Probability-matching in 10-month-old infants

Sarah J. Davis (sdavis@bcs.rochester.edu)
Department of Brain & Cognitive Sciences, University of Rochester,
Rochester, NY 14607 USA

Elissa L. Newport (newport@bcs.rochester.edu)
Department of Brain & Cognitive Sciences, University of Rochester,
Rochester, NY 14607 USA

Richard N. Aslin (aslin@cvs.rochester.edu)
Department of Brain & Cognitive Sciences, University of Rochester,
Rochester, NY 14607 USA

Abstract

Evidence from the probability learning literature indicates that when presented with simple situations that require making predictions, adults tend to probability match whereas children are likely to show maximization (Stevenson & Weir, 1959; Weir, 1964). The reason for this developmental difference is not fully understood, but one possibility investigated here is that children have fewer resources available to differentiate among the probabilities of the competing alternatives. To investigate this hypothesis at its origin, we used an anticipatory eye movement paradigm to gather two-alternative choice responses from 10-month-old infants. In two experiments we presented infants with either an entirely predictable (100-0%) or a probabilistic (70-30%) series of visual events. Infants showed evidence of probability matching rather than maximizing. These results are discussed in the context of alternative explanations for maximizing and the utility of eye-tracking as a window on infants' probability learning.

Keywords: Probability learning; probability matching; eye-tracking; infants

Introduction

As we explore our world, we sample from our environment in order to make predictions about future events and to assess the likelihood of receiving rewards. For example, evidence from the statistical learning literature indicates that adults, infants, and animals can extract information about the distributional properties of visual and auditory stimuli in the absence of a task (Fiser & Aslin, 2001, 2002; Kirkham et al., 2002; Saffran et al., 1996a,b; Toro & Trabalón, 2005). Evidence from the causal learning literature indicates that young children are sensitive to event contingencies (Gopnik et al., 2004).

This ability to track and store information about probabilities allows learners to adjust their behavior to maximize their predictions and their receipt of rewards, even when there is not a perfect correlation between events and their outcomes. When faced with the task of predicting future events in an uncertain environment a learner has two strategies. One is to make predictions that directly match the exposure probabilities observed in the environment, a pattern known as *probability matching*. The other is to always choose the more common outcome, a pattern known as *maximization*.

In the context of reward prediction, an ideal learner should choose the action that maximizes the overall rate of reward. However, evidence from the probability learning literature indicates that adults tend to probability match rather than maximize, at least in simple situations (Gardner, 1957; Weir, 1964, 1972). In the classic probability learning experiment, Gardner (1957) presented participants with two light bulbs and asked them on each trial to predict which light would illuminate. After participants made a choice, one of the bulbs would turn on. One bulb turned on 70% of the time and the other bulb 30% of the time. If the participants were probability matching (i.e., picking the 70% light on 70% of the trials and the 30% light on 30% of the trials), then their overall accuracy would average 58% correct. If, on the other hand, learners chose the 70% light on every trial, their overall accuracy would be 70% correct. In this situation, maximizing on the more probable alternative is the better strategy because it leads to higher overall accuracy. Yet under most circumstances, adults typically probability match.

Whether adults show probability matching or maximizing on a given task can be influenced by a number of factors, including the contingency of the feedback (Weir, 1972) and the number of response alternatives. For example, when the number of alternatives increases, participants are more likely to show maximizing behavior in both visual tasks (Gardner, 1957; Weir 1964) and auditory language learning experiments (Hudson Kam & Newport, 2009).

The age of participants is also related to performance in these tasks, with the youngest children often exhibiting the highest rates of maximizing behavior (Austin & Newport, unpublished manuscript; Hudson Kam & Newport, 2009; Stevenson & Weir, 1959; Weir, 1964). When given access to the same input, why might children act differently than adults? It seems unlikely that they are better strategizers than adults. Rather this behavior could be based on their greater cognitive limitations, either in their representations of the world or their use of those representations.

When a learner comes into an environment where there are two possible outcomes, such as the two light bulb task, it

is plausible that expectations will be equally divided between the two events. With each new piece of data (i.e., which light bulb actually turns on during that trial), expectations will be updated. After sufficient exposure to the environmental events, a skilled learner should have formed expectations that match the frequency of the observed events (i.e., probability matching). One possible explanation for why children are more likely to maximize than adults is that their representation of events in the environment may be less nuanced. Rather than having a representation that matches the frequencies found in their input, they might have a representation that is more weighted towards the most commonly experienced outcome because they do not accurately perceive or store the lower probabilities. If the young learner is biased to focus attention on the more salient (i.e., common) event probability, then this could lead to the maximizing behavior observed in many tasks.

A second possibility for these behavioral age differences is that adults and children may have the same representations for events, but they may have differences in how they select a response based on these representations. Perhaps children require more evidence than adults before they are able to enact a relatively uncommon option. Although the child may have a representation of the less common event occurring 30% of the time, this relatively low rate of occurrence may not be enough to cause them to consider that choice as a viable alternative, or it may not be enough to support the selection of the weaker response tendency. When adults are engaged in a probability learning experiment, there is a clear goal (to pick the correct light) and an overt choice response is required. The benefit of testing infants is that we can explore their natural tendencies to sample from their environment without an explicit task. Observing this behavior may provide insight into the underlying causes for probability matching or maximizing.

A number of studies of infants' abilities to track probabilities and statistics have relied on looking time measures. For example, Xu and Garcia (2008) explored whether 8-month-old infants could use probabilistic information to make generalizations about likely events. After viewing a box containing many more red balls than white balls, the infant saw the experimenter pull 5 balls from the box. Infants showed a significantly longer looking time when the 5 balls at test consisted of four white and only one red ball (an unlikely outcome given the contents of the box) than when there were four red and one white ball (a likely outcome). This result indicated that infants are sensitive to the probabilities in their environment and that they can use this information to make predictions about future events.

This type of looking time measure can be used to indicate what events are surprising to an infant, and thus indicate their expectations. Duration of looking, however, is only able to reveal the infant's behavior (and allow inferences about expectations) after an event has occurred. In order to reveal whether infants probability-match or maximize, a paradigm is required that measures choice responses. Here we provide data from just such a two-choice paradigm using an anticipatory eye movement paradigm (McMurray & Aslin, 2004). A target object disappeared behind an occluder and reappeared at each of two locations with some probability (e.g., 70% left and 30% right). The dependent measure was the proportion of trials in which infants exhibited anticipatory eye-movements to each of these locations *before* the target reappeared from behind the occluder. Our results show that 10-month-old infants probability-match rather than maximize, and that they also modulate their expectations based on the ratio of event probabilities.

Experiment 1: Baseline

The purpose of this experiment was two-fold. First, it was designed to test the feasibility of using an anticipatory eye movement paradigm to track infant eye gaze during a predictive object occlusion task. Second, it provides a baseline measure of anticipatory responses when the outcome is entirely predictable, and thus can be compared to a probabilistic set of outcomes.

Methods

Participants. 12 parents from the Rochester community volunteered their infants. The parents were recruited through mailings, posted flyers, and web ads. The infants ranged in age from 10.1 to 11.1 months (M = 10.6 months) and had no reported hearing or vision deficits. An additional 8 infants were tested but excluded due to fussiness (7) and eye-tracker calibration difficulties (1). Participants received either $10 or a toy as compensation.

Materials. An image of a blue occluder in the shape of an inverted T was presented on a light grey background. The target object was a yellow smiley-face.

Apparatus. Eye-tracking was performed using a table-mounted Tobii 1750 eye-tracker with a 17-inch monitor. The stimuli were presented using the SMART-T program (Shukla, Wen, White & Aslin, in press) through Matlab running on a Mac Mini with an Inter Core 2 Duo processor.

Procedure. Each infant was seated on a parents' lap with the child's eyes approximately 23 inches from the Tobii monitor. Infants viewed a minimum of 20 trials using the McMurray and Aslin (2004) occlusion-based anticipatory eye movement paradigm. This design creates a two-alternative forced-choice (2AFC) procedure in which infants can learn to anticipate the reappearance of a target object after occlusion.

Each trial began with the appearance and looming (to 150% of its original size) of the target object (the smiley-face) below the occluder while a sound played ("Ooo"). During each trial the object moved upward at a rate of 150 pixels per second (3.9 deg/sec) behind the occluder, paused for one second, and then continued to move so as to

reappear in the open space at the top right or top left of the occluder (Figure 1). While following this path, a tonal melody was played as a background sound. Full occlusion of the target occurred for 2400 ms. Once fully visible from behind the occluder, the object loomed again and another sound was played ("Wow").

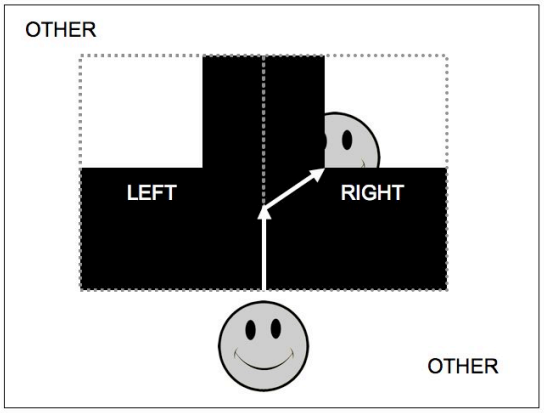

Figure 1: Target object with path on a right side trial, inverted T occluder and coding areas

On the first trial the occluder was semi-transparent, allowing the infant to see the movement of the target behind the occluder. Over the next 4 trials the occluder became less transparent until it was opaque and fully hid the target's movement on the 6th trial.

In this baseline condition all trials ended with the target reappearing on the same side. Half the infants had 100% reappearance on the left (100L) and half had 100% reappearance on the right (100R).

If infants are predicting the trajectory and reappearance of the target, and not just reacting to its reappearance, then the majority of their fixations during the period of target occlusion should be on the side of the screen where they expect the target to reappear.

Results

The first 10 trials of the experiment were treated as the "training" phase and were not included in the analyses. In order to observe the predictions of the infant, fixation location was observed during the period of target occlusion. To ensure that the infant was actively watching the display, only trials in which the infant fixated on the target during its initial loom (before it began to move to become occluded) were analyzed. In addition, only those infants who attended to a minimum of 10 trials (after the training period) were included in the final analyses. Infants completed an average of 32 trials (range: 27-47).

Two areas of interest (AOI) were defined for identifying anticipatory eye movements: the right and left sides of the occluder (see Figure 1). Looks outside these AOIs (labeled as "other" in Figure 1) were not considered to be indicative of a left/right side choice for predicting the reappearance of the target. Infants in the 100L and 100R conditions did not differ on their proportions of looking time to the more common side ($p>0.5$) and so the two groups were collapsed for all analyses.

The proportion of looks to each side was calculated at each time point during the occlusion period for each infant. Figure 2 shows the resulting mean time-course plot for infants in the baseline condition. These results suggest that infants were able to predict the reappearance of the target, spending the majority of the time while the object was occluded looking at the correct AOI.

Two metrics of correct anticipations were calculated for each infant. The first computed the total looking time to each AOI while the target was occluded and expressed the looking time to the correct AOI as a proportion. Infants in this baseline condition spent 75% of their looking time to the dominant AOI (100R or 100L). This performance was significantly greater than 50% ($p<0.001$).

The second metric calculated, for each infant on each trial, to which AOI they spent the majority of their looking time, regardless of the magnitude of the difference between the two AOIs. By this binary-choice metric, infants in the

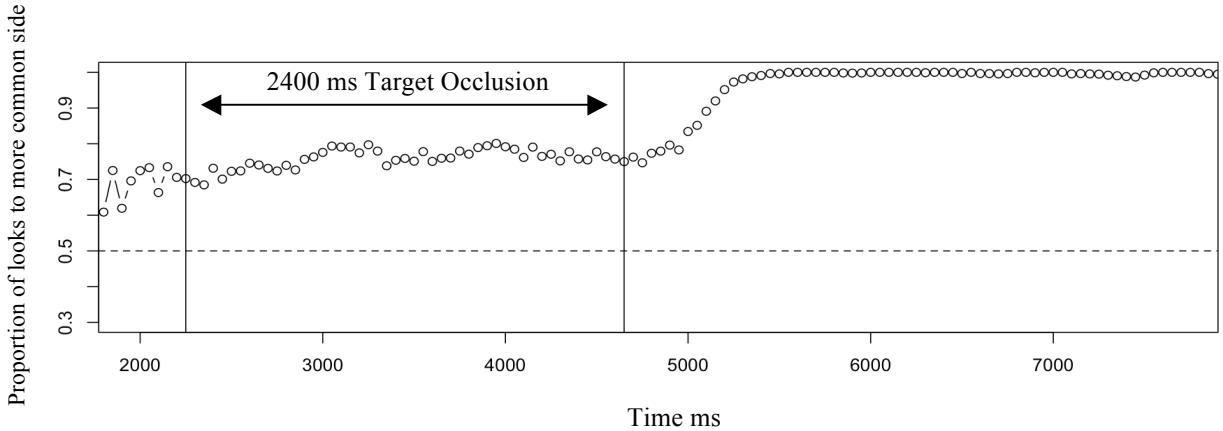

baseline condition chose the more common side on an

Figure 2: Proportion of looks to the more common side during the period when the target object was occluded

average of 76.2% of the trials. This performance was also significantly greater than chance ($p<0.001$).

Even though neither the proportion of looking time nor the binary-choice metrics indicated that infants *solely* attended to the 100% AOI, both measures were significantly above chance, thereby establishing a baseline for comparison with probabilistic designs.

Experiment 2: Probabilistic Exposure

Experiment 1 illustrated that infants can predict the reappearance of a target object when there is only evidence that the target reappears in one of two AOIs. The goal of Experiment 2 is to explore, in a probabilistic design with evidence of some target reappearance in each of the AOIs, whether infants will maximize (show the same looking behavior as in Experiment 1), or whether they will alter their looking behavior in response to the increase in uncertainty about where the target will reappear.

Methods

Participants. 15 parents volunteered their infants and were compensated in the same manner as Experiment 1. The infants ranged in age from 10.0 to 10.6 months (M=10.3). An additional 9 infants were tested but excluded due to experimenter error (1), fussiness (7), and eye-tracker calibration difficulties (1).

Materials and Apparatus. Same as in Experiment 1.

Procedure. The Procedure was identical to Experiment 1 except that across trials the target reappeared on the more common side only 70% of the time and on the other side 30% of the time. The infants were split between the left side more common (70L30R) and right side more common (30L70R) conditions.

Results

Infants completed an average of 32.6 trials (range: 22-52). As in Experiment 1, the proportion of looks to the right and left sides was calculated at each time point during the period of target occlusion. Infants in the 70L30R and 30L70R conditions did not differ in their proportions of looking time to the more common side ($p>0.1$), and so the two groups were collapsed for all analyses. Overall, infants spent 56% of their looking time on the more common (70%) side. This was not significantly greater than chance ($p=0.6$).

The binary choice metric was also calculated for each infant on each trial. This method revealed that infants in the probabilistic condition chose the more common side on an average of 58.6% of trials. This metric was also not significantly greater than chance ($p=0.19$).

Infants in this 70-30 probabilistic design did not spend significantly more than half of their looking time to the more common AOI, even though numerically there was a slight bias in the correct direction.

Comparison of Experiments

The critical comparison is between looking time to the more common side in the baseline (Experiment 1) and probabilistic (Experiment 2) conditions, which indicates whether infants show maximization in the 70-30 condition of Experiment 2. As stated above, infants in the baseline condition spent 75% of their anticipatory time looking to the more common (100%) AOI, whereas infants in the probabilistic (70%-30%) condition anticipated the more common AOI 56% of the time. This was a significant difference ($p<0.05$). Using the binary-choice metric, infants in the baseline condition chose the more common side on 76.2% of trials, whereas infants in the probabilistic condition chose that side on 58.6% of trials. This difference was marginally significant ($p=0.07$). These results suggest that infants in the baseline condition did deploy significantly more of their attention to the more common AOI than did infants in the probabilistic condition.

Discussion

Studies of probabilistic knowledge in infants have, prior to the present report, relied on post-event looking times (i.e., Xu & Garcia, 2008). Although highly informative, these measures do not provide information about expectations before the event occurs. The results of the present eye-tracking studies indicate that infants are able to track the motion of a target object during an occlusion task and that this is a viable paradigm for examining probabilistic expectations.

The critical comparison between performance on the baseline and probabilistic conditions indicates that infants are spending significantly less time on the more common side in the probabilistic condition. This indicates that infants are not maximizing, but are instead responding to the probabilities of their input. Given the behavior of young children on two-choice response tasks (Stevenson & Weir, 1959), it is somewhat surprising that 10-month-olds do not show maximization behavior. It is not clear, however, whether the baseline condition serves as an accurate indicator of how infants interpret deterministic events. Because performance in the 100-0% condition only reached 75%, the lower rate of 56% in the 70-30% condition, which did not differ from 50%, may have under-estimated performance in this probabilistic condition. Nevertheless, it is clear that infants did not show the same 75% rate of responding in the 70-30% condition, and therefore did not show maximizing.

In order to address this issue in the future, the results of Experiment 2 should be compared to a wider range of probability contrasts. For example, the 70-30 condition (Experiment 2) may be too close to 50-50 for infants to detect the probability difference. Thus, a 80-20 condition might reveal performance that is closer to maximizing. In addition, allowing an overt choice response such as reaching for the moving object, would permit comparisons between eye movements and other types of responses.

We know from visual tasks and miniature language learning studies that the number of alternatives available significantly alters performance (Gardner, 1957; Weir, 1964; Hudson Kam & Newport, 2009). By increasing the number of choice locations available in our tasks, we can, for example, compare performance on a 70-30 condition with a 70-15-15 condition. Differences in anticipation rates for the 70% choice across these 2-choice and 3-choice conditions might increase our understanding of how infants treat probabilistic information.

One of the benefits of the anticipatory eye movement paradigm used in the present experiments is that it can be used across a wide range of ages. Our results indicate that 10-month-old infants do not show the type of maximization that can be found in preschoolers. By using a modified version of this task with a wider age range we may be able to determine how eye gaze behavior changes across age groups, especially at ages when we observe maximizing on overt choice tasks (e.g., the light bulb paradigm).

Finally, although we have access to time-course data from the anticipatory eye movements, we are not able from this information alone to determine what representations the infants and children might have during the task. In particular, how might their looking time on a trial-by-trial basis relate to their beliefs about the distribution of events? Developing a computational model that used this eye-gaze data to map ongoing beliefs to behavior would greatly increase our understanding of this relationship.

Together the results of our first two experiments indicate that (1) we can use an anticipatory eye movement paradigm to collect information about expectations of probabilistic events, and (2) infants do not show evidence of maximization in a simple two-choice probabilistic task. In ongoing work we are using this paradigm to explore further questions about probabilistic expectations and response choice.

Acknowledgments

This research was supported by NIH Grants HD037082 to RNA and DC00167 to ELN. We thank Johnny Wen for assistance with analyses; Suzanne Horowitz, Alyssa Thatcher, and the research assistants of the Rochester Baby Lab for help recruiting and scheduling subjects; and Patricia Reeder and Celeste Kidd for advice, comments, and suggestions.

References

Austin, C. A. & Newport, E. L. (unpublished manuscript). Predictable versus unpredictable variation: Regularization in adult and child learners.

Fiser, J. & Aslin, R. R. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science, 12,* 499-504.

Fiser, J. & Aslin, R. R. (2002). Statistical learning of higher order temporal structure from visual shape-sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 130,* 658-680.

Gardner, R. A (1957). Probability-learning with two and three choices. *The American Journal of Psychology, 70*(2), 164-185.

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

Hudson Kam, C. L., & Newport, E. L. (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology, 59*, 30-66.

Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition, 83*(2), B35-B42.

McMurray, B., & Aslin, R. N. (2004). Anticipatory eye movements reveal infants' auditory and visual categories. *Infancy, 6*(2), 203-229.

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

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: The role of distributional cues. *Journal of Memory and Language, 35,* 606-621.

Shukla, M., Wen, J., White, K. S. & Aslin, R. N. (in press). SMART-T: A system for novel fully automated anticipatory eye-tracking paradigms. *Behavior Research Methods.*

Stevenson, H. W. & Weir, M. W. (1959). Variables affecting childrens' performance in a probability learning task. *Journal of Experimental Psychology, 57*(6), 403–412.

Toro, J. M. & Trabalón, J. B. (2005). Statistical computations over a speech stream in a rodent. *Perception& Psychophysics, 67,* 867-875.

Weir, M. W. (1964). Developmental changes in problem-solving strategies. *Psychological Review, 71(*6), 473-490.

Weir, M. W. (1972). Probability performance: Reinforcement procedure and number of alternatives. *The American Journal of Psychology, 85*(2), 261-270.

Xu. F. & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences, 105*(13), 5012.