*, and a low (but non-zero) probability of it being sound *C*.

The complexity of the next sound is quantified by an information theoretic metric—negative log probability— which represents the amount of "surprise" an idealized learner would have on hearing the next event, or, equivalently, the amount of information processing such a learner would be required to do (Shannon 1948). Thus, if the next sound is *A*—a sound that is *highly likely* according to the model's updated belief—the complexity of that event would be low (i.e., the sound would be highly predictable according to the model). The "Goldilocks" hypothesis thus holds that infants would be more likely to terminate their attention at this sound. Conversely, if the next sound is *C*—a sound that is *highly unlikely* according to the model's updated belief—the complexity of that event would be high (i.e., the sound would be highly surprising according to the model). The "Goldilocks" hypothesis holds that infants should also terminate their attention to the sound sequence at this type of event. However, if the next sound is *B*—a sound that is *moderately probable* according to the model's updated belief— the complexity of that event would fall in the intermediate "Goldilocks" range, thus leading infants to be less likely to terminate their attention to the sound sequence.

The example shown in Fig. IV-1 treats each event as statistically independent (a non-transitional model). However, our previous work also

indicated that a model that tracked the bigram probabilities of events (a transitional model) out-performed the non-transitional model. In the present experiment, therefore, we also constructed and tested a transitional model of the auditory stimuli, which captured how likely each sound was to follow each other sound in computing complexity. Note that for either model, if an infant continued to attend to the sound sequence, the predictions of the model would be updated for the next sound in the sequence. Thus, although infants may terminate their attention at different points in different sound sequences, we hypothesize that these attentional terminations (as measured by look-aways) will occur predictably during events with a medial amount of complexity, as defined by the two models.

Materials and Methods.

PARTICIPANTS.

Thirty-four infants (*mean* = 7.7 months, *range* = 7.1 - 8.9) were tested and all were included in the analysis. All infants were born full-term and had no known health conditions, hearing loss, or visual deficits according to parental report.

STIMULI.

We presented each infant with 32 trials consisting of sequences composed of three sounds, with trials presented in a random order across infants. These sequences were constructed to vary in their information-theoretic properties (e.g., entropy, surprisal). Thus, some sound sequences contained many highly predictable events (e.g., *AAAAAAAA*) and others contained many less predictable ones (e.g., *BBACAACAB*).

Each of the sound sequences presented up to three non-social sounds (e.g., door closing, flute note, train whistle). These sounds were selected randomly for each infant and the three sounds in each sequence were unique, such that each infant heard 96 sounds across all 32 trials. The sounds were chosen to be both reasonably familiar, but also maximally memorable and distinct from one another. Each sound sequence was presented while infants viewed a unique scene on each of the 32 trials, generated by a Matlab script. Each scene consisted of a single, uniquely patterned and colored box concealing a single, unique toy at the center of the screen (see Fig. 2 and Video S1). The box was animated to open (1 sec.), thus revealing its contents, then immediately close (1 sec.), so that each reveal lasted 2 sec. Each reveal was accompanied by one sound from the sound sequence. The box continued to open and close continuously, revealing the same toy on that particular trial and each time accompanied by the next sound in the sound sequence. The toy was present to maintain infants' visual fixation, and did not change within a sequence, but was randomized across trials and infants; thus, there were no differences in the visual displays across sounds in a sequence, and look-aways could only be attributed to the auditory portion of the stimulus presentation.

Fig. IV-2. Example of Display Used in the Experiment. A novel toy object (e.g., a little teardrop-shaped figure) in the box was revealed by up-down animation of an occluder (e.g., a yellow-striped box). Also see **Video S1** for example of animated display.

Neither the boxes nor the objects were repeated across the 32 trials, rendering each object-box pair independent and unique. Thus, there were 32 visual stimuli, one for each sound sequence, and each sound sequence was associated with a different, randomized box-object pairing across infants. This design ensured that differences in attentional termination across sound sequences were not driven by differences in visual materials or particular sounds.

PROCEDURE.

Each infant was seated on his or her parent's lap in front of a table-mounted Tobii 1750 eye-tracker. The infant was positioned such that his or her eyes were approximately 23 inches from the monitor, the recommended distance for accurate eye-tracking. At this viewing distance, the 17-inch LCD screen subtended 24 X 32 degrees of visual angle. The box at the center of the screen was 3 X 3 inches. To prevent parental influence on the infants' behavior, the parents were asked to wear headphones playing music, lower their eyes, and abstain from interacting with their infants throughout the experiment.

Each of the 32 trials was preceded by an animation designed to attract the infant's attention to the center of the screen—a laughing and cooing baby. Once the infant looked at the attention-getter, an experimenter who was observing remotely via a wide-angle video camera pushed a button to start the trial. Every infant heard all 32 sound-sequence trials.

For each trial, an animated scene (box opening and closing) for that sound sequence was played. The animated sequence of events—single instances of one of three sounds accompanied by a box opening and closing—continued until the infant looked away continuously for 1 sec., or until the sequence timed out at 60 sec. The Tobii eye-tracking software automatically determined the 1-sec. look-

away criterion for trial termination. If the trial was terminated before the infant actually looked away, as determined after the experiment by a wide-angle video-recording of the infant's face, the trial was labeled by an experimenter as a "false stop" and discarded before the analysis. False stops occurred as a result of the Tobii software being unable to detect the child's eyes continuously for 1 sec., usually due to infants inadvertently moving out-of-range or blocking their own eyes from detection (14.7% of trials). If the infant looked continuously for the entire 60-sec. sequence, the trial was automatically labeled as a "time out" and also discarded (4.4% of trials). Finally, trials in which the infant looked for fewer than four events were also discarded, since we judged such limited observations are likely insufficient for establishing expectations about the distribution of events (40.9% of trials). Changing the look-away criterion to include more data does not affect the general qualitative or quantitative pattern of results, but we report here data based on these exclusion criteria because they match those of Kidd, Piantadosi, & Aslin (2012). This resulted in a mean of 11.5 +/- 5.5 sequences (of at most 32 trials each) from each infant.

The dependent measure for the subsequent computational modeling was the sound at which the infant looked away in each trial (e.g., the specific point in each sequence where the infant looked away from the display for more than 1 consecutive second).

ANALYSIS.

Analysis of the behavioral data followed the approach used in Kidd, Piantadosi, & Aslin (2012) and Piantadosi, Kidd, & Aslin (in press). A Markov Dirichlet-multinomial model first quantified an idealized learner's expectation that each of the three sounds would occur next, at each point in the sequence. This rational model essentially combines a "smoothing" term—or prior expectation of

sound likelihood—with counts of how often each sound has been heard previously in the sequence to predict each sound's probability of occurring next. The model's estimated negative log probability for each sound quantifies the sound's complexity on a scale corresponding to how many bits of information an idealized learner would require to remember or process each sound. We also applied the MDM model to the data under an assumption of event-order dependence. That is, instead of treating every sound as independent, we examined whether look-aways were predicted by the immediately preceding sound (i.e., a transitional model)[1]. In the analysis that relates model-measured complexity to behavior, standard linear or logistic regressions are inappropriate because infants cannot provide additional data on a trial once they have terminated their attention, thus violating the independence assumption required for these analyses. Thus, the obtained complexity measure was then entered as a quadratic term in a stepwise Cox regression of the behavioral data, as employed in Kidd, Piantadosi, & Aslin (2012). The Cox regression is a type of survival analysis that measures the log linear influence of predictors on infants' probability of terminating attention, but respects the fact that infants cannot provide additional trial data once they terminate attention (Hosmer, Lemeshow, & May, 2008; Klein & Moeschberger, 2003). Importantly, the Cox regression allows the significance of a quadratic complexity term (an underlying U-shape) to be tested while controlling for a baseline distribution of look-aways and other factors including whether the current sound was its first occurrence, the number of unheard sounds, and whether the sound was an immediate sequential repeat.

---

[1] We note that the models imperfectly assume that infants know how many sounds are possible on each display. This simplification keeps the analysis in line with Kidd, Piantadosi, & Aslin (2012) and Piantadosi, Kidd, & Aslin (in press); further, and more importantly, it is the most reasonable of several possible imperfect analysis options. It is likely that infants would learn that only three sounds occur per sequence within the first few trials. Other analyses that model uncertainty in the number of sounds per trial (e.g. a Chinese restaurant process) lead to implausible assumptions, such as that the first sound always has probability of 1 (meaning no other sound was possible).

Results.

Fig. IV-3 shows infants' probability of terminating attention, as a function of the negative log probability of a sound according to the non-transitional model.
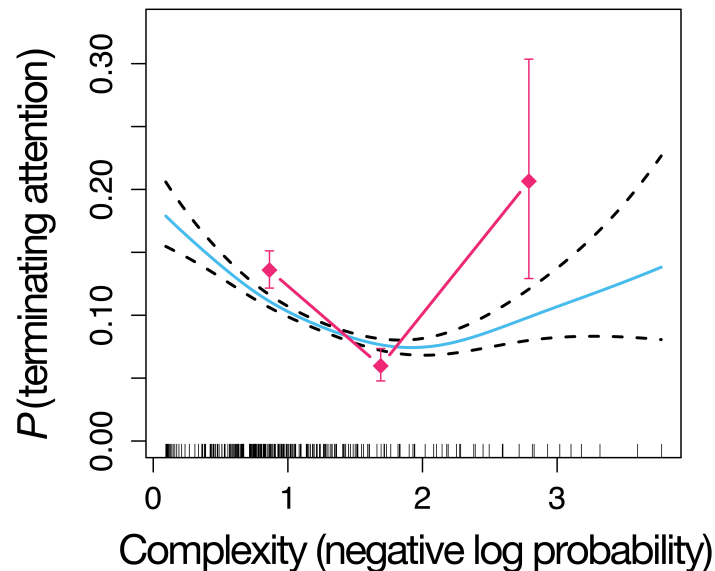
Fig. IV-3. U-Shaped Curve for the Non-Transitional Model. The blue solid curve represents the fit of a Generalized Additive Model (GAM) (Hastie & Tibshirani, 1990) with binomial link function, relating complexity according to the MDM model (*x*-axis) to infants' probability of terminating attention (*y*-axis). The dashed curves show standard errors according to the GAM. The GAM fits include the effect of complexity (negative log probability) and the effect of position in the sequence. Note, the error bars and GAM errors do not take into account subject effects. Vertical spikes along the x-axis represent data points collected at each complexity value. The fuchsia diamonds represent the raw probabilities of terminating attention binned along the *x*-axis.

The plot collapses across infants, sequences, and sequence positions. The diamonds represent the raw probability of terminating attention with complexity divided into 3 discrete bins. The smooth curve represents the fit of a Generalized Additive Model (Hastie & Tibshirani, 1990) with logistic linking function, which fits a continuous relationship between complexity and probability of terminating attention. The figure shows a U-shaped relationship between infants' probability of attentional termination and the model-based estimate of sound event complexity.

This indicates that infants were more likely to terminate attention at a sound in the sequences with either very low or very high complexity (i.e., ones that are very predictable or very surprising, according to the model). There is a "Goldilocks" value of complexity around 2 bits, corresponding to infants' preferred rate of information in this task. However, the Cox regression analysis revealed that this U-shaped trend was not significant controlling for the baseline look-away distribution ($\beta = 0.008$, $z = 0.325$, $p > 0.7$), suggesting that other factors contributed to the U-shape.

Fig. IV-4 shows the outcome of the same analysis, but now applied to successive pairs of events. This transitional model also yields a U-shaped function.
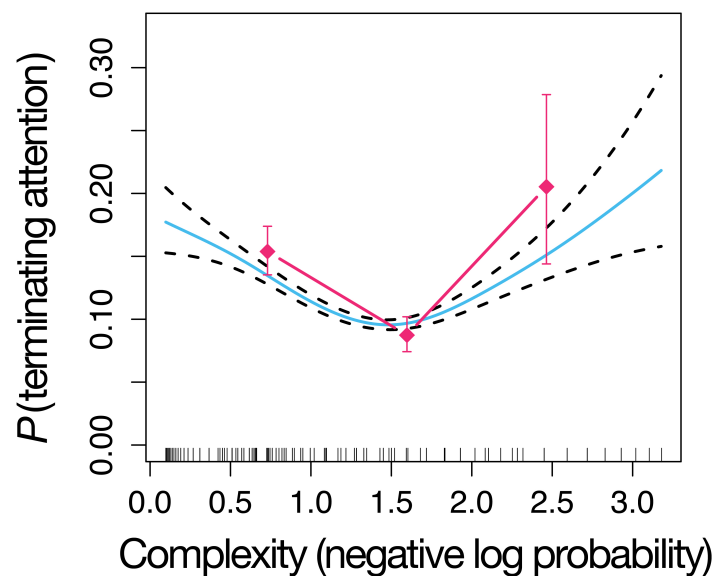
Fig. IV-4. U-shaped curve for the transitional model. The blue solid curve represents the fit of a GAM, relating complexity as measured by the transitional MDM (*x*-axis) to probability of terminating attention (*y*-axis). Dashed curves show GAM standard errors. The GAM fits include the effect of complexity (negative log probability) and the effect of position in the sequence. Note, the error bars and GAM errors do not take into account subject effects. Vertical spikes along the *x*-axis represent data points collected at each complexity value. The fuchsia diamonds represent the raw probabilities of terminating attention binned along the *x*-axis.

The complexity measure—along with a number of control covariates that could plausibly influence infant attentional termination—were entered into the Cox regression using a stepwise procedure that only added variables that improved model fit. The control variables included trial number, whether or not the sound had occurred before in the sequence, and whether or not the sound was the same as the last one that had played in the sequence (Table III-1). This stepwise procedure revealed a highly significant effect for squared complexity ($\beta = 0.136$, $z = 2.91$, $p < 0.01$). This indicates that the U-shape observed in Fig. 4 is statistically significant, even after controlling for an overall baseline look-away distribution and the other potentially confounding variables. The magnitude of this effect can be understood by exponentiating the coefficient for squared complexity ($e^{0.136} = 1.15$). This number quantifies how much more likely infants are to terminate attention at events that are one standard deviation from the experiment's overall mean complexity. In this case, infants are 1.15 times more likely to terminate attention at such high- or low-complexity sounds. This effect is relatively small, though statistically reliable. This analysis also revealed an effect of trial number ($\beta = 0.031$, $z = 5.76$, $p < 0.001$) and first occurrence of a sound ($\beta = 0.523$, $z = 2.23$, $p < 0.05$), suggesting an overall tendency to look away at earlier sounds during later trials and on sounds which are occurring for the first time in the sequence.

Table 1. Cox Regression Coefficients (Transitional model)

Covariate	Coef.	*exp*(coef.)	Std. Error	Z-statistic	P-value
Squared complexity	0.136	1.15	0.047	2.91	0.004 **
Trial number	0.031	1.03	0.005	5.76	8.61e-09 ***
First occurrence	0.523	1.69	0.235	2.23	0.026 *
All transitional-model variables added by the stepwise procedure, which only added variables that improved model fit according to the Akaike information criterion (Akaike, 1974). These results reveal significant quadratic effects of complexity. Both the *complexity* and *squared complexity* variables were shifted and scaled to have a mean of 0 and standard deviation of 1 before they were entered into the regression.					

Discussion.

Our results from the transitional MDM model suggest that infants seek to maintain intermediate rates of complexity when allocating their auditory attention to sequential sounds. This is consistent with the hypothesis that infants employ an implicit strategy of attentional allocation in the auditory modality that is very similar to attention in the visual modality.

Interestingly, the results from the non-transitional model for auditory stimuli were not significant—in contrast to the robust results of the non-transitional model reported for *visual stimuli* in Kidd, Piantadosi, and Aslin (2012). Dissimilarly, the *transitional model* for auditory stimuli showed robust evidence of the U-shaped function, even after controlling for a number of other factors, including a baseline look-away distribution. This notable difference across models could indicate that effects of non-transitional learning are weak or non-existent for auditory stimuli. In other words, attention to auditory stimuli could rely more heavily on temporal order information than does attention to visual stimuli. If so, this would have interesting implications for potential cross-modality differences in infants' attentional systems and learning. For example, though children certainly show sensitivity to frequency differences for auditory stimuli, this apparent sensitivity could arise as the result of learning about transitional statistics (e.g., children's learning about the transitional probabilities between words could yield apparent phrase-frequency sensitivity as in Bannard & Matthews, 2008). It could be that the transient nature of auditory stimuli leads attention to be directed more to successive differences rather than to raw frequencies of occurrence, something that may be less relevant in the visual modality. Alternatively, tracking of the transitional probabilities of auditory stimuli may either be easier or more crucial for developing useful expectations about the auditory world. This is arguably true in language learning, where the

meanings of words are composed not of single events but rather sequences of sounds, and the meanings of utterances tend to be composed not of single words, but of sequences of words. If this were the case, it could be relevant to determine whether this is an innate bias of humans to process auditory stimuli in this way, or whether this attentional pattern might develop over time as infants begin to acquire language[2].

Conclusions.

We hypothesized that infants' probability of terminating their auditory attention would be greatest on sounds whose complexity (negative log probability), according to an idealized learning model, was either very low or very high. We found evidence that this was true for the transitional version of the model, but the trend in the non-transitional version was not significant after controlling for other factors. This may indicate that transitional statistics are more readily tracked by infants in the auditory modality. In general, our results are further evidence for a principle of infant attention that may have broad applicability: infants implicitly seek to maintain intermediate rates of information absorption and avoid wasting cognitive resources on overly simple or overly complex events—in both visual and auditory modalities.

---

[2] It may also be the case that the non-transitional model regression was insignificant because the effects of non-transitional complexity were too highly correlated with the baseline looking distribution. In this case, we might not have had enough power to find an effect of non-transitional sound event complexity while controlling for the baseline distribution.

Acknowledgements.

References.

Akaike, H. (1974). A New Look at the Statistical Model Identification. *IEEE Transactions on Automatic Control, 19*(6), 716–723.

Bannard, C. & Matthews, D.E. (2008). Stored Word Sequences in Language Learning: The Effect of Familiarity on Children's Repetition of Four-Word Combinations, *Psychological Science*, *19*, 241-248.

Cohen, M.A., Horowitz, T.S., & Wolfe, J.M. (2009). Auditory Recognition Memory Is Inferior to Visual Recognition Memory. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(14), 6008-6010.

Conway, C.M., & Christiansen, M.H. (2009). Seeing and Hearing in Space and Time: Effects of modality and presentation rate on implicit statistical learning. *European Journal of Cognitive Psychology*, *21*, 561-580.

DeCasper, A.J., & Fifer, W.P. (1980). Of Human Bonding: Newborns prefer their mothers' voices. *Science*, *208*(4448), 1174-1176.

Friedlander, B.Z. (1970). Receptive Language Development in Infancy: Issues and Problems. *Merrill-Palmer Quarterly of Behavior and Development, 16(*1), 7-51.

Gerken, L., Balcomb, F.K., & Minton, J.L. (2011). Infants Avoid 'Labouring in Vain' by Attending More to Learnable than Unlearnable Linguistic Patterns. *Developmental Science, 14*(5), 972-9.

Hastie T., & Tibshirani R. (1990). *Generalized Additive Models.* Boca Raton: Chapman and Hall/CRC.

Horowitz, A.B. (1972). Habituation and memory: Infant cardiac responses to familiar and unfamiliar stimuli. *Child Development, 43*, 45-53.

Hosmer, D.W., Lemeshow, S., & May, S. (2008). *Applied Survival Analysis: Regression Modeling of Time-to-Event Data, Second Edition.* Hoboken, NJ: John Wiley and Sons.

Jeffrey, W.E., & Cohen, L.B. (1971). Habituation in the human infant. In H. Reese (Ed.), Advances in Child Development and Behavior, Vol.6 (pp. 63-97). New York: Academic Press.

Kidd, C., Piantadosi, S.T., & Aslin, R.N. (2010). The Goldilocks Effect: Infants' preference for visual stimuli that are neither too predictable nor too surprising. *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society* (pp. 2476-2481).

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*(5), e36399.

Kinney, D.K. & Kagan, J. (1976). Infant Attention to Auditory Discrepancy. *Child Development, 47*(1), 155–164.

Klein, J., & Moeschberger, M. (2003). *Survival Analysis: Techniques for censored and truncated data, Second Edition.* New York: Springer-Verlag.

Kuhl, P.K. (2004). Early Language Acquisition: Cracking the speech code. *Nature Reviews Neuroscience, 5*(11), 831-843.

Melson, W.H., & McCall, R.B. (1970). Attentional Responses of Five-Month Girls to Discrepant Auditory Stimuli. *Child Development, 41*, 1159-1171.

Piantadosi, S.P., Kidd, C., & Aslin, R.N. (In press). Rich Analysis and Rational models: Inferring individual behavior from infant looking data. *Developmental Science.*

Saffran J.R., Aslin R.N., & Newport E.L. (1996). Statistical Learning by 8-Month-Old Infants. *Science, 274*, 1926.

Shannon, C.E. (1948). A Mathematical Theory of Communication. *Bell Systems Technical Journal, 27*, 379-423.

Spence, M.J. (1996). Young Infants' Auditory Memory: Evidence for changes in preference as a function of delay. *Developmental Psychology, 29*(8), 685-695.

Sokolov, E.N. (1969). *Mekhanizmy Pamyat.* Moskva: Isdatelstvo Moskovskogo Universiteta.

Ross-Sheehy, S., Oakes, L.M., & Luck, S.J. (2003). The Development of Visual Short-Term Memory Capacity in Infants. *Child Development, 74*, 1807-1822.

Werner, L.A. (2002). Infant Auditory Capabilities. *Current Opinion in Otolaryngology and Head and Neck Surgery, 10*, 398–402.

Zelazo, P.R., & Komer, M.J. (1971). Infant Smiling to Nonsocial Stimuli and the Recognition Hypothesis. *Child Development, 42*(5), 1327-1339.

V. Curious George.

Intrinsic curiosity and information-seeking behavior in monkey learners.

Celeste Kidd, Tommy Blanchard, Richard N. Aslin, & Benjamin Y. Hayden

Curiosity motivates intelligent organisms to behave in ways that reduce uncertainty in a complex world. Recent research findings in the domains of behavioral economics, memory, and motivation have provided insights into some of the neurological mechanisms that underlie curiosity-driven behavior. However, it remains unclear whether the brain possesses mechanisms that track informational complexity and deploy attention adaptively. Here we demonstrate that juvenile rhesus monkeys, like human infants, allocate attention according to the statistical properties of stimuli in their environments. Our results suggest that the ability to monitor the statistical properties of incoming information streams and selectively maintain or terminate attention based on these properties is a general property of intelligent organisms.

Introduction.

Curiosity is a function that drives intelligent creatures—including humans, apes, rats, felines, and canines—to reduce the uncertainty that is inherent in a complex world. Curiosity motivates these creatures to seek out unknowns, and thus guides them towards novel learning material—but this cannot be the whole story. A novelty-based account of curiosity could likely prove sufficient for a learner in a small, static environment composed of simplistic stimuli; however, the true amount of information in the world is vast, voluminous, and complex. Learners are thus constantly confronted with any number of novel stimuli that they must chose among to explore—and they must also consider the possibility that stimuli that they have previously explored may since have changed.

Recent research has focused predominately on understanding curiosity at the level of its underlying biological processes. Under this framework, a predominant theory is that curiosity is a state of increased arousal whose termination is rewarding and facilitates memory (e.g., Berlyne, 1960; Jepma et al., 2012). In other words, curiosity is a negative reinforcer (or *aversive condition*) that motivates exploration and discovery. In the past decade, the field has made significant progress in better understanding the chemical processes associated with curiosity in the reward system of the brain. Desires for novel information are associated with dopamine activation along the mesolimbic pathway, and discoveries of novel information stimulate pleasure via opioid activity in the nucleus accumbens (e.g., Litman, 2005). A unifying feature of this recent work is that a wide variety of researchers attempt to explain behavior through neurobiological processes—typically, the processes associated with desire, motivation, discovery, reward, and memory.

A comprehensive understanding of the exploratory mechanisms that govern curiosity, however, entails more than just the associated biological processes. Rather than asking what biological processes drive exploratory behavior, here we ask what high-level function the exploratory behavior serves—an inquiry falling squarely within Marr's computational level of analysis (Marr, 1982). As Marr observed was true of any information processing system, the mechanisms that govern curiosity and exploration must be understood at three distinct, but complementary, levels: computational, representational, and implementational. The computational level asks what the system does, and why it does these things. The representational level asks what representations and processes the system employs to accomplish these functions. And finally, the implementational level concerns how the system is physically realized. Recent work in behavioral neuroscience has focused a great deal of effort on this third level—the physical mechanisms that drive exploratory behavior—while leaving the higher levels largely neglected.

Here, we tackle two major computational-level questions. First, what does curiosity do, exactly? Does it motivate exploration of any available novel stimulus in a random manner, or perhaps only the most informative stimuli? Addressing this question requires understanding the relationship between intrinsic curiosity and relevant features of the stimulus. This relationship would, in essence, elucidate the high-level design features that govern attentional selection and termination—what makes something inherently interesting? Second, we ask why biology should have given rise to the particular set of exploratory mechanisms that govern attentional selection, as opposed to any other? What greater purpose do exploratory mechanisms serve?

Once we have obtained evidence for a high-level theory that relates informational value to curiosity, we can use that high-level theory to guide our

low-level inquiries about the physical mechanisms of the system. Without a high-level theory, it is difficult to know in what computations neurons are likely to be engaged. Thus, it remains unclear whether the brain possesses mechanisms that track the informational values of stimuli in the environment and deploy attention adaptively.

WHAT IS INHERENTLY INTERESTING AND WHY?

Many theories of curiosity-driven behaviors posit that they aim to guide exploration in order to maximize learning or learning efficiency (e.g., Berlyne, 1954; Dember & Earl, 1957; Fantz, 1964; Kinney & Kagan, 1976; Min Jeong et al., 2009; Ofer & Durban, 1999). According to these theories, the ultimate goal of the learner is to obtain as accurate of a mental model of the world as possible so that the learner may function near optimally. Accurate mental models yield accurate predictions, thus enabling the learner to capitalize on forthcoming rewards while avoiding imminent penalties. If curiosity is in fact a means to these ends, curiosity should seek out information that is most useful for learning. Thus, we will conduct a computational Marr-level test of this principle.

To test the theory that curiosity serves to maximize learning, we need a method for quantitatively relating the stimulus to the information that learners have or have not yet acquired. For this, the informational value of a stimulus crucially depends upon the learner's existing knowledge. A particular stimulus is of very high information value if it is entirely different from everything in the learner's existing knowledge banks; however, that same stimulus would be of very low information value if it (or something very similar) has previously been encoded into the learner's knowledge. With a computationally well-defined method for relating stimuli to knowledge, we can then set up a test to ask whether the learner's

behavior follows the patterns we would expect from curiosity that maximizes learning.

We explore this and other questions using a combination of behavioral and modeling methods with juvenile rhesus monkeys. Our results are the first to demonstrate that monkeys' intrinsic curiosity—in the absence of any external rewards or task goals—is governed by the statistical properties of the stimuli in their environments. Taken in combination with our previous work on the attentional patterns of human infants, our results suggest that the ability to monitor the statistical properties of incoming information streams and selectively maintain or terminate attention based on these properties is a general property of intelligent organisms. Further, they provide strong evidence that curiosity functions to drive learning, both in infants and monkeys.

PREVIOUS WORK ON INFANT ATTENTION.

Our previous work with infants suggested that curiosity—acting through key attentional mechanisms—filters environmental stimuli to provide infants with data that are "just right" for learning (which we referred to as a "Goldilocks" effect). This work (Kidd, Piantadosi, & Aslin, 2010, 2012, under review; Piantadosi, Kidd, & Aslin, in press) explored attentional behavior in 7- and 8-month-old infants. We showed infants visual event sequences of varying informational complexity, as measured by an idealized learning model, and measured which sequential events prompted infants to terminate their attention to the display. We found that infants' probability of looking away was greatest to low surprisal (highly predictable) events—but also to very high surprisal (highly unexpected) ones. Further work examined the robustness and utility of the "Goldilocks" principle for young learners. This attentional strategy holds in multiple types of visual displays (Chapter 3, Kidd, Piantadosi, & Aslin, 2010,

2012), for auditory stimuli (Chapter 4, Kidd, Piantadosi, & Aslin, under review), and even within individual infants (Piantadosi, Kidd, & Aslin, in press). These results suggest that infants implicitly allocate their attention towards events with intermediate surprisal values. In so doing, infants are likely reserving their attention for intermediately informative events and avoiding wasting time and cognitive resources on those events that are overly predictable/simple or overly unexpected/complex.

EXPERIMENT AND MODELING APPROACH.

In the present experiment, we measured monkeys' visual attention to sequential visual events that varied in their information theoretic properties (e.g., *surprisal*), as determined by an idealized learning model. Both the experiment and modeling approach were based on our earlier studies of infant visual and auditory attention (Kidd, Piantadosi, & Aslin, 2010, 2012, under review; Piantadosi, Kidd, & Aslin, in press).

Fig. V-1 illustrates the logic of the experiment and our analysis approach in which one of three possible objects "pop-up" from behind an occluding box in a sequence that varies in predictability (surprisal value) across a series of trials.
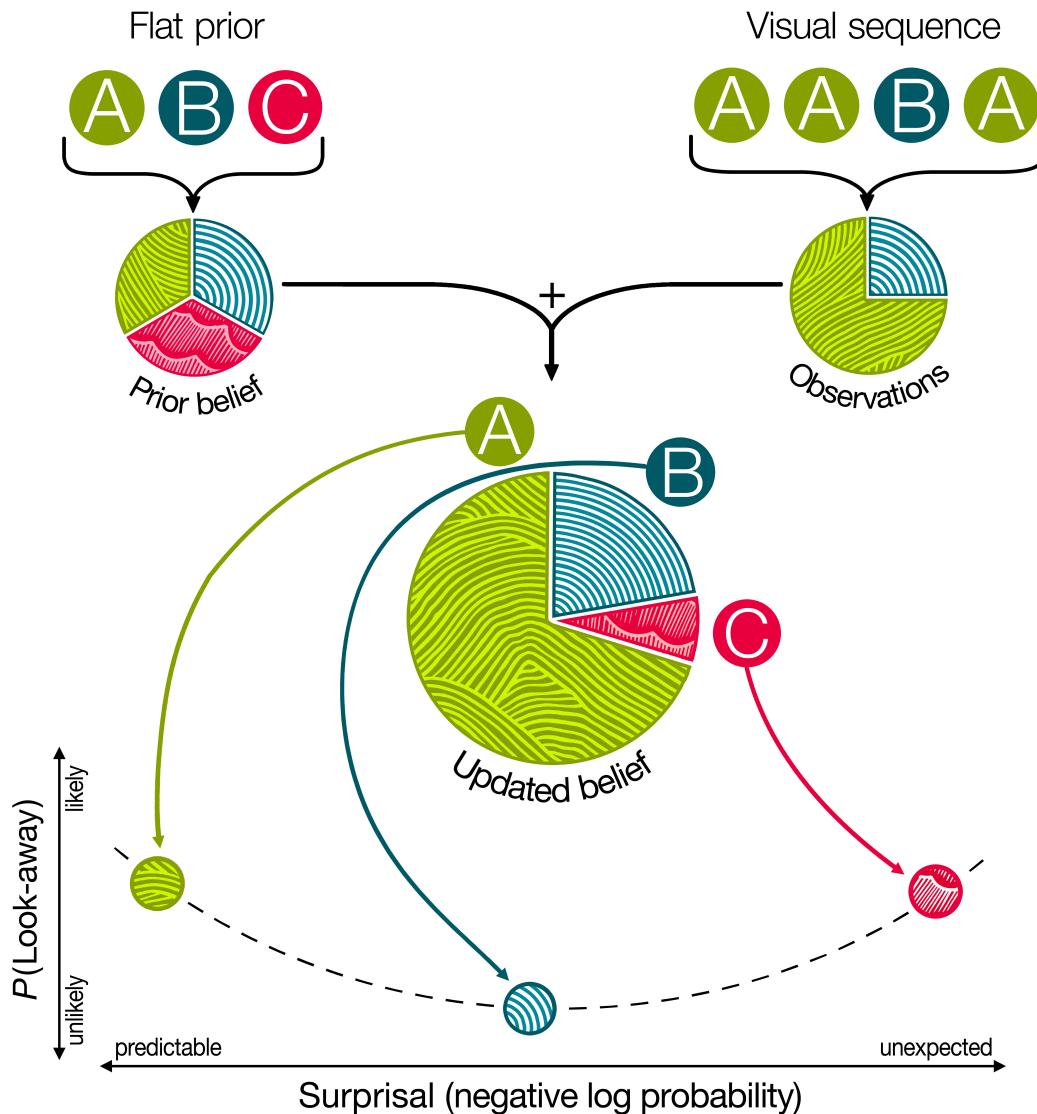
Fig. V-1. Idealized Learning Model Schematic. Schematic showing an example of how the idealized learning model forms probabilistic expectations about the expectedness of the next event in a sequence. The model begins with a simple prior corresponding to the beliefs a learner possesses before beginning to make any observations. By using a flat (or uninformative) prior, we assume that the learner begins the sequence presentation with the belief that each of the three possible objects are equally likely to pop-up from behind their occluding boxes. Once sequence presentation begins, the model estimates the surprisal value of the current event at each item in the sequence. To do this, it combines the simple prior with the learner's previous observations from the sequence in order to form a posterior or updated belief. The next object pop-up event then conveys some surprisal value according to the probabilistic expectations of the updated belief.

In Kidd, Piantadosi, & Aslin (2010, 2012, under review), infants showed a "Goldilocks" effect for attentional termination: they were most likely to terminate their attention to the sequence at events that were either highly predictable or highly unexpected, according to the model. Thus, pop-up events to which the updated belief assigns either a very high probability (e.g., object *A*) or a very low probability (e.g., object *C*) would be expected to be more likely to cause infants to terminate their attention to the sequence than those to which it assigns an intermediate probability (e.g., object *B*).

In this example trial, the learner has seen a sequence composed of three *A* objects and one *B* object. The key question is how monkey learners will allocate their attention upon observing the next "pop-up" event in the sequence. For example, to test whether a stimulus' surprisal value predicts monkey's look-away behavior (as it does for infants), we could examine whether monkeys terminate the trial (as human infants did) upon observing the next "pop-up" event in the sequence. The observed object pop-ups (*AAAB*) comprise the observed data, which are combined with a flat, uninformative prior—essentially a smoothing term to avoid zero probabilities—to form an updated (posterior) belief. In this example, the updated belief leads to an expectation that the next event has a high probability of being object *A*, a moderate probability of it being object *B*, and a low (but non-zero) probability of it being object *C*. Thus, the predictability of the next pop-up event is quantified by an information theoretic metric—*surprisal,* which is simply the negative log probability the event's occurrence, as estimated by the model. The *surprisal* value represents the amount of "surprise" an idealized learner would have on seeing the next event, or, equivalently, the amount of information processing such a learner would be required to encode the event (Shannon, 1948). Thus, if the next object to pop-up is object *A*—an event that is highly likely according to the model's updated belief—the surprisal of that event would be low (i.e., the object pop-up would be highly predictable according to the model). The "Goldilocks" hypothesis thus holds that learners would be more likely to terminate their attention at this visual event. Conversely, if the next object to pop-up is object *C*—an event that is highly *unlikely* according to the model's updated belief—the surprisal of that event would be high (i.e., the object pop-up would be highly surprising according to the model). The "Goldilocks" hypothesis holds that learners should also terminate their attention to the visual sequence at this type of event. However, if the next object to pop-up is *B*—an event that is moderately

probable according to the model's updated belief—the complexity of that event would fall in the intermediate "Goldilocks" range, thus leading learners to be less likely to terminate their attention to the visual sequence.

A more formal presentation of the model is presented in Appendix V-1.

WHAT OUR BEHAVIORAL MEASURES OF CURIOSITY REVEAL.

Here we aim to discover whether a rational statistical model that is similar to those used with our infant data can explain monkeys' attentional patterns. In doing this, we hope to discover the factors that influence monkeys' intrinsic curiosity and the mechanisms that govern attentional allocation and learning.

First, we test for **predictive looks** in monkeys. We define predictive looks as saccades to the spatial area where the next event will take place, in advance of the onset of that event. Such anticipatory looking behavior has been demonstrated to reflect adult humans' ability to rapidly update linguistic expectations during language comprehension tasks (e.g., Altmann & Kamide, 2007). Our results reveal that monkeys exhibit clear anticipatory looking behavior. These results further indicate that monkeys are able to 1) monitor the statistical properties of incoming information streams, 2) rapidly form probabilistic expectations based on those incoming stimulus statistics, and 3) reallocate their attention accordingly.

Second, we collect a behavioral measure of **look-aways** for the monkeys, as we did with the infants[1]. Results of this analysis suggest that monkeys, like human infants, exhibit a U-shaped pattern of preference for events of intermediate values of surprisal / predictability. Monkeys are more likely to terminate attention

---

[1] We note that the look-away measure employed here with monkeys differs somewhat from that used with infants. Namely, in the infant experiments, the displays terminated contingently when infants looked away continuously for one second. Here, though we compute when look-aways occurs based on monkeys' visual fixations, the displays continue regardless of monkeys' attentional patterns. In our previous work with infants, the observed preference for intermediate complexity could depend upon being in a situation in which exploration is possible.

to highly predictable (low surprisal) events—and also highly unexpected (high surprisal) events. Thus monkeys also exhibit a "Goldilocks" effect in their visual attention behavior.

Finally, we examine **reaction times (RTs)** in monkeys. Previous studies in human adults have suggested that RTs reflect processing difficulty, as they exhibit a linearly increasing relationship with stimulus surprisal and a linearly decreasing relationship with stimulus discriminability (e.g., Grice et al., 1982; Sternberg, 1969; Teichner & Krebs, 1974). Typically, human adults have been reported to respond rapidly to highly predictable stimuli, but very slowly to that which is highly unexpected (e.g., Brysbaert et al., 2011). Our results instead reveal a different—and quite surprising—significant trend. The relationship between RTs and monkeys' expectations about stimulus predictability is U-shaped, with monkeys exhibiting the fastest RTs for intermediately predictable stimuli.

These results reflect the first behavioral evidence that monkeys' intrinsic curiosity, in the absence of any external rewards or explicit task goals, is governed by the statistical properties of stimuli in their environment. More importantly, these results are the first to demonstrate that monkeys' curiosity reflects their probabilistic beliefs and knowledge about the world, which are rapidly updated on a continuous basis as they make new observations. We take these results, along with results from previous work with human infants, as strong evidence that the ability to monitor the statistical properties of incoming information streams and selectively maintain or terminate attention based on these properties is a general property of intelligent organisms.

We note that the example shown in Fig. V-1 treats each event as statistically independent (a *unigram* model). However, our previous work in infant visual and auditory attention indicated that a model that tracked the conditional bigram probabilities (a *transitional* model) out-performed the non-transitional models—

both for attention to visual and auditory sequences (e.g., Kidd, Piantadosi, & Aslin, 2012, under review). In the present experiment, therefore, we also constructed and tested a transitional model of the visual sequence, which captured how likely each object pop-up was to follow each other object pop-up in the sequence. Note that for either model, the predictions of the model would be updated for each pop-up event in the sequence.

As we will see, monkeys exhibit a very different behavioral pattern. In contrast to infants, the *unigram* statistics are far more robust than the transitional statistics for monkey learners The robustness of the unigram statistics in predicting each of the behavioral measures may suggest an important difference between monkeys and human learners: infants may possess a sensitivity to transitional probabilities that monkeys lack. This difference is particularly intriguing in light of that fact that infants eventually acquire language, a serially ordered process, while monkeys do not.

Materials and Methods.

SUBJECTS.

We tested five juvenile male rhesus monkeys (*Macaca mulatta*) from the University of Rochester colony. Each subject had a small head-holding prosthesis for collecting high-resolution measurements of eye movements. Prior to these experiments, the subjects had been habituated to the lab and trained to perform oculomotor tasks for liquid rewards through standard reinforcement training (with only positive fluid rewards).

STIMULI.

We designed the displayed stimuli to be easily captured by a simple statistical model. Each trial featured one of 80 possible visual-event sequences (Appendix V-2). All sequences were presented to all subjects in a different randomized order. Only one sequence was presented per trial, and each was presented in the form of a unique animated display generated by a Matlab script.

Each animated display featured three identical white boxes in three distinct spatial locations on the screen (see Fig. V-2 for example display).
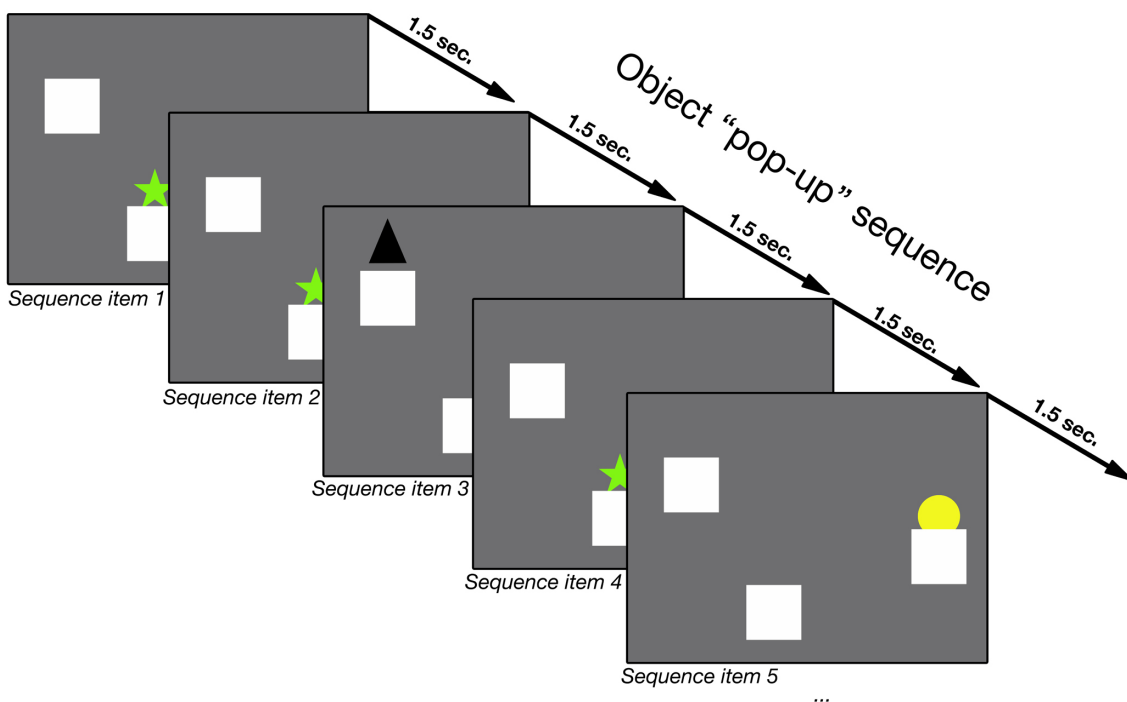
Fig. V-2. Example of Sequential Visual Display. The illustration shows four different time-points in the sequence. Each display featured three boxes, each occluding a unique geometric object (e.g., a green star). At each event in the sequence, one of the three objects popped up from behind one of three boxes.

The locations of the three boxes for a given sequence were chosen randomly, but remained static throughout the sequence. The spatial locations of the boxes in the display were randomized across trials; thus, each sequence was presented to each

subject with the three boxes in different randomized spatial configurations[2]. Each of the three boxes concealed one unique geometric object (e.g., a yellow triangle, a red star, and a blue circle). Like the spatial locations, the geometric objects remained associated with their respective boxes throughout the sequence, but were randomly selected from a large set of possible geometric objects across trials. Thus, each sequence was presented to each subject with the three boxes containing a different randomized set of geometric objects for each trial.

The experimental trials for this study were interspersed infrequently between experimental trials for other studies. Some of the studies employed positive reinforcement-learning paradigms, which reward subjects for certain responses with an increased amount of liquid.[3] In some of these studies, certain stimulus colors and shapes were associated with different-sized liquid rewards (e.g., Blanchard, Pearson, & Hayden, under review). We carefully selected the geometric objects used in this study to avoid features that had been previously associated with rewards in other studies. To achieve this goal, we composed stimuli by pairing only colors and shapes that had never previously been used in reward-reinforced experimental tasks with our subjects.

Each of the 80 sequences was conveyed by the order in which objects "popped-up" out of boxes in the displays. Each event within a sequence consisted of one of the three unique objects popping out from behind one of the three boxes (750 ms), and then back into the box (750 ms). Thus, the total duration of each object pop-up was 1.5 sec., and these pop-ups were presented sequentially with no overlap or delay. The 80 unique event sequences were composed such that the pop-

---

[2] The script that selected random spatial locations for boxes implemented a constraint to ensure that no two boxes (nor the objects popping out of them) occupied overlapping areas of the display within a single trial.

[3] Only positive reinforcement—not aversive conditioning—is ever used in any experimental lab task.

up probabilities of each of the three objects varied[4]. Some sequences contained many predictable pop-up events (e.g., *AAAAAAAAA*) while others contained many more surprising ones (e.g., *AAAAABABBAC*). Each of the 80 sequences was presented to each subject in full, though the order of sequence presentation was randomized across subjects.

The exhaustive randomization and counterbalancing of all extraneous variables (e.g., sequence order, object identity, color, shape, spatial location) across trials and subjects served to control for uncertainty about—and variation across—subjects' existing mental representations, processing speeds, and biases for stimulus salience. This design also enabled us to evaluate the effect of the predictability of each event in a sequence independently of either group or individual-subject biases. If, for example, all subjects had a bias for red objects, red objects could reduce the probability that subjects would look-away during these events; however, because red objects occur randomly at different points across different sequences across different subjects, the effect of both redness and sequence-order can be observed and evaluated.

FACILITIES AND EQUIPMENT.

The testing facility was built specifically for primate studies, located about 100 feet from the university colony in the same building. The room featured a computer monitor for stimulus presentation, floor plate for firmly mounting an ergonomic primate chair (Crist), and 1,000-Hz EyeLink infrared eye-tracking

---

[4] The 80 sequences comprised the same 32 sequences used in the infant attention experiments in (e.g., Kidd, Piantadosi, & Aslin, 2012; Piantadosi, Kidd, & Aslin, in press) and 48 additional sequences that were twice as long. The longer sequences made it possible to test the monkeys on events with higher surprisal values than would have been possible with the shorter sequences alone. An event can only have a high surprisal value when its occurrence violates strongly held probabilistic expectations on the part of the learner; thus, high surprisal events are only possible after an learner has made a enough observations to form the strong probabilistic expectations necessary for a new observation to violate them.

system (SR Research, Osgoode, ON) for sampling horizontal and vertical eye positions. We presented stimuli at a 1024 x 768 resolution on the computer monitor that was placed 144.8 cm (57 inches) in front of the subjects. We presented stimuli using a computer running Matlab (Mathworks, Natick, MA) with Psychtoolbox (Brainard, 1997; Kleiner et al, 2007; Pelli, 1997;) and EyeLink Toolbox (Cornelissen, Peters, & Palmer, 2002).

A standard solenoid valve controlled the duration of water delivery during the experiment (Parker). We measured fluid volumes associated with solenoid open time in order to ensure that fluid amounts were linearly proportional to the values instructed by the program. We confirmed that water delivery volume was constant regardless of the volume of water in the reservoir over the ranges used in this experiment. Fluid access was controlled outside of experimental sessions.

PROCEDURE.

The monkeys were eye-tracked as they watched the sequence of object pop-ups on our Matlab-generated visual displays (Fig. V-1). The automated liquid-delivery system delivered a 53µL water reward when each object was at its peak (every 1.5 sec.), regardless of where or whether the monkey was looking. The head-stabilized eye-tracking system restricted monkeys' head movements during the experiment, but the monkeys were still able to fixate off-screen or close their eyes during stimulus presentation. Regardless of monkeys' gaze behavior, each sequence (one per trial) was displayed in full.

As mentioned in the description of the experimental stimuli above, the 80 experimental trials for this study were interspersed infrequently between experimental trials for other studies. The rate of sequence presentation was

between 0 and 2 trials per day, interspersed within 400 - 2,000 unrelated trials for other studies.

We recorded all spatial and temporal details of the randomized, Matlab-generated sequential displays we presented, as well as all monkey visual fixations during the stimulus presentations. After the data were collected, we used a second set of scripts to compute *three dependent behavioral measures* from the raw timecourse data: **predictive looks, look-aways**, and **reaction times**. The computation of each of these attentional metrics involved measuring monkeys' fixations to and away-from certain active locations in which object pop-ups occurred on the display. These regions-of-interest were defined as including the area of the occluding box and the space immediately above it where an object appeared when it popped up (Fig. V-3).
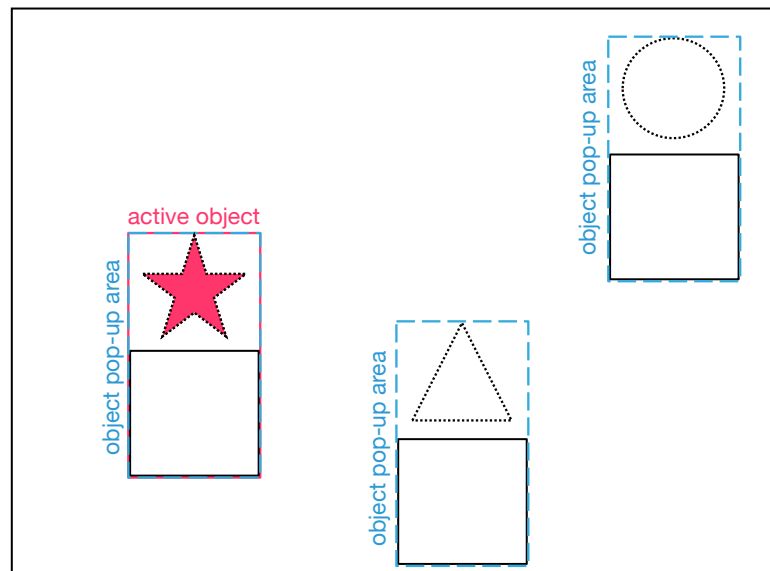
Fig. V-3. Regions of Interest in Sequential Visual Displays. This schematic shows the bounds of areas-of-interest used to compute the dependent attentional measures from the raw eye-tracking data. Object pop-up areas are outlined in blue. Only one object ever popped up at a time, so for each item in the sequence, only one object was active (area outlined in fuchsia).

First, we measured **predictive looks**. Predictive looks indicate whether the monkey was already looking at the current object area when it first became active (before the pop-up), excluding events in the sequence that were immediate repeats of the same object that was most recently active. Next, we computed **look-aways** according to an *a priori* criterion that was picked to indicate a significant drop in attention to the stimulus sequence. For the purpose of our analyses, we defined look-aways as less than 50% looking at active areas of the display (i.e., box and object areas) during a single pop-up event in a sequence. Thus, the look-away criterion was met whenever more than 50% of a monkey's fixations during a pop-up event were either to blank areas of the display or off-screen, or whenever a monkey closed their eyes for more than 50% of a pop-up event[5]. Most commonly, monkeys met the look-away criterion through off-screen looks to just below the center of the monitor display. The final dependent measure, **reaction times**, encoded monkeys' saccade latency (in ms) before first visually fixating a currently active pop-up event. We excluded reaction times that indicated monkeys were already fixating the target area at the onset of the pop-up event from our analyses since these values do not truly represent saccade latencies.

Analysis.

Analysis of the behavioral data followed the approach used for the infant studies in Kidd, Piantadosi, & Aslin (2010, 2012, under review) and Piantadosi, Kidd, & Aslin (in press). A Markov Dirichlet-multinomial model first quantified an idealized learner's expectation that each of the three objects would pop-up next,

---

5 Monkeys could theoretically meet the look-away criterion through a combination of blank looks, off-screen looks, and eye-closing acts that totaled more than 50% of the pop-up event's duration. This did not, however, seem to usually be the case; sustained off-screen looks were the most common way in which monkeys met the look-away criterion.

at each point in the sequence. This rational model essentially combines a "smoothing" term—or prior expectation of object pop-up likelihood—with counts of how often each object has previously popped-up in the sequence to predict each object's probability of popping-up next. The model then takes the negative log of that estimated probability for each object in order to a compute its *surprisal,* corresponding to how *predictable* or *surprising* an idealized learner would find that event if the object popped up. This measurement has also been referred to as *information content*, as the resulting value corresponds to how many bits of information an idealized learner would require to remember or process each pop-up event (Shannon, 1948). Since this model assumes event-order independence (i.e., it only tracks the zero-order statistics), we refer to it as the *unigram* model.

We also applied the MDM model to the data under an assumption of event-order *dependence.* That is, instead of treating every object pop-up event as independent, we examined whether attentional behavior was predicted by a probability conditioned on the immediately preceding event (i.e., given that the last pop-up was object *A*, how surprising is it to see object *A* again?). This version of the model will allow us to determine to what degree subjects track *transitional statistics*—and to what degree these transitional statistics influence their attentional behavior. We refer to this version as the *transitional* model.

The surprisal estimates obtained from the *unigram* and *transitional* models were then entered into mixed effect linear and logistic regressions of the behavioral data[6]. Surprisal was entered as both a *linear surprisal* and a *quadratic surprisal*

---

[6] We note that we used stepwise Cox regressions in Kidd, Piantadosi, & Aslin (2010, 2012, under review) rather than the more standard regression varieties used here. The reason is that the infant experimental designs employed necessarily violated the independence assumption of standard regressions (Hosmer, Lemeshow, & May, 2008; Klein & Moeschberger, 2003); the monkey experimental designs commit no such violation.

term. In previous work in human adults, surprisal is linearly related to, for instance, reading time on words (e.g., Smith & Levy, 2013), where surprisal is measured according to a word's predictability in context. In our previous work in human infants, we demonstrated that *quadratic surprisal* of a visual stimulus predicts when infants are most likely to terminate their attention to a sequence (the "Goldilocks" effect). Infants exhibited a U-shaped preference for stimuli with intermediate surprisal values; infants were most likely to terminate their attention at points in a sequence characterized by very low or very high surprisal values, corresponding to stimulus items that were highly predictable or highly unexpected. This pattern differs critically from Smith & Levy (2013) in that it used an attentional termination behavioral measure, rather than a reaction time measure. To our knowledge, no work has examined the relationship between surprisal and behavioral measures in primates, using any kind of behavioral measure. As a result, we focus on both exploratory and hypothesis-testing analyses that could potentially recover either linear or quadratic effects (or others) across all of the monkey behavioral measures we collect. These analyses allowed us to test the significance of the linear and quadratic surprisal terms while controlling for a number of other factors that could plausibly influence subjects' attentional behavior.

Table V-1 lists descriptions of each of the control covariates. These included whether the current object is popping up for the first time, the number of unobserved objects, and temporal factors (e.g., sequence item, trial number). The regressions also included random intercepts, and both linear and quadratic predictability slopes by subject.

Table V-1. Factors and Control Covariates.

Factor	Description
Surprisal	*Quantity expressing how unexpected the current object pop-up was according to a subject's current probabilistic beliefs about the distribution of pop-up events, as estimated by the ideal observer model; the negative log probability of the current object pop-up according to the updated/posterior model beliefs*
Squared surprisal	*Quadratic term to test for a U-shaped relationship between surprisal and attentional behavior*
First occurrence	*Whether the current object is popping up for the first time (binary)*
Unseen objects	*The number of unobserved objects (0, 1, or 2)*
Repeat	*Whether the current object pop-up was an immediate sequential repeat (binary)*
Trial number	*Number of trials into the current experiment (not including interspersed trials from other experiments); also indicates number of observed sequences since each trial features only one sequence (1 – 64)*
Sequence item	*Number of items (event pop-ups) into the current sequence (1 – 60)*
Distance	*Linear distance between the previous and current object pop-ups, as measured in pixels from the box centers (0 – 926)*
All continuous and ordinal factors were standardized—shifted and scaled to have a mean of 0 and standard deviation of 1—before they were entered into the regressions in order to make the resulting coefficients easy to interpret.	

The general approach of relating the surprisal estimate from the computational model to the monkey behavioral data is to visualize the data using a Generalized Additive Model (Hastie & Tibshirani, 1990). Statistical significance was evaluated using mixed effect linear and logistic regressions with random intercepts, and linear and quadratic surprisal slopes by subject. This approach enables testing for the significance of linear and U-shaped trends while controlling for variation between subjects, as well as a number of control factors also expected to influence attentional behavior. The data visualization is critical for ensuring that significant quadratic trends correspond to U-shapes over the range of the tested data. This approach prevented us from erroneously conjecturing that the true relationship between the model's estimates of surprisal and monkeys' attentional behavior was U-shaped when it was actually better described by some other

function (e.g. the function 1/*x*, which might yield a significant quadratic effect given sufficient statistical power, but which is asymptotically flat—not a U).

Summary of Main Results.

Here we outline the main results of our analysis, which aimed to test several factors expected to influence monkeys' intrinsic curiosity and the mechanisms that govern attentional allocation and learning. Exhaustive reporting of regression results for all models and factors appears in the *Detailed Results by Behavioral Method* section, which follows.

ANTICIPATORY LOOKING TOWARDS EXPECTED EVENTS.

Monkeys were more likely to predictively look at more predictable object pop-up events, according to the significant linear surprisal term of the controlled unigram predictive-looks regression ($\beta = -0.25$, $z = -3.89$, $p < 0.0001$). Moreover, this effect was quite robust. The transitional version of the model also yielded similar trends (Table V-2 and Fig. V-5). These result suggest that, like adult humans, monkeys can rapidly update their expectations in accordance with the statistical properties of incoming information streams and reallocate their attention appropriately.

PREFERENTIAL ATTENTION TOWARDS UNKNOWNS.

Monkeys also exhibited increased curiosity for unknowns in the visual displays. The controlled unigram predictive-looks regression revealed that monkeys produced more predictive looks when there were more previously unseen objects

($\beta = 0.23$, $z = 5.97$, $p < 0.0001$) and on an object's first appearance ($\beta = 0.41$, $z = 6.78$, $p < 0.0001$). These results suggest that monkeys exhibited increased visual interest in boxes that had not yet revealed their contents. The transitional-model regression also revealed the indicators of information-seeking—more predictive looks for previously unseen objects ($\beta = 0.30$, $z = 7.82$, $p < 0.0001$) and on an object's first appearance ($\beta = 0.55$, $z = 9.76$, $p < 0.0001$).

PREFERENCE FOR EVENTS OF INTERMEDIATE SURPRISAL.

Monkeys are more likely to terminate attention to highly predictable (low surprisal) events—and also highly unexpected (high surprisal) events—as estimated by the unigram GAM analysis. Fig. V-5 shows the U-shaped relationship between look-away probability and surprisal, as estimated by the unigram GAM. This quadratic trend is statistically significant in the regression that considers only the surprisal and squared surprisal measures ($\beta = 0.10$, $z = 3.44$, $p < 0.001$, Table V-3 and Fig. V-5). However, a more conservative regression that includes other predictors yields only a statistically robust linear trend ($\beta = -0.16$, $z = -1.88$, $p < 0.06$) and no statistically robust quadratic trend ($\beta = 0.02$, $z = 0.69$, $p = 0.49$), most likely due to data sparsity in the highest surprisal range (right side) of the U. An additional by-subject analysis revealed U-shaped functions for the majority of subjects, in further support of this theory (Fig. V-10). Taken together, these results suggest that the true relationship may be U-shaped, although this conclusion is preliminary. Thus, like human infants, monkeys also appear to exhibit a "Goldilocks" effect in their visual attention behavior.

ROBUSTNESS OF UNIGRAM OVER TRANSITIONAL STATISTICS.

Unlike human infants—for whom the look-away model that tracked the conditional bigram probabilities out-performed the unigram model—the *unigram* statistics are for more robust than transitional statistics for monkey learners. (See Table V-3 and Fig. V-6.) This difference may indicate that infants possess sensitivity to transitional probabilities that monkeys lack. This difference could be relevant in explaining cross-species differences in linguistic capabilities. (Infants eventually acquire language, while monkeys do not.) The robustness of the unigram statistics over the transitional ones was evident not only in monkeys' look-away behavior, but also for monkeys anticipatory looking and reaction times (see *Detailed Results by Behavioral Measure*.)

REACTION TIMES AS MEASURE OF INTEREST.

In stark contrast to previous work with human adults, the relationship between stimulus predictability and monkeys' reactions times is U-shaped. In previous studies, adult human RTs appeared to reflect processing difficulty. Human adults have been widely reported to respond rapidly to highly predictable stimuli, but very slowly to that which is highly unexpected (e.g., Brysbaert et al., 2011; Smith & Levy, 2013). Our results instead reveal a different—and quite surprising—significant effect. The relationship between RTs and monkeys' expectations about stimulus predictability is U-shaped, with monkeys exhibiting the fastest RTs for *intermediately predictable* stimuli (as revealed by the significant quadratic factor in the controlled regression for the unigram model, $\beta = 5.23$, $z = 2.31$, $p < 0.02$). These results may suggest that rather than serving as a measure of processing difficulty, as previous studies of adult humans have suggested, RTs may also encode the learner's *level of arousal* or interest in a stimulus.

SIGNIFICANCE OF RESULTS.

These findings represent the first to our knowledge to demonstrate that monkeys' curiosity reflects their probabilistic beliefs and knowledge about the world. Further, these findings provide evidence that monkeys rapidly and continuously update their probabilistic beliefs, and that these updates are reflected in their gaze behavior. More broadly, these results suggest that the ability to monitor the statistical properties of incoming information—and the ability to selectively maintain or terminate attention based on these properties—is a general characteristic of intelligent organisms.

Detailed Results by Behavioral Measure.

PREDICTIVE-LOOK MODELS.

Table V-2 contains the predictive-looking measures. The predictive-looking measures indicate whether the monkey was already looking at the current object when it first became active (before the object actually popped up), excluding events in the sequence that are immediate repeats of the same object that was most recently active.

Table V-2. Surprisal Term Coefficients for Predictive-Looks Regression.

		Linear		Quadradic (U)		GAM Trend
		Coef.	P-value	Coef.	P-value	
Unigram	*(a)* RAW	−0.51	***	0.09	***	Decreasing
	(b) CONTROLLED	−0.25	***	0.01	ns	Decreasing
Transitional	*(c)* RAW	−0.15	**	−0.01	ns	Decreasing
	(d) CONTROLLED	−0.08	ns	−0.05	**	Decreasing
Significance codes: ' *** ' < 0.001 ' ** ' < 0.01 ' * ' < 0.05 '·' < 0.1 'ns' < 1						

Unigram model for predictive looks.

Estimated GAM curves relating predictive looks to the unigram surprisal measure is shown in Fig. V-4, for both the raw data (*a*) and the highly conservative controlled version (*b*). The results of the related regression analyses appear in Table V-2*a-b*.
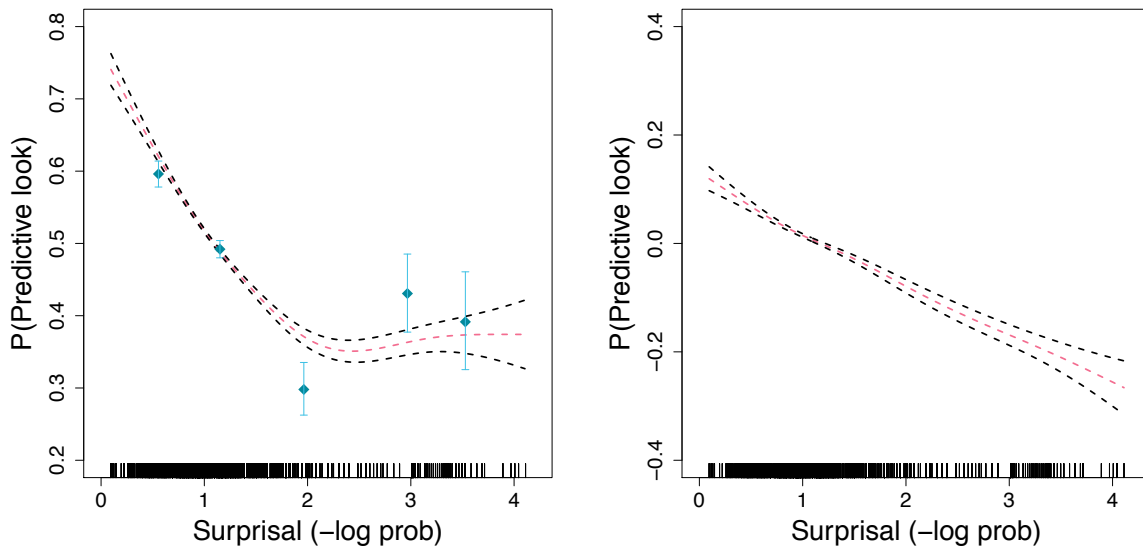
Fig. V-4. Predictive-Look Probability as a Function of Unigram Surprisal. *(a)* Monkeys' probability of predictively looking to the next active object (*y*-axis) as a function of surprisal (*x*-axis) as measured by the unigram model; the smooth curve shows the fit of a generalized additive model with standard errors. *(b)* Predictive-look probability (*y*–axis) and unigram surprisal (*x*-axis), while controlling for all factors described in Table V-1**.** Both plots depict significant decreasing linear trends, with monkeys looking more often at the most predictable events according to the unigram model.

These figures reveal a very clear linearly decreasing trend, with a higher probability of predictive looks for the lowest surprisal values. This indicates that monkeys were more likely to predictively look at more predictable object pop-up events. The trend is evident in the raw data (Fig. V-4*a*) and becomes even more transparent in the conservative controlled version (Fig. V-4*b*). This pattern of results suggests that the effect of surprisal value on predictive looks is quite robust. The regression analyses (Table V-2*a-b*) provide further evidence for this claim. The raw regression

reveals significant linear and quadratic terms ($\beta = -0.50$, $z = -8.49$, $p < 0.0001$; $\beta = 0.09$, $z = 4.97$, $p < 0.0001$), corresponding to the non-linear monotonic decreasing trend in Fig. V-4*a*. The controlled regression reveals only a highly significant effect for the decreasing linear term ($\beta = -0.25$, $z = -3.89$, $p < 0.0001$), which is likely indicative of the true relationship between predictive looks and surprisal. The slight uptick on the right side of the function in Fig. V-4*a* thus likely reflects correlations with other factors, which is why the quadratic term is no longer significant in the controlled regression.

The regression also revealed a number of other significant effects for the unigram model (Appendix V-3). Monkeys produced more predictive looks in later trials ($\beta = 0.25$, $z = 11.96$, $p < 0.0001$) and for later sequence items ($\beta = 0.55$, $z = 20.80$, $p < 0.0001$). Interestingly, monkeys also produced more predictive looks when there were more previously unseen objects ($\beta = 0.23$, $z = 5.97$, $p < 0.0001$) and on an object's first appearance ($\beta = 0.41$, $z = 6.78$, $p < 0.0001$). These effects are consistent with the ideas that monkeys may have exhibited indictors of curiosity for unknowns in the visual display. These results may reflect that monkeys exhibited increased visual interest in boxes with as-of-yet unknown contents. Greater predictive looks for more unseen objects and on an object's first appearance could indicate that monkeys were visually checking in on boxes that had not yet revealed their contents.

Finally, there was a very small but significant effect of distance ($\beta = 0.09$, $z = 2.90$, $p < 0.01$), unintuitively indicating that monkeys produced more predictive looks for objects at a greater distance.

Transitional model for predictive looks.

Estimated GAM curves relating predictive looks to the transitional surprisal measure are shown in Fig. V-5 for both the raw data (*a*) and the highly

conservative controlled version (*b*). The results of the related regressions appear in Table V-2*c-d*.
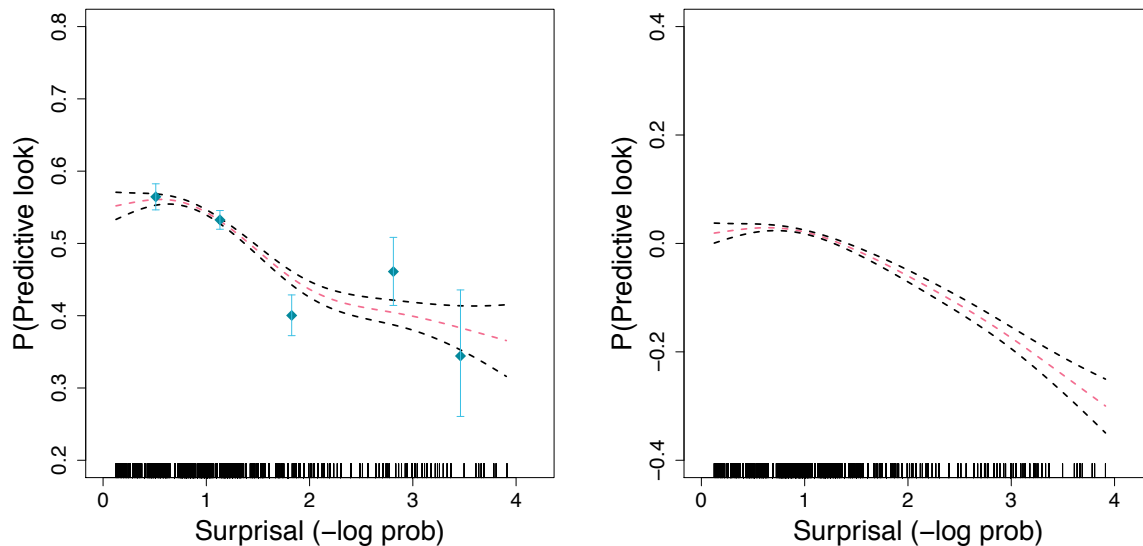
Fig. V-5. Predictive-Look Probability as a Function of Transitional Surprisal. *(a)* Monkeys' probability of predictively looking to the next active object (*y*-axis) as a function of surprisal (*x*-axis) as measured by the transitional model; the smooth curve shows the fit of a generalized additive model with standard errors. *(b)* Predictive-look probability (*y*–axis) and transitional surprisal (*x*-axis), while controlling for all factors described in Table V-1. As in the unigram version (Fig. V-8), both these predictive-looking plots depict decreasing linear trends, with monkeys looking more often at the most predictable events according to the unigram model.

Like the unigram versions, these transitional plots suggest a non-linear monotonic decreasing trend, with decreasing predictive looks for the most unexpected (high surprisal) pop-up events. The trend is evident in both the raw data (Fig. V-5*a*) and remains apparent in the conservative controlled version (Fig. V-5*b*). These results suggest that the effect of surprisal value on predictive looks is quite robust for both unigram and transitional statistics. The trends for both the linear and quadratic terms in the regressions (Table V-2*c-d*) also suggest monotonic decreasing trends in predictive looks as the transitional surprisal value increases. The raw regression (Table V-2*c*) yielded a significant linear effect for the transitional surprisal term ($\beta = -0.15$, $z = -2.85$, $p < 0.005$), but a non-significant quadratic effect ($\beta = -0.005$,

$z = -0.28$, $p = 0.78$). The controlled regression (Table V-2d) yielded a seemingly opposite pattern of results: a significant quadratic effect ($\beta = -0.05$, $z = -2.59$, $p < 0.01$), but a non-significant linear effect that trends in the expected decreasing direction ($\beta = -0.08$, $z = -1.27$, $p = 0.21$). As is apparent in Fig. V-5b, the overall trend of the function is decreasing and thus the significant quadric term in the control regression reflects the non-linearity of this decreasing function, not a U.

The transitional-model regression also revealed all of the same significant predictors of predictive looking as found for the unigram version (Appendix V-3). As before, monkeys produced more predictive looks over time—in later trials ($\beta = 0.25$, $z = 11.80$, $p < 0.0001$) and for later sequence items ($\beta = 0.58$, $z = 22.80$, $p < 0.0001$). Again, monkeys also exhibited indicators of information-seeking behavior: more predictive looks for previously unseen objects ($\beta = 0.30$, $z = 7.82$, $p < 0.0001$) and on an object's first appearance ($\beta = 0.55$, $z = 9.76$, $p < 0.0001$). There was again a small but significant effect of distance in the same direction as before ($\beta = 0.09$, $z = 2.88$, $p < 0.005$), with monkeys producing more predictive looks for objects at a greater distance.

LOOK-AWAY MODELS.

The look-away measure encodes whether a monkey looked away from the object-relevant areas of the display for more than the *a priori* criterion of 50% of the duration of an object pop-up event[7]. Table V-3 shows a summary of the

---

[7] Though we chose 0.5 as the value for our *a priori* look-away criterion, we note that the particular value chosen for this criterion does not significantly change either the GAM trend or the regression analyses. As the criterion value approaches 1, the stricter look-away criterion yields fewer overall look-aways. As it approaches 0, it becomes more lax and thus yields a greater number of overall look-aways. However, our analyses, which relate look-aways to stimulus surprisal, are designed to uncover the linking function that best describes the relationship between these two measures. Thus, they are not differentially impacted by having overall more or fewer looks meet criterion.

logistic regressions predicting this measure of attentional termination from the surprisal values across both unigram and transitional measures.

Table V-3. Surprisal Term Coefficients for Look-Away Regression.

		Linear		Quadradic (U)		GAM Trend
		Coef.	P-value	Coef.	P-value	
Unigram	*(a)* RAW	−0.44	***	0.10	***	U
	(b) CONTROLLED	−0.16	·	0.02	ns	U (sparse right)
Transitional	*(c)* RAW	−0.21	***	0.05	*	U
	(d) CONTROLLED	0.00	ns	−0.03	·	Decreasing
Significance codes: ' *** ' < 0.001 ' ** ' < 0.01 ' * ' < 0.05 '·' < 0.1 ' **ns**' < 1						

Unigram model for look-aways.

Fig. V-3*a* shows the relationship between the unigram measures and the raw look-away measure, not controlling for other variables. This plot shows the GAM model's estimated relationship between look-away probability and surprisal. In addition, the bars in this plot show binned raw data (look-away probability in each binned value of surprisal). This relationship exhibits a clear U-shape, such that monkeys are least likely to look away for intermediate surprisal values: events that are particularly predictable or particularly unexpected according to the unigram Markov Dirichlet-Multinomial model are most likely to trigger criterion inattention to the box displays. As Table V-3*a* shows, this quadratic trend is statistically significant in the regression that considers only the surprisal and squared surprisal measures.

Fig. V-6. Look-Away Probability as a Function of Unigram Surprisal. *(a)* Monkeys' probability of looking away (*y*-axis) as a function of surprisal (*x*-axis) as measured by the unigram model. The points and error bars show the raw probability of looking away; the smooth curve shows the fit of a generalized additive model with standard errors. Vertical tick marks show values of surprisal attained in the experiment. A very clear U-shape exists between unigram predictability and *p*(look-away), with monkeys least likely to look away at moderately surprising events. *(b)* The relationship between look-away probability (*y*–axis) and unigram surprisal (*x*-axis), while controlling for all factors described in Table V-1. This still exhibits a clear U-shaped pattern**.**

Fig. V-6*b* shows a GAM visualization in which the other covariates are controlled. This plot shows the GAM's estimated relationship between surprisal and look-away probability (note that the that *y*-value may be negative since other covariates are also included). This plot suggests the presence of a U-shaped relationship even when the other factors are controlled. However, the logistic regression (Table V-3*b*) reveals only a statistically robust linear trend ($\beta = -0.16$, $z = -1.88$, $p < 0.07$) and no statistically robust quadratic trend ($\beta = 0.02$, $z = 0.69$, $p < 0.49$). The controlled GAM plot shows that the sparsity of data on the right may be responsible for this result. Taken together, these results could indicate that the true relationship may be U-shaped but that we do not have a sufficient number of data points in a high enough surprisal-value range to yield significance for the right half of the U. Consistent with this account, a by-subject analysis revealed U-shaped functions for the majority of monkey subjects (Fig. V-10).

Interestingly, this regression also reveals a number of significant effects of the other covariates (Appendix V-4). Monkeys are significantly less likely to terminate attention on repeated events ($\beta = -0.09$, $z = -2.40$, $p < 0.02$). They are more likely to terminate attention to items on their first appearance ($\beta = 0.47$, $z = 8.84$, $p < 0.001$) and when there are many unseen objects ($\beta = 0.29$, $z = 9.93$, $p < 0.001$). They are more likely to terminate attention in later trials ($\beta = 0.54$, $z = 28.05$, $p < 0.001$) and later in sequences ($\beta = 0.70$, $z = 28.3$, $p < 0.001$), likely reflecting patterns of within-trial boredom and cross-trial fatigue respectively. There was no effect of distance ($\beta = 0.05$, $z = 1.46$, $p = 0.15$) on look-aways. This suggests that look-aways are not driven by physical constraints (e.g. having to saccade very far), but are instead based on the abstract, statistical properties of the stimuli.

The significance of the U-shape and surprisal measures in the raw regression but not the controlled one may suggest that the observed U-shape is an artifact of other variables—perhaps the other factors in the regression drive monkey's look-away behavior and once these are controlled, there is no effect of surprisal. On the other hand, the opposite is also possible: the true causal force may be the U-shape, and this happens to be correlated to some degree with other variables; once these variables are controlled, there is no significant U-shaped trend. We note that there are no results or theory *a priori* to predict that these other variables matter, and for some—like the number of unseen objects—the direction of effect is even opposite of what might have been expected *a priori*. From this point of view, the controlled regression is conservative, as it checks for an effect once every other plausibly important variable we can think of has been accounted for. However, as revealed by the controlled GAM plot, even when these variables are controlled, there is still a trend towards a U-shape, indicating that the true function may be U-shaped, but

just not highly robust. The by-subject analysis for the unigram model in Fig. V-10 is in line with this theory.

In either case, the linear trend on surprisal is relatively robust ($p < 0.06$ in the controlled regression), suggesting that model-based measures of predictability do influence looking. This decreasing trend goes in the most plausible direction, with monkeys least likely to terminate their attention to highly surprising events.

Transitional model for look-aways.

Results from the transitional model are shown in Table V-3*c-d* and Fig. V-7. Here, there is a clear U-shaped relationship in the raw model and model fits (Fig. V-7*a*, Table V-3*c*). However, this trend disappears entirely when other variables are controlled, both in the GAM model (Fig. V-7*b*) and the logistic regression (Table V-3*d*). Indeed, the regression does not find a significant linear component, likely because the estimated GAM curve is not monotonic, but largely flat. This suggests that the transitional model does not well-predict monkey's look-away patterns, in contrast to the robust linear and likely quadratic trends exhibited by the unigram model. This conclusion is further supported by the by-subject analysis for the transitional model in Fig. V-11.
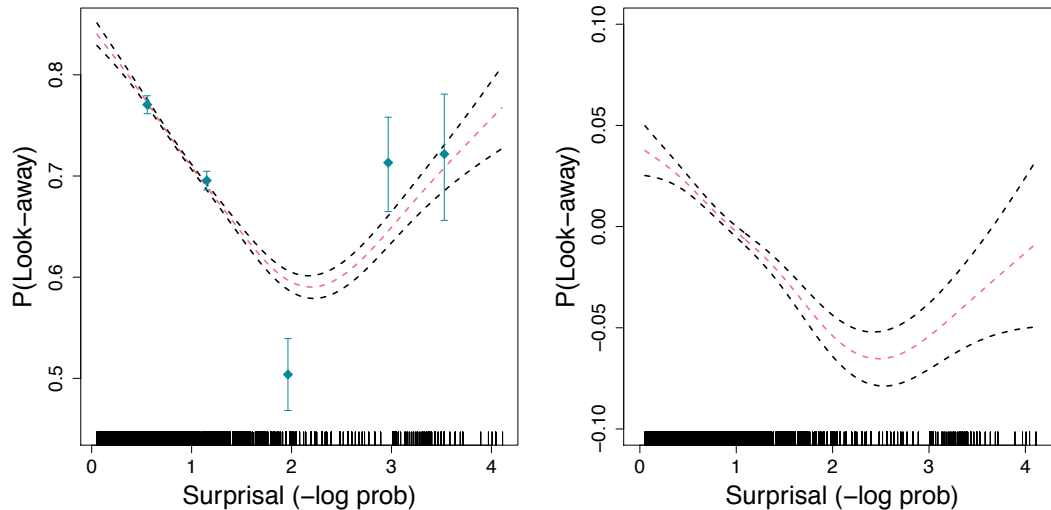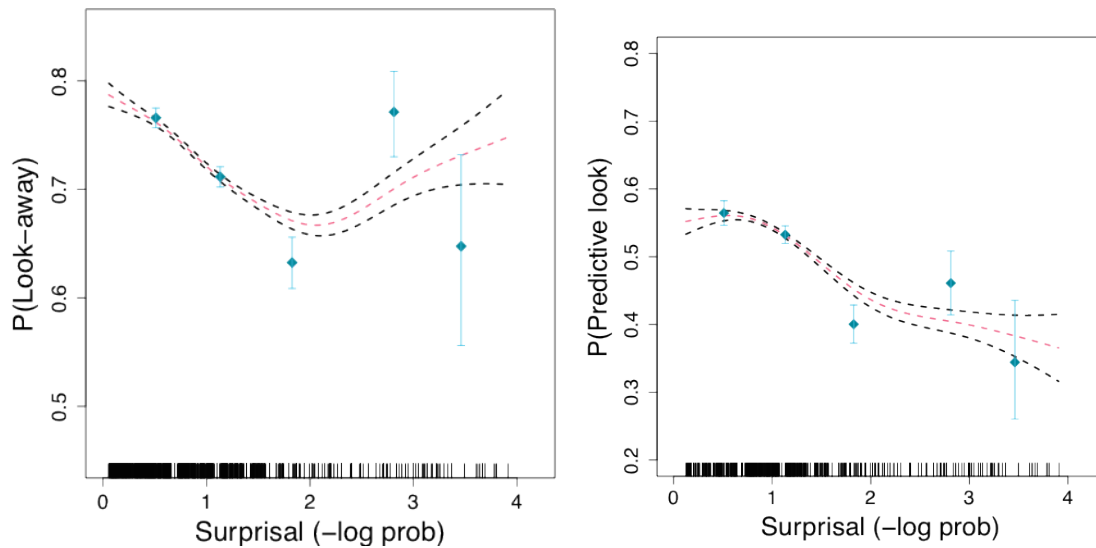
Fig. V-7. Look-Away Probability as a Function of Transitional Surprisal. *(a)* Monkeys' probability of looking away (*y*-axis) as a function of surprisal (*x*-axis) as measured by the transitional model. The points and error bars show the raw probability of looking away; the smooth curve shows the fit of a generalized additive model with standard errors. *(b)* Look-away probability (*y*–axis) and unigram surprisal (*x*-axis), while controlling for all factors described in Table V-1. Though the function on the left appears U-shaped, the function becomes a largely flat, decreasing curve once other factors are controlled (right).

As with the unigram model, this analysis revealed robust effects of other covariates in the same directions and with nearly identical magnitudes (Appendix V-4): monkeys were less likely to look away on repeated events for this analysis ($\beta = -0.11$, $z = -3.08$, $p < 0.005$), more likely to look away on the first appearance of an object ($\beta = 0.58$, $z = 11.94$, $p < 0.001$), more likely to look away later in trials ($\beta = 0.54$, $z = 27.89$, $p < 0.001$) and later with sequences ($\beta = 0.74$, $z = 29.97$, $p < 0.001$), and more likely to look away when there were more unseen objects ($\beta = 0.36$, $z = 13.08$, $p < 0.001$). There was no effect of saccade distance on look-away probability ($\beta = 0.04$, $z = 1.31$, $p = 0.19$).

REACTION-TIME MODELS.

We next report reaction time measures, shown in Fig. V-8 and Fig. V-9, and in Table V-3. The reaction time measures the monkeys' latency to arrive at the object that is currently active, looking only at events in the sequence in which the monkey is not already looking at the object (and thus, must react to view it).

Table V-4. Surprisal Term Coefficients for Reaction-Time Regression.

		Linear		Quadradic (U)		GAM Trend
		Coef.	P-value	Coef.	P-value	
Unigram	*(a)* RAW	−64.1	***	14.2	***	U
	(b) CONTROLLED	−23.8	**	5.2	*	U
Transitional	*(c)* RAW	−33.2	**	7.3	*	Shallow U
	(d) CONTROLLED	−4.9	ns	1.1	ns	Flat
Significance codes: '***' < 0.001 '**' < 0.01 '*' < 0.05 '·' < 0.1 'ns' < 1						

Unigram model for reaction times.

Estimated GAM curves relating look-away probability to the unigram surprisal measure is shown in Fig. V-8, for both the raw data (*a*) and the controlled version (*b*). Results of the related regression analyses appear in Table V-4*a-b.*

Fig. V-8. Reaction Time as a Function of Unigram Surprisal. *(a)* Monkeys' reaction time (latency) to fixate the active object (y-axis) as a function of surprisal (x-axis) as measured by the unigram model; the smooth curve shows the fit of a generalized additive model with standard errors. *(b)* RT (*y*–axis) and unigram surprisal (*x*-axis), while controlling for all factors described Table V-1. These plots depict significant U-shaped functions in both the raw data and once other factors are controlled.

These figures reveal that the raw data exhibits a very clear U-shaped relationship, with latency lowest for moderately surprising events. The U-shape holds, even when we examine the highly conservative estimate that includes all control predictors, though the U becomes more shallow (Fig. V-8*b*). We note that though the controlled plot spans a different numerical range (due to the fact that other variables are controlled), the relative range and scaling of the axes is identical. In addition to the U-shaped trend in the controlled GAM, there is also a slightly decreasing linear trend, with monkeys slower to fixate events that are more unexpected (higher surprisal). The regressions here (Table V-4*a-b*) reveal significant linear and quadratic trends for both the raw and controlled unigram model. The significance of the quadratic terms likely corresponds to a genuine U over the range of surprisal, especially in light of the fact that the significance holds even in the controlled GAM. The significance of the linear terms indicates that

there is also an increasing trend in RT for higher surprisal values for the curve shown in Fig. V-8*b*.

As with the regression predicting look-aways, the regression reveals a number of other significant effects, including faster looks to repeated items ($\beta = -46.56$, $t = -6.21$, $p < 0.001$), slower looks to items on their first appearance ($\beta = 36.00$, $t = 3.68$, $p < 0.001$), slower looks for later trials ($\beta = 15.51$, $t = 4.02$, $p < 0.001$) and later sequence items with a trial ($\beta = 24.34$, $t = 4.90$, $p < 0.001$). There was a marginally significant effect of distance, with longer distances taking longer to respond to trials ($\beta = 13.76$, $t = 1.95$, $p < 0.06$), and no effect of how many items were unobserved ($\beta = 7.96$, $t = 1.32$, $p = 0.20$). (See Appendix V-5 for further details.)

Transitional model for reaction times.

Curves showing the relationship between reaction time and transitional predictability are shown in Fig. V-9, for both the raw data (Fig. V-9*a*) and the controlled version (Fig. V-9*b*). The results of the related regression analyses appear Table V-4*c-d*.
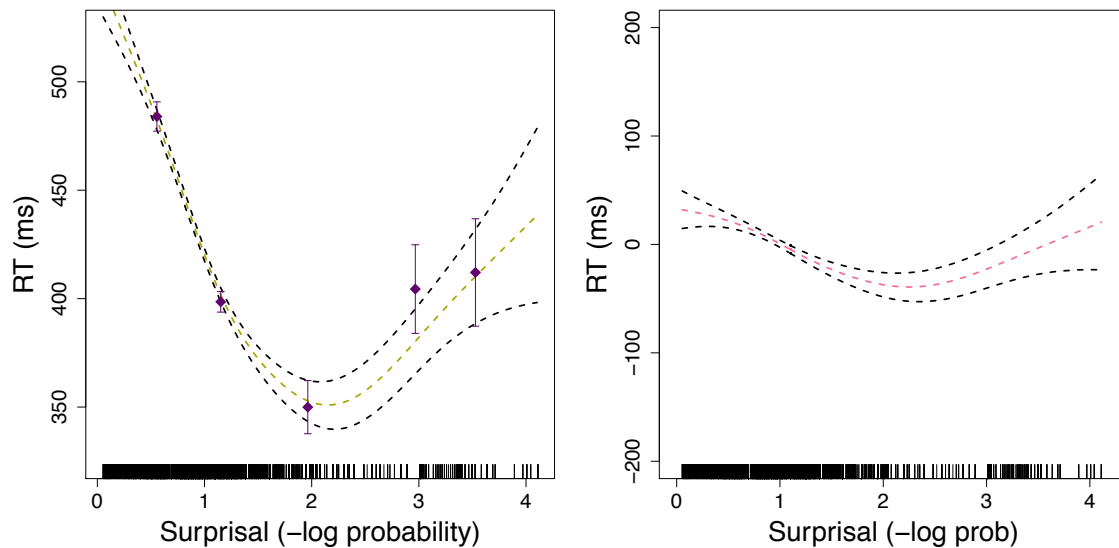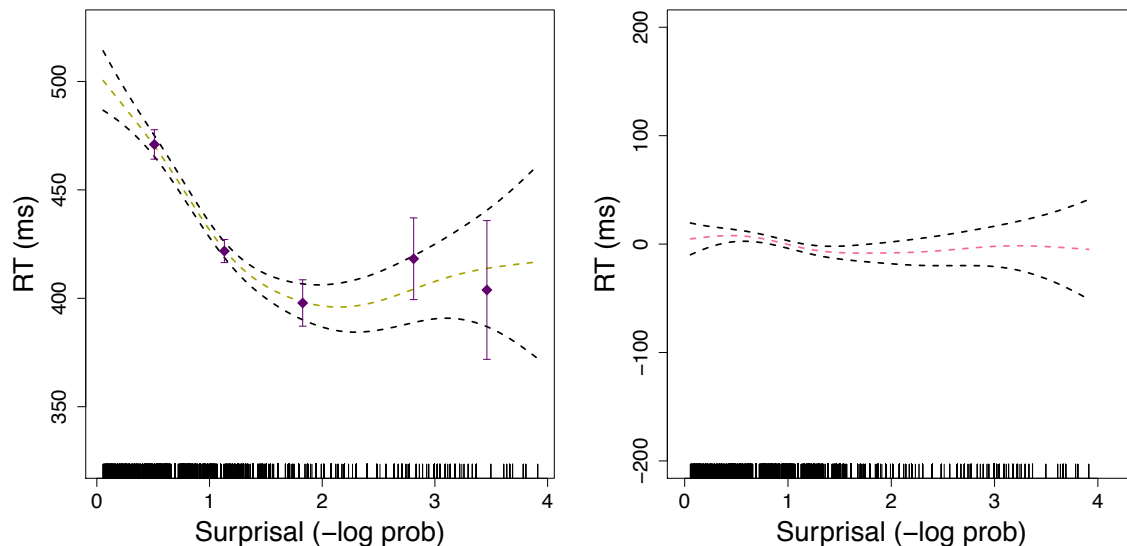
Fig. V-9. Reaction Time as a Function of Transitional Surprisal. *(a)* Monkeys' reaction time (latency) to fixate the active object (*y*-axis) as a function of surprisal (*x*-axis) as measured by the transitional model; the smooth curve shows the fit of a generalized additive model with standard errors. *(b)* RT (*y*–axis) and transitional surprisal (*x*-axis), while controlling for all factors described in Table V-1. As in the transitional version of the look-away analysis, the function relation RT to transitional surprisal becomes flat once other factors are controlled.

The GAM curve in Fig. V-9*a* reveals that the raw data exhibits an overall non-linear monotonic decreasing trend, with lower latencies for higher surprisal values. Consistent with the GAM visualization, the regression analysis for the raw transitional RT data reveals significant linear and quadratic trends (Table V-4*c*). This pattern of results suggests that as the transitional probabilities between pop-up events became more unexpected, monkeys were faster to fixate. Such a pattern of results might occur, for example, if monkeys were bored with the most predictable transitions between objects; faster RT might occur due to increased interest in more surprising events. However, the controlled GAM curve in Fig. V-9*b* suggests otherwise. Once we include all control predictors, the curve becomes exceptionally flat. This pattern is further supported by the controlled regression for the transitional RT data (Table V-4*d*), which failed to find significance for either the linear or quadratic terms. The dramatic difference between the raw and

controlled plots suggests that the decreasing trend in the raw data does not reflect the true relationship between RT and TP surprisal, but was rather driven by correlations with other factors.

Consistent with this idea, the controlled regression reveals a number of other significant effects, including all of the factors that were significant in the unigram version of the model (Appendix V-5). These include faster looks to repeated items ($\beta = -50.77$, $t = -6.86$, $p < 0.001$), slower looks to items on their first appearance ($\beta = 47.11$, $t = 5.39$, $p < 0.001$), slower looks for later trials ($\beta = 16.78$, $t = 4.32$, $p < 0.001$) and later sequence items ($\beta = 27.89$, $t = 5.71$, $p < 0.001$). As in the unigram version, there was a similar marginally significant effect of distance, with longer distances associated with slower RTs ($\beta = 13.51$, $t = 1.91$, $p < 0.06$). In contrast to the unigram version, here there was a significant effect of how many items were unobserved ($\beta = 15.83$, $t = 2.82$, $p < 0.005$), with more unseen items yielding higher RTs.
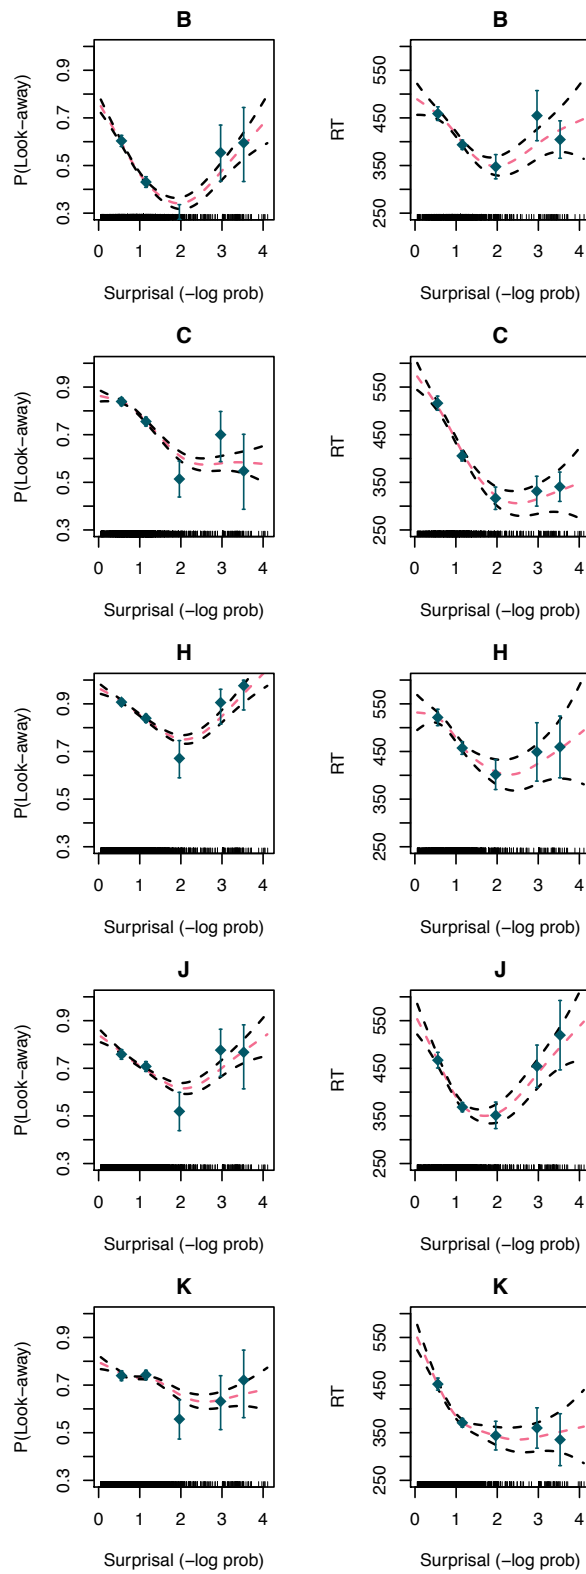
Fig. V-10. By-Subject Plots for the Unigram Model. For both look-away and RT, most monkey subjects exhibit behavior that shares a clear U-shaped relationship with surprisal. Interestingly, the preferred surprisal value across both behavioral measures appears to be approximately equivalent within subjects.
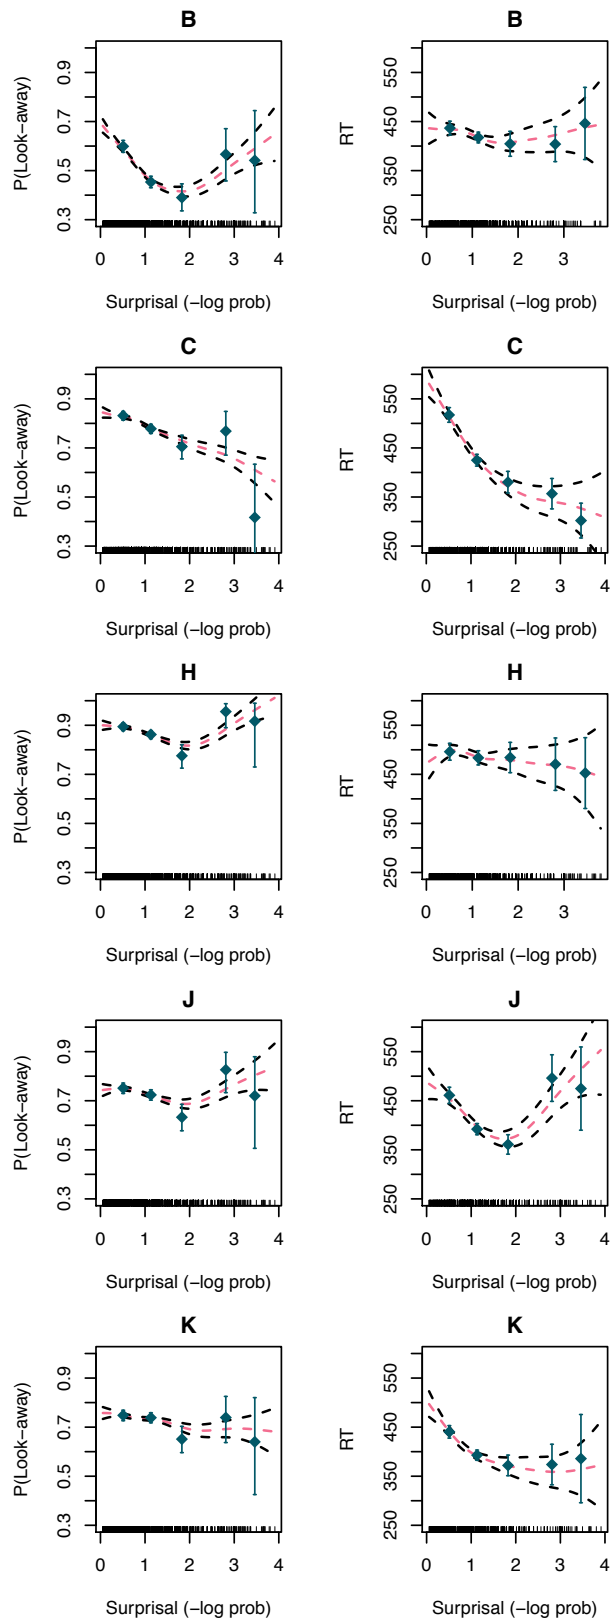
Fig. V-11. By-Subject Plots for the Transitional Model. The by-subject analyses for the transitional version of the model are less clear, with generally flatter—or at least shallower—trends for individual subjects. These results are consistent with the theory that monkeys' behavior relies more heavily on unigram statistics than transitional ones.

Conclusions.

The model's significant relationship with the behavioral measures is strong evidence that monkeys are able to track the statistics of the displays and that these statistics influence their attention. Our analysis has sought to discover the relationship between estimations of statistical probability and attentional behavior. This has revealed a range of behavior across measures and types of analyses. A U-shaped relationship was observed primarily in the unigram analysis of reaction times, and other measures showed significant linear trends in predicted directions.

The robustness of the unigram statistics in predicting each of the behavioral measures possibly suggests an important difference between monkeys and human learners. In general, unigram statistics were a more robust predictor of behavior than transitional statistics. This contrasts with infant attentional behavior. In tests of attention to sequential visual (Kidd, Piantadosi, & Aslin, 2012) and auditory (Kidd, Piantadosi, & Aslin, under review) stimuli transitional models out-performed unigram models. This difference may suggest that infants possess a sensitivity to transitional probabilities that monkeys lack. Differences in how each species processes sequential stimuli are plausible given that only human infants eventually acquire language.

Acknowledgements.

References.

Akaike, H. (1974). A New Look at the Statistical Model Identification. *IEEE Transactions on Automatic Control, 19*(6), 716–723.

Altmann, G.T.M., & Kamide, Y. (2007). The Real-Time Mediation of Visual Attention by Language and World Knowledge: Linking anticipatory (and other) eye movements to linguistic processing. *Journal of Memory and Language, 57*, 502-518.

Aslin, R.N. (2007). What's in a Look? *Developmental Science, 10*, 48–53.

Berlyne, D.E. (1954). A Theory of Human Curiosity. *British Journal of Psychology, 45*, 180-191.

Berlyne, D.E. (1960). Conflict, Arousal, and Curiosity, Vol. 331. New York: McGraw-Hill.

Brainard, D.H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*, 433-436.

Brysbaert, M., Buchmeier, M., Conrad, M., Jacobs, A.M., Bölte, J., & Böhl, A. (2011). The Word Frequency Effect: A review of recent developments and implications for the choice of frequency estimates in German. Experimental Psychology, 58(5), 412-24.

Conway, C.M., & Christiansen, M.H. (2009). Seeing and Hearing in Space and Time: Effects of modality and presentation rate on implicit statistical learning. *European Journal of Cognitive Psychology, 21*, 561-580.

Cornelissen, F.W., Peters. E., & Palmer, J. (2002). The Eyelink Toolbox: Eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments & Computers, 34*, 613-617.

Dember, W.N., & Earl, R.W. (1957). Analysis of Exploratory, Manipulatory, and Curiosity Behaviors. *Psychological Review, 64*, 91–96.

Fantz, R.L. (1964). Visual Experience in Infants: Decreased attention to familiar patterns relative to novel ones. *Science, 146*, 668.

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*, 15822.

Friedlander, B.Z. (1970). Receptive Language Development in Infancy: Issues and Problems. *Merrill-Palmer Quarterly of Behavior and Development, 16(*1), 7-51.

Genzel D, & Charniak E (2002). Entropy Rate Constancy in Text. *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Stroudsburg, PA: Association for Computational Linguistics. pp. 199–206.

Gerken, L., Balcomb, F.K., & Minton, J.L. (2011). Infants Avoid 'Labouring in Vain' by Attending More to Learnable than Unlearnable Linguistic Patterns. *Developmental Science, 14*(5), 972-9.

Grice, G.R., Nullmeyer, R., Spiker, V.A. (1982). Human Reaction Time: Toward a General Theory. *Journal of Experimental Psychology (General)*, *111*, 135–53.

Hastie T., & Tibshirani R. (1990). Generalized Additive Models. Boca Raton: Chapman and Hall/CRC.

Hosmer, D.W., Lemeshow, S., & May, S. (2008). Applied Survival Analysis: Regression Modeling of Time-to-Event Data, Second Edition. Hoboken, NJ: John Wiley and Sons.

Jepma, M., Verdonschot, R.G., van Steenbergen, H., Rombouts, S.A.R.B., & Nieuwenhuis, S. (2012). Neural Mechanisms Underlying the Induction and Relief of Perceptual Curiosity. *Frontiers in Behavioral Neuroscience*, *6*, 5.

Kidd, C., Piantadosi, S.T., & Aslin, R.N. (2010). The Goldilocks Effect: Infants' preference for visual stimuli that are neither too predictable nor too surprising. *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society* (pp. 2476-2481).

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*(5), e36399.

Kidd, C., Piantadosi, S.P., & Aslin, R.N. (Under review). The Goldilocks Effect in Infant Auditory Cognition.

Kinney, D.K. & Kagan, J. (1976). Infant Attention to Auditory Discrepancy. *Child Development, 47*(1), 155–164.

Klein J., & Moeschberger, M. (2003). Survival Analysis: Techniques for Censored and Truncated Data, Second Edition. New York: Springer-Verlag.

Kleiner, M., Brainard, D.H., & Pelli, D.G. (2007). What's New in Psychtoolbox-3? *Perception 36 ECVP Abstract Supplement.*

Litman, J.A. (2005). Curiosity and the Pleasures of Learning: Wanting and liking new information. *Cognition and Emotion*, *19*(6), 793-814.

Marr, D. (1982). *Vision.* San Francisco: W.H. Freeman.

Maye, J., Werker, J.F., & Gerken, L.A. (2002). Infant Sensitivity to Distributional Information Can Affect Phonetic Discrimination. *Cognition, 82*, B101–B111.

Min Jeong, K., Ming, H., Krajbich, I.M., Loewenstein, G., McClure, S.M., Wang, J., & Camerer, C.F. (2009). The Wick in the Candle of Learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, *20*(8), 963-973.

Ofer, G., & Durban, J. (1999). Curiosity: Reflections on its nature and functions. *American Journal of Psychotherapy*, *53*(1), 35-51.

Pelli, D.G. (1997). The VideoToolbox Software for Visual Psychophysics: Transforming numbers into movies. *Spatial Vision, 10*, 437-442.

Piantadosi, S.P., Kidd, C., & Aslin, R.N. (In press). Rich Analysis and Rational models: Inferring individual behavior from infant looking data. *Developmental Science.*

Piantadosi, S.T., Tily, H., & Gibson, E. (2011). Word Lengths Are Optimized for Efficient Communication. *Proceedings of the National Academy of Sciences of the United States of America 108*(9), 3526–3529.

Shannon, C.E. (1948). A Mathematical Theory of Communication. *Bell Systems Technical Journal*, *27*, 379-423.

Smith, N.J., & Levy, R. (2013). The Effect of Word Predictability on Reading Time is Logarithmic. *Cognition.*

Sternberg, S. (1969). The Discovery of Processing Stages: Extensions of Donders' method. In W. G. Koster (Ed.), *Attention and performance II. Acta Psychologica*, *30*, 276-315.

Teichner, W.H. & Krebs, M.J. (1974). Laws of visual choice reaction time. *Psychological Review*, *81*, 75-98.

VI. Summary.

Rational Species- and Modality-General Principles of Learning.

This work investigated both implicit and overt measures of the choice behavior of both young children and monkeys in order to understand the decision-making mechanisms that guide the acquisition of knowledge. This work, which encompassed behavioral experimentation with young children and non-human-primates across multiple domains—including visual attention and overt choice—aimed to better understand the efficacy and limitations of rational cognitive theories. In this thesis, I presented empirical evidence that suggests that naïve learners rely on rational utility maximization both to build complex models of the world starting from very little knowledge and, more generally, to guide their decisions and behavior.

Appendices

Chapter II (Rational Snacking) Appendices

Appendix II-1: Additional Scripted Dialogue between Experimenter and Child

Onset of experiment: *"So, today we have a very exciting art project planned for you! Upstairs, we have everything we'll need for you to make your own cup like this one! And you'll be able to take it home with you! Does that sound like something you'd like to do?"*

Art Material Choice (Choice 1): *"To decorate your cup, you have a choice of what art supplies to use. You could use these [crayons] right now. Or—if you can wait for me to go get them from another room—you can use our big set of art supplies instead. The big set has markers, pens, colored pencils—a lot of cool stuff. How does that sound? [Response.] Okay, I'm going to go get the big set of art supplies from the other room. You should stay right here in that chair. Can you do that? [Response.] I'll leave these [crayons] right here, and if you haven't used them when I come back, you can use our big set of art supplies instead!"*

Sticker Choice (Choice 2): *"Would you like to add a sticker to your picture? [Response.] For stickers, you have a choice. You can use this [sticker] right now. Or—if you can wait for me to go get them from the other room—you can have a bunch of stickers to use instead. How does that sound? [Response.] Okay, I'm going to go get more stickers from the other room. You should stay right here in that chair. Can you do that? [Response.] I'll leave this [sticker] here and if you haven't used it when I come back, you can have a bunch of stickers to use instead!"*

Appendix II-2: Detailed Subject Data

EXPERIMENTAL CONDITION	SUBJECT ID	GENDER	AGE IN MONTHS	WAIT-TIME (S)			WAITED 15
				Coder 1	*Coder 2*	*Diff*	
unreliable	1	m	42.7	17	17	*0*	N
unreliable	3	f	48.2	**21**	19	*2*	N
unreliable	5	m	48.2	7	7	*0*	N
unreliable	7	m	48.7	15	15	*0*	N
Unreliable	9	f	48.7	10	10	*0*	N
Unreliable	11	f	49.2	31	31	*0*	N
Unreliable	13	m	51.9	496	**498**	*-2*	N
unreliable	15	f	52.6	900	900	*0*	**Y**
unreliable	17	m	53.7	457	457	*0*	N
unreliable	19	m	58.3	72	**73**	*-1*	N
unreliable	21	m	62.8	18	18	*0*	N
unreliable	23	m	63.6	150	150	*0*	N
unreliable	25	f	67.3	195	195	*0*	N
unreliable	27	m	67.4	149	**150**	*-1*	N
Unreliable Group Means		***9m, 5f***	***54.52***	***181.57***			***7.14%***
reliable	2	m	43.4	900	900	*0*	**Y**
reliable	4	m	43.8	785	785	*0*	N
reliable	6	f	44.1	**431**	430	*1*	N
reliable	8	f	48.3	900	900	*0*	**Y**
reliable	10	m	48.3	59	59	*0*	N
reliable	12	f	48.6	144	**145**	*-1*	N
reliable	14	m	53.8	900	900	*0*	**Y**
reliable	16	f	54	900	900	*0*	**Y**
reliable	18	m	54.9	900	900	*0*	**Y**
reliable	20	m	57.4	900	900	*0*	**Y**
reliable	22	m	59.1	594	594	*0*	N
reliable	24	f	65.1	900	900	*0*	**Y**
reliable	26	m	68.8	900	900	*0*	**y**
reliable	28	m	70.1	900	900	*0*	**y**
Reliable Group Means		***9m, 5f***	***54.26***	***722.43***			***64.29%***

Appendix II-2: Raw data and group means. The 28 participants were randomly assigned to one of two conditions (*unreliable* or *reliable*). The two groups were matched for age and gender. Wait-time until first taste (i.e., lick or bite) were judged by two naïve coders watching video recordings that were blinded for condition. The coders' timing judgments were checked against one another to ensure validity, and where timing judgments differed slightly, the later judgment was used (and appears in bold above). The judgments of the two coders were found to differ by at most by 2 seconds. Children's waiting behavior was also coded in terms of a binary outcome measure corresponding to whether or not they waited the entire 15 min. without tasting the marshmallow or not (as indicated in the "Waited 15" column above). The percentages in this column reflect the portion of the group that waited the full 15 min.: 7.14% in the *unreliable* condition and 64.29% in the *reliable* one.

Appendix II-3: Analysis of Mood Variables

We used three control variables to investigate the potential influence of mood on children's wait times: *contentedness*, *smiling*, and *fidgeting*. Each measurement was based on a portion of each child's video data—the first 30 sec. of the waiting period.

1. **CONTENTEDNESS:** Two naïve coders rated each child's apparent contentedness on a scale from 1 to 9, with 1 indicating *very sad* and 9 indicating *very happy*. We computed z-scores for each coder's judgments, and then a mean z-score for each child.

2. **SMILING:** Two naïve coders measured for how long each child smiled, in seconds. We computed the mean of these two judgments.

3. **FIDGETING:** A Python script automatically computed an estimate of each child's movement. The script computed the mean number of pixel changes frame-to-frame for each child, above a noise threshold (diff > 50). The threshold served to control for pixel changes caused by the noise inherent in digital frame-to-frame comparisons of this type (caused by, for example, small differences in compression and subtle lighting changes). Thus, the threshold enabled us to measure only changes caused by the body movements of each child.

Wilcoxan rank sum tests indicated that these variables did not significantly differ across conditions in our sample population. Independent samples t-tests ($\alpha_{2\text{-}tail} = 0.05$) also failed to detect a significant difference across conditions.

	GROUP MEANS		WILCOXAN RANK SUM TEST	INDEPENDENT SAMPLES T-TEST
	Unreliable *(N = 14)*	*Reliable* *(N = 14)*		
CONTENTEDNESS *Mean z-scores*	0.03 *(sd = 0.89)*	-0.03 (*sd = 0.89)*	*W* = 106.5, *p* > 0.71	*t* = 0.178, *df* = 26, *p* > 0.85
SMILING *Mean seconds*	3.16 *(sd = 3.68)*	4.45 *(sd = 6.53)*	*W* = 96.5, *p* > 0.96	*t* = 0.644, *df* = 26, *p* > 0.52
FIDGETING *mean interframe pixel change*	0.61 *(sd = 0.36)*	0.61 *(sd = 0.39)*	*W* = 97, *p* > 0.98	*t* = 0.000, *df* = 26, *p* = 1.00

Chapter V (Curious George) Appendices

Appendix V-1: Markov Dirichlet-Multinomial Model (Ideal Learner Model)

Intuitively, learners observe how many times each event occurs in the world, and then use these event counts to infer an underlying probability model of their observations. In the experiment, there are three possible events corresponding to which of three objects appears from behind its box.

An observer who sees only a single event happen would not likely infer that the single observed event is the only one possible (i.e., has probability of 1); instead, observers likely bring expectations to this learning task. In the MDM model used here, this prior expectation is parameterized by a single free parameter, α, which controls the strength of the learner's prior belief that the distribution of events is uniform. As α gets large, the model has strong prior beliefs that the distribution of events in the world is uniform; as α approaches zero, the model believes more strongly that the true distribution closely resembles that of the empirically observed event counts. In modeling, we chose a value of $\alpha = 1$, corresponding to a uniform prior expectation about the distribution of events (33-33-33). However, the qualitative results—in particular, the U-shaped relationship between surprisal and look-away probability—do not depend strongly on the choice of α.

Formally, suppose there are N events, $x_1, x_2, \ldots, x_N$ and the ith event has been observed c_i times. We are interested in estimating (or *scoring*) a multinomial distribution parameterized by $\theta = (\theta_1, \theta_2, \ldots, \theta_N)$ where θ_i is the true (unobserved) probability of event x_i. Under a Dirichlet-Multinomial model,

(1)

$$P(\theta | c_1, \ldots c_N, \alpha) = \frac{1}{B} \prod_{i=1}^{N} \theta_i^{\alpha + c_i - 1}$$

where B is a normalizing constant that depends on the c_i and α. That is, after observing each event type occur some number of times, the infant may form a representation, θ, of their guess at the true distribution of events. Every distribution can be scored according to Equation 1, allowing one to compute how strongly a learner should believe that any particular θ is the correct one. We predict that infants' likelihood of looking away at a current event will depend upon the surprisal of that current event, which is determined by both the previously observed events and the identity of the current event. We predict that events of either very low surprisal (highly predictable) or very high surprisal (highly unexpected) will be more likely to trigger a look-away than events with moderate surprisal.

When the ith event occurs, the main variable of interest here is its negative log probability according to the model. We compute this by integrating over the above posterior distribution on θ. This corresponds to a measure of the information conveyed by observing event i according to an ideal Bayesian learner who had seen all previous events. We predicted that infants would be more likely to look away during events that contained either too little or too much information, giving a U-shaped (quadratic) relationship between this negative log probability measure and the actual observed look-away probability.

Appendix V-2: Visual sequences for eye-tracking experiment

Table V-2A. Short sequences (30 items)

Sequence ID	\ Sequence item 1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
1	2	1	2	2	2	1	3	3	3	1	3	1	1	2	2	1	1	1	2	3	2	2	1	2	1	1	1	3	2	2
2	3	3	3	3	1	3	3	1	3	3	3	2	3	3	3	1	1	1	3	3	1	3	1	2	1	1	1	2	2	2
3	1	1	1	1	1	1	1	3	3	3	2	2	1	2	3	3	3	2	1	2	2	3	1	1	2	3	1	1	1	3
4	3	2	1	2	2	1	2	1	1	3	2	2	1	1	2	2	2	1	1	1	2	1	1	1	2	2	1	2	2	2
5	3	1	3	3	3	3	3	3	3	2	2	3	1	3	3	2	2	2	1	3	2	3	1	3	3	1	2	3	2	3
6	1	1	1	1	1	1	1	1	2	2	2	1	2	3	3	1	2	2	1	2	3	2	3	3	2	2	1	3	3	1
7	2	1	1	2	1	1	3	3	1	2	1	1	2	2	1	2	2	2	2	2	2	1	1	2	2	2	2	2	2	1
8	1	1	1	1	1	3	1	1	2	3	3	1	1	1	3	1	1	1	2	3	2	1	3	3	2	2	3	2	2	2
9	1	3	1	1	3	1	3	2	1	3	3	1	2	1	1	1	1	1	3	3	2	1	1	1	1	1	2	1	1	2
10	2	1	1	3	2	3	3	2	1	3	2	3	3	1	1	1	1	3	1	3	1	1	3	3	3	1	3	3	1	3
11	1	1	3	2	1	2	2	3	3	2	1	1	2	2	1	2	1	1	3	2	1	1	1	2	2	1	2	1	2	1
12	1	2	1	2	3	1	2	2	3	1	2	1	2	3	2	2	2	2	2	1	1	3	3	3	2	3	3	1	3	2
13	3	3	2	3	1	2	2	1	1	2	3	3	3	3	2	1	1	2	3	3	3	2	1	1	1	2	1	2	1	2
14	3	2	2	3	3	2	2	2	3	3	2	3	3	2	2	3	3	2	3	2	2	2	3	2	3	3	1	1	1	3
15	2	3	1	2	1	1	2	1	1	2	2	1	3	2	2	1	2	1	2	1	2	3	2	1	1	3	1	1	3	3
16	2	2	2	2	2	1	2	2	1	2	1	1	2	1	1	1	1	2	3	2	2	2	2	3	3	3	1	3	3	1
17	1	2	3	1	2	3	3	1	2	3	2	3	1	2	2	2	3	1	3	3	1	2	2	3	1	3	1	1	3	2
18	2	1	2	1	1	2	1	1	1	1	3	2	2	2	2	2	3	2	1	2	2	2	3	2	3	1	2	3	3	2
19	2	3	3	1	2	2	2	2	3	3	2	3	3	3	2	2	2	2	2	2	2	2	3	2	3	3	2	2	2	3
20	1	3	3	1	3	3	2	2	1	2	2	3	2	2	3	1	1	2	2	3	1	1	2	2	3	2	1	2	2	2
21	3	1	1	1	1	3	1	3	3	3	1	3	3	1	3	3	1	3	1	1	3	1	2	2	3	2	1	1	2	3
22	2	1	2	2	1	3	1	2	2	2	2	1	3	2	2	2	2	1	3	2	2	2	1	3	2	2	1	2	1	3
23	3	2	1	3	2	1	1	3	1	2	3	2	1	1	2	1	2	1	1	1	2	1	1	1	2	1	1	2	1	2
24	2	1	1	1	3	3	1	3	1	1	3	3	3	1	3	3	1	3	2	1	1	2	2	3	1	2	1	1	2	2
25	3	2	1	3	1	3	2	2	1	3	1	3	1	2	3	1	1	3	3	2	3	2	1	3	3	3	3	2	1	2
26	3	3	3	3	3	3	3	2	2	1	2	1	2	1	1	1	2	3	1	1	3	1	3	1	2	1	3	3	2	2
27	2	2	1	2	3	1	1	3	2	2	3	3	2	2	3	1	1	3	1	3	1	1	3	1	1	1	1	1	3	2
28	2	3	1	3	3	1	1	1	2	2	3	3	1	3	3	1	2	2	3	2	2	3	3	2	1	2	1	3	3	2
29	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
30	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2
31	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3
32	1	2	1	2	1	2	1	2	1	2	1	2	3	1	2	1	2	3	3	3	3	3	3	3	3	3	3	3	3	3

Table V-2*B*. Long sequences (60 items)

Sequence item (columns 1–60)

Sequence ID (rows 33–80)

Appendix V-3: Predictive-Look Regression Tables

Table V-3A. Unigram Raw Predictive-Look Regression.

```
Generalized linear mixed model fit by the Laplace approximation
Formula: already_there ~ std_surprisal + std_sq_surprisal + (1 +
std_surprisal +      std_sq_surprisal | subj)
Data: d[d$repeated == 0, ]

AIC   BIC   logLik deviance
14824 14890  -7403    14806

Random effects:
Groups Name            Variance  Std.Dev. Corr
subj   (Intercept)      0.0270622 0.164506
std_surprisal           0.0121747 0.110339 -0.016
std_sq_surprisal        0.0010764 0.032808 -0.316 -0.920
Number of obs: 10920, groups: subj, 5

Fixed effects:
                Estimate Std. Error z value Pr(>|z|)
(Intercept)      0.08854    0.07641   1.159    0.247
std_surprisal   -0.50861    0.05988  -8.494  < 2e-16 ***
std_sq_surprisal 0.08801    0.01770   4.971 6.65e-07 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
           (Intr) std_sr
std_surprsl -0.026
std_sq_srpr -0.267 -0.893
```

Table V-3*B*. Unigram Controlled Predictive-Look Regression.

```
Generalized linear mixed model fit by the Laplace approximation
Formula: already_there ~ firstappear + std_trial + std_seq_item +
std_dist +      std_unseen + std_surprisal + std_sq_surprisal + (1 +
std_surprisal +      std_sq_surprisal | subj)
Data: d[d$repeated == 0, ]

AIC   BIC   logLik deviance
13396 13497  -6684    13368

Random effects:
Groups Name              Variance    Std.Dev. Corr
subj   (Intercept)       0.03573834 0.189046
std_surprisal            0.01354808 0.116396 -0.160
std_sq_surprisal         0.00092291 0.030379 -0.239 -0.899
Number of obs: 10503, groups: subj, 5

Fixed effects:
                 Estimate Std. Error z value Pr(>|z|)
(Intercept)      -0.219164   0.105796  -2.072 0.038305 *
firstappear1      0.410469   0.060565   6.777 1.22e-11 ***
std_trial         0.253667   0.021216  11.956  < 2e-16 ***
std_seq_item      0.546424   0.026267  20.803  < 2e-16 ***
std_dist          0.089913   0.031019   2.899 0.003748 **
std_unseen        0.232972   0.039021   5.970 2.37e-09 ***
std_surprisal    -0.251501   0.064651  -3.890 0.000100 ***
std_sq_surprisal  0.006215   0.017611   0.353 0.724151
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
            (Intr) frstp1 std_tr std_sq_t std_ds std_ns std_sr
firstapper1 -0.510
std_trial    0.007 -0.007
std_seq_itm  0.057 -0.020 -0.065
std_dist    -0.250  0.039 -0.009  0.027
std_unseen  -0.223  0.541 -0.053  0.444    0.061
std_surprsl -0.185  0.180 -0.017  0.118   -0.020  0.195
std_sq_srpr -0.182 -0.005  0.016 -0.216    0.024 -0.155 -0.861
```

Table V-3C. Transitional Raw Predictive-Look Regression.

```
Generalized linear mixed model fit by the Laplace approximation
Formula: already_there ~ std_bi_surprisal + sq_std_bi_surprisal + (1 +
std_bi_surprisal + sq_std_bi_surprisal | subj)
Data: d[d$repeated == 0, ]

AIC   BIC   logLik deviance
14425 14491  -7204    14407

Random effects:
Groups Name           Variance  Std.Dev. Corr
subj   (Intercept)    0.0288117 0.169740
std_bi_surprisal      0.0100174 0.100087  0.243
sq_std_bi_surprisal 0.0011827 0.034391 -0.661 -0.838
Number of obs: 10503, groups: subj, 5

Fixed effects:
                 Estimate Std. Error z value Pr(>|z|)
(Intercept)       0.14090    0.07883  1.7874  0.07388 .
std_bi_surprisal   -0.15302    0.05378 -2.8453  0.00444 **
sq_std_bi_surprisal -0.00526    0.01869 -0.2814  0.77837
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
          (Intr) std_b_
std_b_srprs  0.198
sq_std_b_sr -0.563 -0.818
```

Table V-3*D*. Transitional Controlled Predictive-Look Regression.

```
Generalized linear mixed model fit by the Laplace approximation
Formula: already_there ~ firstappear + std_trial + std_seq_item +
std_dist +       std_unseen + std_bi_surprisal + sq_std_bi_surprisal +
(1 +       std_bi_surprisal + sq_std_bi_surprisal | subj)
Data: d[d$repeated == 0, ]

AIC   BIC   logLik deviance
13407 13508  -6689    13379

Random effects:
Groups Name              Variance     Std.Dev.     Corr
subj   (Intercept)       0.0374518    0.193525
std_bi_surprisal         0.0130942    0.114430     0.198
sq_std_bi_surprisal      0.0011663    0.034152     -0.647       -0.861
Number of obs: 10503, groups: subj, 5

Fixed effects:
                  Estimate Std. Error z value Pr(>|z|)
(Intercept)        -0.32041    0.10494  -3.053  0.00226 **
firstappear1        0.54786    0.05616   9.755  < 2e-16 ***
std_trial           0.25072    0.02125  11.800  < 2e-16 ***
std_seq_item        0.58481    0.02565  22.798  < 2e-16 ***
std_dist            0.08927    0.03103   2.877  0.00402 **
std_unseen          0.29740    0.03802   7.822 5.19e-15 ***
std_bi_surprisal   -0.07680    0.06061  -1.267  0.20510
sq_std_bi_surprisal -0.04949    0.01908  -2.593  0.00950 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
            (Intr) frstp1 std_tr std_s_ std_ds std_ns std_b_
firstapper1 -0.463
std_trial    0.006 -0.013
std_seq_itm  0.059 -0.021 -0.070
std_dist    -0.252  0.038 -0.009  0.033
std_unseen  -0.197  0.539 -0.064  0.417  0.063
std_b_srprs  0.091  0.137 -0.037  0.032 -0.023  0.137
sq_std_b_sr -0.472 -0.011  0.039 -0.140  0.025 -0.105 -0.822
```

Appendix V-4: Look-Away Regression Tables

Table V-4A. Unigram Raw Look-Away Regression.

```
Generalized linear mixed model fit by the Laplace approximation
Formula: wouldbelookaway ~ std_surprisal + std_sq_surprisal + (1 +
std_surprisal +      std_sq_surprisal | subj)
   Data: d

  AIC   BIC logLik deviance
 21813 21884 -10897    21795

Random effects:
 Groups Name             Variance  Std.Dev. Corr
 subj   (Intercept)      0.3833228 0.619131
        std_surprisal    0.0287164 0.169459 -0.059
        std_sq_surprisal 0.0041159 0.064155 -0.047 -0.901
Number of obs: 19890, groups: subj, 5

Fixed effects:
                Estimate Std. Error z value Pr(>|z|)
(Intercept)      0.94801    0.27750   3.416 0.000635 ***
std_surprisal   -0.43882    0.07966  -5.508 3.62e-08 ***
std_sq_surprisal 0.10336    0.03001   3.444 0.000573 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
           (Intr) std_sr
std_surprsl -0.052
std_sq_srpr -0.052 -0.884
```

Table V-*4B.* Unigram Controlled Look-Away Regression.

```
Generalized linear mixed model fit by the Laplace approximation
Formula: wouldbelookaway ~ repeated + firstappear + std_trial +
std_seq_item + std_dist + std_unseen + std_surprisal + std_sq_surprisal
+ (1 + std_surprisal + std_sq_surprisal | subj)
   Data: d

   AIC   BIC logLik deviance
 18967 19085  -9468    18937

Random effects:
 Groups Name             Variance Std.Dev.  Corr
 subj   (Intercept)       0.503506 0.709582
 std_surprisal            0.031810 0.178353  0.021
 std_sq_surprisal         0.003412  0.058412 -0.046 -0.868
Number of obs: 19473, groups: subj, 5

Fixed effects:
                 Estimate Std. Error z value Pr(>|z|)
(Intercept)       0.79457    0.32187   2.469   0.0136 *
repeated1        -0.08929    0.03715  -2.403   0.0162 *
firstappear1      0.47413    0.05364   8.839   <2e-16 ***
std_trial         0.54284    0.01935  28.047   <2e-16 ***
std_seq_item      0.70397    0.02487  28.308   <2e-16 ***
std_dist          0.05262    0.03616   1.455   0.1456
std_unseen        0.28996    0.02919   9.933   <2e-16 ***
std_surprisal    -0.16137    0.08605  -1.875   0.0607 .
std_sq_surprisal  0.01930    0.02814   0.686   0.4930
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
          (Intr) reptd1 frstp1 std_tr std_sq_t std_ds std_ns std_sr
repeated1   -0.019
firstapper1 -0.151  0.033
std_trial    0.010  0.004  0.018
std_seq_itm  0.029 -0.032  0.000  0.002
std_dist     0.005 -0.840  0.021 -0.012  0.023
std_unseen  -0.045 -0.013  0.440 -0.029  0.512   0.031
std_surprsl  0.008 -0.061  0.134 -0.013  0.093  -0.006  0.175
std_sq_srpr -0.057  0.035  0.017  0.013 -0.132   0.007 -0.130 -0.842
```

Table V-4C. Transitional Raw Look-Away Regression.

```
Generalized linear mixed model fit by the Laplace approximation
Formula: wouldbelookaway ~ std_bi_surprisal + sq_std_bi_surprisal + (1
+      std_bi_surprisal + sq_std_bi_surprisal | subj)
   Data: d

   AIC   BIC logLik deviance
 21337 21408 -10660    21319

Random effects:
 Groups Name                 Variance  Std.Dev. Corr
 subj   (Intercept)          0.3898874 0.624410
 std_bi_surprisal            0.0093706 0.096802  0.451
 sq_std_bi_surprisal         0.0017557 0.041901 -0.454 -0.798
Number of obs: 19473, groups: subj, 5

Fixed effects:
                    Estimate Std. Error z value Pr(>|z|)
(Intercept)          1.01748    0.27992   3.635 0.000278 ***
std_bi_surprisal    -0.20923    0.04850  -4.314  1.6e-05 ***
sq_std_bi_surprisal  0.05175    0.02096   2.469 0.013546 *
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
           (Intr) std_b_
std_b_srprs  0.409
sq_std_b_sr -0.419 -0.762
```

Table V-4*D.* Transitional Controlled Look-Away Regression.

```
Generalized linear mixed model fit by the Laplace approximation
Formula: wouldbelookaway ~ repeated + firstappear + std_trial +
std_seq_item +       std_dist + std_unseen + std_bi_surprisal +
sq_std_bi_surprisal +      (1 + std_bi_surprisal + sq_std_bi_surprisal
| subj)
   Data: d

   AIC   BIC logLik deviance
 19011 19129  -9490    18981

Random effects:
 Groups Name                   Variance  Std.Dev. Corr
 subj   (Intercept)            0.5112778 0.715037
 std_bi_surprisal              0.0169621 0.130239  0.387
 sq_std_bi_surprisal           0.0012296 0.035065 -0.505 -0.860
Number of obs: 19473, groups: subj, 5

Fixed effects:
                    Estimate Std. Error z value Pr(>|z|)
(Intercept)          0.75129    0.32360   2.322  0.02025 *
repeated1           -0.11315    0.03680  -3.075  0.00211 **
firstappear1         0.57708    0.04833  11.941  < 2e-16 ***
std_trial            0.53995    0.01936  27.891  < 2e-16 ***
std_seq_item         0.73896    0.02466  29.965  < 2e-16 ***
std_dist             0.04742    0.03613   1.312  0.18941
std_unseen           0.35537    0.02717  13.082  < 2e-16 ***
std_bi_surprisal     0.00413    0.06385   0.065  0.94843
sq_std_bi_surprisal -0.03134    0.01900  -1.649  0.09905 .
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
          (Intr) reptd1 frstp1 std_tr std_s_ std_ds std_ns std_b_
repeated1   -0.022
firstapper1 -0.135  0.079
std_trial    0.010  0.003  0.015
std_seq_itm  0.029 -0.005 -0.001 -0.004
std_dist     0.005 -0.847  0.020 -0.009  0.021
std_unseen  -0.037  0.040  0.410 -0.037  0.496  0.033
std_b_srprs  0.345 -0.040  0.100 -0.028  0.088 -0.010  0.138
sq_std_b_sr -0.437  0.014  0.037  0.033 -0.195  0.015 -0.141 -0.788
```

Appendix V-5: Reaction-Time Regression Tables

Table V-5A. Unigram Raw Reaction-Time Regression.

```
Linear mixed model fit by REML
Formula: time_till_target ~ std_surprisal + std_sq_surprisal + (1 +
std_surprisal +      std_sq_surprisal | subj)
Data: d[d$time_till_target > 0, ]

AIC         BIC          logLik       deviance     REMLdev
117565      117635       -58773       117561       117545

Random effects:
Groups    Name           Variance     Std.Dev.   Corr
subj      (Intercept)     1074.731    32.7831
std_surprisal             102.695     10.1339    0.441
std_sq_surprisal           10.282      3.2066   -0.976   -0.234
Residual                 108592.987  329.5345
Number of obs: 8144, groups: subj, 5

Fixed effects:
                Estimate Std. Error t value   p value
(Intercept)      422.152     15.182  27.807    < 2e-16     ***
std_surprisal    -64.055      6.823  -9.387    < 2e-16     ***
std_sq_surprisal  14.215      2.188   6.498    < 8e-11     ***

Correlation of Fixed Effects:
           (Intr) std_sr
std_surprsl  0.309
std_sq_srpr -0.682 -0.527
```

Table V-5B. Unigram Controlled Reaction-Time Regression.

```
Linear mixed model fit by REML
Formula: time_till_target ~ repeated + firstappear + std_trial +
std_seq_item +       std_dist + std_unseen + std_surprisal +
std_sq_surprisal +       (1 + std_surprisal + std_sq_surprisal | subj)
Data: d[d$time_till_target > 0, ]

AIC         BIC         logLik      deviance    REMLdev
112088      112200      -56028      112103      112056

Random effects:
Groups    Name          Variance    Std.Dev.    Corr
subj      (Intercept)   1.0249e+03  32.0146
std_surprisal           1.2532e+02  11.1948      0.203
std_sq_surprisal        6.2208e+00   2.4942     -1.000       -0.183
Residual                1.0906e+05 330.2406
Number of obs: 7764, groups: subj, 5

Fixed effects:
                Estimate Std. Error t value    p value
(Intercept)      423.636     17.397  24.351    < 2e-16     ***
repeated1        -46.556      7.499  -6.208    < 6e-10     ***
firstappear1      35.995      9.781   3.680    0.00032     ***
std_trial         15.507      3.857   4.020    < 7e-05     ***
std_seq_item      24.339      4.969   4.898    < 1e-06     ***
std_dist          13.764      7.077   1.945    0.05180     .
std_unseen         7.963      6.020   1.323    0.18587
std_surprisal    -23.889      8.415  -2.839    0.00453     **
std_sq_surprisal   5.234      2.266   2.310    0.02091     *

Correlation of Fixed Effects:
          (Intr)   reptd1 frstp1 std_tr std_sq_t std_ds std_ns std_sr
repeated1   -0.090
firstapper1 -0.499  0.030
std_trial    0.025  0.015 -0.009
std_seq_itm  0.103 -0.037 -0.005 -0.126
std_dist     0.015 -0.802  0.032 -0.015  0.037
std_unseen  -0.190 -0.001  0.506 -0.082  0.498    0.047
std_surprsl  0.013 -0.161  0.318 -0.015  0.213   -0.010  0.379
std_sq_srpr -0.499  0.119 -0.045  0.015 -0.354    0.009 -0.336 -0.585
```

Table V-5C. Transitional Raw Reaction-Time Regression.

```
Linear mixed model fit by REML
Formula: time_till_target ~ std_bi_surprisal + sq_std_bi_surprisal + (1
+      std_bi_surprisal + sq_std_bi_surprisal | subj)
Data: d[d$time_till_target > 0, ]

AIC          BIC            logLik       deviance     REMLdev
112321       112391         -56151       112318       112301

Random effects:
Groups    Name                   Variance   Std.Dev. Corr
subj     (Intercept)             1211.32    34.804
std_bi_surprisal                  557.69    23.615   0.348
sq_std_bi_surprisal                25.07     5.007  -0.829 -0.813
Residual                       111980.36   334.635
Number of obs: 7764, groups: subj, 5

Fixed effects:
                   Estimate Std. Error t value   p value
(Intercept)         433.250     16.130  26.859 < 2e-16     ***
std_bi_surprisal    -33.254     11.606  -2.865 0.00417     **
sq_std_bi_surprisal   7.318      2.922   2.504 0.01243     *

Correlation of Fixed Effects:
           (Intr) std_b_
std_b_srprs  0.326
sq_std_b_sr -0.682 -0.748
```

Table V-5*D*. Transitional Controlled Reaction-Time Regression.

```
Linear mixed model fit by REML
Formula: time_till_target ~ repeated + firstappear + std_trial +
std_seq_item +      std_dist + std_unseen + std_bi_surprisal +
sq_std_bi_surprisal +      (1 + std_bi_surprisal + sq_std_bi_surprisal
| subj)
Data: d[d$time_till_target > 0, ]

AIC    BIC       logLik deviance REMLdev
112089 112200    -56028   112104  112057

Random effects:
Groups Name               Variance    Std.Dev.    Corr
subj(Intercept)           1065.575    32.6431
std_bi_surprisal           612.121    24.7411     0.279
sq_std_bi_surprisal         24.666     4.9665    -0.666     -0.902
Residual                109018.601   330.1797
Number of obs: 7764, groups: subj, 5

Fixed effects:
                  Estimate Std. Error t value   p value
(Intercept)          418.241     17.218  24.291 < 2e-16     ***
repeated1            -50.770      7.400  -6.861 < 8e-12     ***
firstappear1          47.112      8.745   5.387 < 8e-08     ***
std_trial             16.777      3.881   4.323 < 2e-05     ***
std_seq_item          27.894      4.889   5.705 < 2e-08     ***
std_dist              13.505      7.081   1.907 0.05655     .
std_unseen            15.834      5.611   2.822 0.00478     **
std_bi_surprisal      -4.935     12.325  -0.400 0.68916
sq_std_bi_surprisal    1.106      3.013   0.367 0.71362


Correlation of Fixed Effects:
          (Intr)  reptd1 frstp1 std_tr std_s_ std_ds std_ns std_b_
repeated1   -0.106
firstapper1 -0.452  0.092
std_trial    0.023  0.021 -0.021
std_seq_itm  0.104  0.002 -0.012 -0.145
std_dist     0.015 -0.811  0.033 -0.013  0.034
std_unseen  -0.162  0.067  0.471 -0.108  0.477  0.046
std_b_srprs  0.196 -0.049  0.116 -0.036  0.087 -0.015  0.140
sq_std_b_sr -0.509  0.027  0.000  0.049 -0.250  0.016 -0.179 -0.799
```