Published in final edited form as:

Cogn Dev. 2018 ; 48: 105–116. doi:10.1016/j.cogdev.2018.06.008.

Probability Learning in an Uncertain World: How Children Adjust to Changing Contingencies

Sarah J. Starling^{1,2}, Patricia A. Reeder³, and Richard N. Aslin¹

¹ Department of Brain and Cognitive Sciences, University of Rochester, Rochester, NY, 14604

² Department of Social Sciences, Psychology Program, DeSales University, Center Valley, PA, 18034

³ Department of Psychological Science, Gustavus Adolphus College, St. Peter, MN, 56082

Abstract

We regularly make predictions about future events, even in a world where events occur probabilistically rather than deterministically. Our environment may even be non-stationary such that the probability of an event may change suddenly or from one context to another. 4–6 year olds and adults viewed 3 boxes and guessed the location of a hidden toy. After 80 trials with one set of probabilities assigned to the 3 boxes, the spatial distribution of these probabilities was altered. Adults easily responded to this change, with participants who *maximized* in the first half (by choosing the most common location at a higher rate than it was presented) being the fastest at making this shift. Only the older children successfully switched to the new location, with younger children either partially switching, perseverating on their original strategy, or failing to learn the first distribution, suggesting a fundamental development in children's response to changing probabilities.

Keywords

Probability learning; Non-stationary environment; Maximization; Predicting events

1. Introduction

1.1 Predicting future events

As learners, we are faced with the difficulty of extracting and interpreting information from a highly complicated environment. At any moment we must choose, from the wealth of possible cues available, the ones that are the most meaningful and reliable. There is not, however, always a perfect correlation between cues and their consequences, due to inconsistencies in how they are causally related. This may lead to classic induction problems

Corresponding author: Sarah J. Starling, Ph.D., Assistant Professor of Psychology, DeSales University, 2755 Station Avenue, Center Valley, PA 18034, 610-282-1100, ext. 1781, Sarah.Starling@desales.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

where, due to limited or conflicting information, the data available to a learner may support a range of differing hypotheses about how the world works. To add to this confusion, the efficacy of any particular cue as a learning tool may change across time and context. In order to successfully navigate such an environment, learners must find a way to respond to these varied forms of unpredictability in their input.

One way to guide our learning is to explore our environment in search of regularities. Rather than dividing our attention across all of the possible sources of information, efficient learners should direct their attention to the most commonly occurring and potentially predictive information available. Much evidence from the past few decades has demonstrated that adults (Saffran, Newport & Aslin, 1996; Fiser & Aslin, 2001, 2002), infants (Saffran, Aslin & Newport, 1996; Maye, Werker & Gerken, 2002), and animals (Toro & Trobalón, 2005) extract information about the distributional properties of stimuli, even in the absence of an explicit task or direct feedback about how cues and consequences are linked (see review by Aslin & Newport, 2012). In addition, a wealth of recent evidence has demonstrated that not only are human infants and children sensitive to this distributional information, but they can utilize it to make inferences about the likelihood of event outcomes.

For example, young children are highly sensitive to the causal relationships between events (Gopnik et al., 2004). By 8 months of age infants are able to determine the likelihood of potential outcomes and then use this information to make predictions about what future events should and should not occur (Tégla et al., 2007; Xu & Garcia, 2008). Moreover, children may use a mature, rational strategy for making inferences about causal events in the absence of feedback (Denison, Bonawitz, Gopnik, & Griffiths, 2013).

In an ideal world, one would want to predict *specific* events, but that ability is quite rare because most events are not cued with perfect reliability. For example, we can be certain that sunrise will follow sunset, but we are much less certain about whether sunrise will be followed by a sunny or a cloudy sky. We can, however, make *general* predictions by gathering information about base rates. For example, over the course of a year, we might observe that the ratio of sunny to cloudy days is 5:1 (San Diego) or 1:5 (Rochester). This base-rate estimate plays an important role in how one would prepare to greet the day: carry an umbrella in Rochester or apply sunscreen in San Diego. Thus, knowledge about distributions of events, in a given context, can influence our predictions and lead to successful outcomes. However, very few outcomes are predicted by a single cue. The presence of clouds is not the only cue to the likelihood of needing an umbrella, especially when the base rate of clouds is high.

Thus, in many domains, the information available to us when we need to predict future events may be inconsistent or contradictory. But in addition to this unpredictability is the fact that the distributions of events in our environment may change over time. Our future behavior will be influenced by whether we believe that our probabilistic environment is stationary or non-stationary. Stationarity assumes that the relevant probabilities stay the same over time, at least in a given context. So although we cannot perfectly predict upcoming events, the distribution of events will not change. If we expect a nonstationary

environment, however, then we know that the probabilities that we have learned thus far may shift. For example, as winter ends and spring begins, the likelihood of a sunny day increases and thus we need to update our expectations and behaviors accordingly. One methodology that is particularly well suited to exploring how learners interpret these types of inconsistencies is *probability learning*, which requires participants to predict future events in a probabilistic task.

1.2 Behavioral strategies in probability learning tasks

When faced with the task of predicting future events in a non-deterministic environment, a learner seeking to maximize accuracy or reward could employ one of two main 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 nearly always choose the more common outcome, a pattern known as *maximization* (c.f., Estes & Straughan, 1954). In several classic experiments, participants were presented with two light bulbs and on each trial were asked to predict which light would illuminate (e.g., Neimark, 1956; Gardner, 1957, 1958; Weir 1972). After participants made a choice, one of the bulbs would turn on. For example, one bulb turned on 70% of the time and the other bulb 30% of the time. In this situation, maximizing on the more probable alternative is the better strategy because it leads to higher overall accuracy. If the participants were probability matching (i.e., picking the 70% light on 70% of the trials and picking the 30% light on 30% of the trials), then their overall accuracy would average 58% correct (49% + 9% respectively). If, on the other hand, learners chose the 70% light on every trial, their overall accuracy would be 70% correct (70% + 0%). For this reason, maximization is the best behavioral pattern if (1) the environment is truly probabilistic (i.e., there is no deterministic pattern to the order of the lights), (2) the goal is to correctly choose the location of the light as often as possible and (3) the environment is stationary, meaning that there is never any change in the presented probabilities. It is not obvious, however, what the best approach would be in a non-stationary environment if our goal is not only to maximize reward in the short term but also to recognize a global shift in probabilities so that the learner can adjust their response pattern to optimally match the updated probabilities.

Studies of probability learning have demonstrated that highlighting the majority location, either by increasing its cue-salience (Gardner, 1957) or by increasing the number of minority alternatives (Gardner, 1957; Weir 1964, 1972), promotes the selection of the majority location above the level of probability matching. This same phenomenon has been found in auditory language learning experiments (Hudson Kam & Newport, 2009). This tendency to over-predict the majority choice may partially result from the fact that as the number of choices increases, the likelihood of each of the minority choices being correct decreases. This maximizing tendency is a rational response by adults to the memory demands of keeping track of multiple alternatives, especially when choices are based on a sparse sampling of the input.

Although maximizing in a stationary environment leads to an overall higher level of accuracy, adults tend to probability match rather than maximize in most simple choice-tasks (Gardner, 1957; Weir, 1964, 1972) and in language learning experiments (Austin &

Newport, 2012; Hudson Kam & Newport 2005, 2009). Children, however, are more likely than adults to show maximization or boosting behavior that enhances the choice of the majority location (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, such as poorer memory for the outcomes of past choices when there are multiple locations to keep track of. This same reliance on memory for past outcomes could form the basis for the influence of complexity on maximizing behavior in adults when they are faced with 3 or more choices (Gardner, 1957; Weir 1964, 1972). It could also be based on the fact that children require more data than adults to be confident that further exploration is not necessary to maximize performance on the task (as in Denison et al., 2013).

Evidence in support of these explanations for developmental differences in probability learning tasks comes from a study using a partial reinforcement design, in which one correct option is reinforced at a specific rate, while the other options never lead to reward. Weir (1964) tested participants from ages 3–18 on a three-choice test where the majority location was reinforced at 33% (33–0–0) or 66% (66–0–0). Across the final 20 trials, both the youngest and the oldest participants maximized more often than children aged 7–15, but they reached this behavior by different strategies. The 3–5 year old children reached asymptotic behavior very early, focusing on the one option that produced a reward. The danger of a short exploration period is that the explorer cannot immediately tell whether only one location is ever rewarded, as in this task, or whether they have just been unlucky in their first few guesses. The 18 year olds, on the other hand, came to this behavior slowly after exploring all of their options, possibly through a desire to find the best reward rate and discovering over time that the other two locations were unlikely to produce a reward. Additionally, school aged children are the most likely of any age group to pick a successive choice pattern, such as Left-Middle-Right or Right-Middle-Left and stick with it (Weir, 194; Craig and Meyers, 1963; Derks & Paclisanu, 1967). This suggests that young children will quickly settle on whichever location led to a reward, middle-aged children make use of repetitive patterns to explore the space but are unable to move past these attempts, and the oldest children are able to test multiple hypotheses until coming to the best solution.

1.2 Responding to shifts in the environment

The learning situation changes, however, when the assumption of stationarity is relaxed and a learner is exposed to a shift in probabilities. Consider a foraging task in which an animal has previously visited three food sources. Early in the season it may be that bush A is much more likely than bush B or bush C to have the good tasting berries (70–15–15). But as the year goes on, it is possible that this might change such that bush B is now the better bet (15–70–15). In order to make a rational choice about where to go for food, the animal must notice and respond to this change in the environment. If the animal is sampling according to the original 70–15–15 distribution, then once the majority location has lost its goodtasting berries, their choices must yield significantly inferior outcomes compared to past choices, and the decision to abandon the majority location must be followed by sampling either (or both) of the minority locations. This shift in sampling from the three locations is subject to considerable uncertainty, yet failing to shift runs the risk of grossly undershooting the

optimal sampling strategy. If the input distribution undergoes shifts only rarely, then staying with the initial sampling strategy has relatively little cost – the change will be discovered eventually as more data are gathered. However, if the input distribution shifts repeatedly after short periods of stability, then maximizing is a suboptimal strategy – it prevents the learner from sampling the current minority location(s) that could quickly become the majority location. Thus, learners should strike a balance between sampling and exploiting the alternatives (Kamil & Roitblat, 1985). If they sample the input extensively so that they have high confidence in the underlying contingencies, they will delay achieving the maximum set of rewards. If, on the other hand, they quickly settle on one choice and do not sample enough to explore other choices, they might miss out on the best choice in the future (Kearar, Rashkovich, Cohen & Shmida, 2002).

Given the differences seen in adult and child responses on probabilistic tasks, we can ask how adults and children will behave in a non-stationary environment. Data from the reversal-shift literature indicates that adults and children respond differently when given evidence that the relative importance of cues has changed. In a reversal-shift task a participant is given a discrimination task with a pair of stimuli that differ on a single dimension. For example objects may differ in both color and height. In the first part of the task color may be the discriminative cue with white leading to reward. In the second part of the task, a change is made in the rules such that the new rule to follow is either the opposite of the previous rule (reversal-shift) or is now based on a new feature of the task (nonreversal-shift). In a reversal-shift change, color might still be the discriminating factor, but now black, rather than white, leads to reward. In a nonreversal-shift, some other dimension becomes relevant such as the height of the object. Although young children (Jeffrey, 1965; Kendler, Kendler & Wells, 1960) and rats (Kelleher, 1956) find nonreversal-shifts easier to learn than reversal shifts, the opposite is found for adults and older children (Kendler & Mayzner, 1956; Kendler, Kendler & Learnard, 1962). Sanders (1971) demonstrated that second graders were much more willing than preschoolers to abandon a previously reinforced response. But if this reversal shift were verbally acknowledged with instructions explaining that a shift had occurred, even preschoolers could follow the reversal. Our ability to respond to changes in the importance of an environmental cue therefore depends not only on developmental factors but also on the type of change that occurs and the salience of the cue. This suggests that pre-school aged children should have a more difficult time responding to an unannounced shift in contingencies during a probability-learning task than would older children and adults.

Our innate curiosity about how the world works may be enough to discourage any behavioral strategy that does not allow for exploration. One explanation for the high rate of matching behavior in human adults is that probability matching is not a strategy per se but that it is a consequence of our search for patterns in the world (Wolford, Newman, Miller & Wig, 2004). That is, in order to fully maximize on one choice, a learner must be willing to give up the chance to explore other options. Given that we live in an environment that is often patterned, probability matching is an adaptive response that encourages the learner to constantly search for a potentially predictive pattern (Gaissmaier & Schooler, 2008). The adaptive benefit of such a selective pattern search hypothesis has found support in studies demonstrating that individuals who probability match are more likely to successfully exploit a pattern when they encounter it than those who maximize (Gaissmaier & Schooler, 2008). It

is possible, therefore, that a matching strategy (as is commonly seen in adults) may be beneficial for recognizing an unannounced change in the large-scale statistics of a probability-learning task. Alternatively, probability matching may be an efficient implicit strategy for sampling the alternatives without running the risk of failing to detect new structures, such as a shift in the underlying probability distribution of the majority alternative.

The claim that matching may serve us best in a changing environment does, however, assume that there is not perfect access to information about the environment. The type and availability of feedback may also influence choice behavior. In probability learning tasks, either a response-feedback or a choice-feedback design can be used, and this determines how much information is given on each trial. In a response-feedback probability learning design the participant is only informed about whether they were correct or incorrect, but when they guess incorrectly they are not told what would have been the correct choice. When there are three or more options available, if you are incorrect you don't know which of the other possible choices was the correct answer. In a choice-feedback design, however, no matter which choice is made, the participant is always informed of what the correct answer had been. Both children and adults choose the majority location at a significantly higher rate in choice-feedback than in response-feedback designs (Weir, 1972; Witting & Weir, 1971). The main explanation offered for this difference is that in the response-feedback task, because incorrect guesses can't signify which was the correct choice, it takes longer to build up the information necessary to form a representation of the statistics of the task.

Foraging situations are similar to response-feedback designs in that after making a bad choice (i.e., pick a flower that is not producing well) the forager is not informed which flower they *should* have visited. Additionally, if the previously poor-performing flower suddenly perks up, they have no way of knowing unless they have been sampling from it (thus making the compromise of cutting back on access to the best food source). In contrast, in a choice-feedback situation they aren't forced to sample from the minority location(s) in order to quickly become aware that the statistics of the environment have changed. Because they are always informed of the correct choice on each trial, a learner can safely maximize because they only need to pay attention to the feedback on each trial to see when a previously unrewarded choice is now reinforced. For that reason, probability matching, particularly if it is the result of pattern searching, might be counterproductive in a choice-feedback, or ideal learning environment. A learner who is busy searching for a pattern might be slow to recognize that the entire distribution of rewards has shifted, and such a pattern search is not needed in a choice-feedback design.

1.3 Goals of the present study

The goal of the present study was to examine how the behavioral strategies used by adults and children relate to their ability to recognize and respond to changes in the probabilities of a task structure. Adults and children (aged 4–6) were exposed to a 3-Alternative Forced Choice learning task in which they had to guess the location of a hidden object. The object location on each trial was probabilistically determined such that one of the three locations contained the hidden object on 70% of the trials, with the other two locations hiding the

object on the remaining 30% of trials (i.e., 15%, 15%). After 80 trials the distribution changed without warning such that a new location became the most common (at 70%). We refer to this as the majority location. By examining the choice behavior of participants before and after this change, any age related differences in the ability to recognize and respond to changes in a probabilistic structure were explored. Based on age related differences seen in the reversal shift literature (Jeffrey, 1965; Kendler, Kendler & Wells, 1960; Kendler & Mayzner, 1956; Kendler, Kendler & Learnard, 1962) we predicted that our adult participants and oldest children would be the most successful at detecting the change in the distribution and responding accordingly and they should show a preference for choosing the new majority location after the change in distribution. It was predicted, however, that the younger children might not be able to make this shift and instead perseverate on the rule learned in the first distribution.

Unlike many classic studies that rely on a response-feedback design in which participants are only informed of the accuracy of their guess on each trial and not the “correct” choice, we used a choicefeedback design in the present study. By informing the participant about the location of the hidden object on each trial, regardless of their choice, we hoped to speed up the learning process. This also means that a participant is able to maximize on the most common location while still learning about the relative occurrence of each of the less common choices. Thus it was predicted that maximization, in addition to probability matching, would be a successful strategy for recognizing the change, unlike in traditional foraging situations where maximization is a suboptimal strategy. To further explore this, the relationship between speed of response to the switch and the behavioral strategy before the switch was examined. It was predicted that participants who maximized in the first half of the experiment would be faster at recognizing the change in the distribution of events than participants who chose the first majority location at a lower rate. Additionally a three-choice task was used because the large difference in reward level between the majority and the two minority locations before and after the switch should also help to highlight the change in the spatial distribution of the probabilities. These design features should allow for a shorter task in which young participants are able to build a representation of the first distribution and have experience with the second distribution before becoming too tired or bored.

2. Method

2.1 Participants

Children were recruited through a (a) database of former infant participants, (b) at two local daycares, and (c) from a summer daycare and activity program. Adults were paid \$5 for their participation. After completion of the study children tested in the lab were offered \$10 or a small prize (such as a t-shirt or tote bag), and children tested at a daycare or summer program were given a small bag of toys (such as bouncy balls and bubbles). Informed consent was obtained from the adults immediately before participation in the study. Consent was obtained for the children either immediately before the study (if completed in the lab) or previously (if tested in a daycare or summer program).

The final sample consisted of 35 children (21 female) with a mean age of 59.1 months (range 48.2–81.8 months) and 16 adults (9 female, age range 18–29). All adults either had

completed their undergraduate degree or were currently enrolled in college. An additional 12 children were removed from analyses because of loss of interest (9), computer malfunction (1), extreme inattentiveness to the task (1), or because they had participated in a previous version of the experiment (1).

2.2 Testing Equipment

Adults were tested on one of two Dell Touchscreen-enabled computers in the laboratory. Children were tested either in a child-friendly testing room in the laboratory or in a quiet room at a local daycare or summer program. In the laboratory, the experiments were run on a Dell Touchscreen, and at the daycare or summer program the experiments were run on a Lenovo ThinkPad Touchscreen Tablet Laptop.

2.3 Stimuli

To provide children with an interesting task that would maintain their attention for dozens of trials, we used a total of 160 color images of children's toys (see examples in Figure 1). The toys were virtually hidden inside of square boxes colored red, yellow or blue. A "game show ding" sound was also used to indicate when the participant correctly guessed the toy's location.

2.4 Procedure

Participants were seated in front of the touchscreen computer and were told that they would be playing a guessing game. To start the game, participants were instructed to press the white "Go!" button on the touchscreen. This immediately started the first trial and brought up three colored boxes (red, yellow and blue) as shown in the left panel of Figure 2. The participants were informed that these were toy boxes and that on each trial a toy would be hidden inside one of the boxes. Their task on each trial was to guess where they believed a toy was hiding. Upon pressing a box, a toy emerged out of the correct box on that trial, by slowly rising vertically, and then stopped above the box as shown in the right panel of Figure 2.

If the choice by the participant was the correct box (i.e., where the toy actually was on that trial), it was additionally indicated by an auditory "ding" sound. If the incorrect location was chosen, the toy still rose out of the correct box, but no sound was played. The participant then had to click on the toy to make it disappear, and the "Go!" button immediately reappeared to start the next trial. There was no time limit and participants could take as long as they wanted to make their choice on each trial.

An experimenter was present for the entire time that the child participants performed the task. For the first few blocks or until the child clearly understood the procedure, the experimenter prompted the child to "guess where a toy is hiding". To further ensure that the child understood that the toys were actually coming out of one of the boxes on the screen, the experimenter also commented on the outcome of the first few trials such as "Look, the blocks were in the red box" or "You found the car in the yellow box".

Participants completed a total of 8 blocks of 20 trials. The correct location for each trial within each block was pre-set and pseudo-randomized across blocks. Every block represented the overall statistics of the experiment (70%–15%–15%), but the order of toy location was constrained such that there were no strings of five or more trials in a row for which the same box was the toy location. Participants completed four blocks (80 trials) with a fixed majority box location (i.e., 70% at Red), followed by four blocks with a different majority box location (i.e., 70% at Blue). The boxes themselves did not move or change color or position across all 160 trials. Thus, only the statistics of the task changed after the first 80 trials, such that although the same 70%–15%–15% percentages were used, one of the 15% boxes was now at 70% and the former majority box dropped to 15%. All six possible combinations of first and second location for the majority box were used across participants.

In order to help maintain children's focus and interest, there was a short break between each block during which the participants were given a sticker. Children were given a background picture (such as a backyard scene with a doghouse or a picture of a fishbowl) and an appropriate set of stickers (such as dogs or fish) to attach to the background. They were told that every time a purple star appeared on the touchscreen (the signal that a block had ended) they could put a sticker on their background picture. Adults were also provided with sticker rewards at the end of each block, but in addition were given information about what percentage of the experiment they had completed (e.g., 12.5%, 25%, etc.).

3. Results

3.1 Behavioral strategies after learning each probability distribution

To explore the behavioral strategies used by adults and children at the end of the 80 trials where the majority box location was invariant, each participant was classified according to their choice behavior on blocks 3–4, the final 40 trials of the first probability distribution, and blocks 7–8, the final 40 trials of the second probability distribution. We focused on these final 2 blocks so as to give participants time to reach their final behavior. Their percent choice of the majority location was used to categorize them as: Maximizers (>85%), Probability Matchers (55%–85%), or Low Choice (<55%).² A 2×3 Fisher's exact test found that the categorization of behavioral types in the final two blocks of the first distribution across age groups was not significantly different ($p=0.13$) (see Figure 3). This suggests that overall choice of the majority box across participants in blocks 3–4 was similar (68.2% of trials for children and 75.8% for adults). While 15/16 (93.8%) adults matched or over-predicted the majority location, 27/35 (77.1%) of children did so, with the other 8 children (22.9%) choosing the majority location less often than it occurred. A similar analysis of the last 40 trials from the second probability distribution yielded an age related difference in behavior. A 2×3 Fisher's exact test found that the categorization of behavioral types across

²A range of behaviors that is called matching is not well defined in the literature, and thus the lower bound of our matching range was chosen based on the following reasoning. A participant who chose the majority location on 55% of trials and split their remaining choices equally between the two minority locations would end up with a 55%–22.5%–22.5% split. This would mean that the majority location would not only be selected on the majority of trials but it would also be chosen with at least a 2-to-1 ratio over either minority location. Thus, selecting the majority location less than 55% was labeled Low Choice.

age groups for blocks 7–8 was significantly different ($p=0.04$). While 12/16 (75%) adults matched or over-predicted the majority location, only 14/35 (40%) of children did so.

Using data from the last two blocks of the two probability distributions, a two-way repeated measures analysis of variance with age group and distribution (pre or post-switch) as factors was also run. There was a main effect of age, with adults choosing the relevant majority location more often than children ($F(1,49) = 6.38, p=0.01, \eta^2 = .12$), and a main effect of distribution, with participants choosing the relevant majority location more often before the switch than after ($F(1,49)=23.25, p<0.0001, \eta^2 = .32$). An age $\times$ pre/post switch interaction did not reach significance ($F(1,49)=3.25, p=0.08, \eta^2 = .06$), but it may suggest that children were more likely than adults to have their pre-switch choices interfere with their ability to move on to the post-switch distribution. This possibility is addressed in further analyses.

To further examine this trending interference effect, terminal performance in both halves of the task was compared to indicate how well individual participants were able to fully recover after the change in the probability distributions. Average choice of the majority location in the last two blocks of each distribution (blocks 3–4 and 7–8) was calculated for each participant and the correlation between these values was determined. On average adults showed a slightly higher level of choice of the majority location in the last two blocks before the switch than in the final two blocks after the switch (75.8% and 67.2% respectively). A correlation revealed that the extent of choice of the majority location before and after the switch was significantly related, $r = .547, N = 16, p = 0.029$, two tailed. However, inspection of the data revealed that one individual (who happened to be categorized as “Low Choice” in both halves) drove this correlation. When this participant is removed, the correlation was no longer significant, $r = .481, N = 15, p = 0.069$, two tailed.

Overall children showed a much higher level of choice of the majority location in the final two blocks of the first distribution than in the second distribution (68.2% vs. 44.4%). A two-tailed t-test found this difference in choice level to be significant, $t(18) = 3.09, p < .01, d = 0.86$. However, unlike adults, the correlation between choice of the majority location in the final 40 trials before and after the switch was significant, $r = .35, N = 35, p = 0.04$, two tailed. While level of choice for the end of the first block was not predictive of level of choice for the second distribution for adults, it was so for children. Thus the level of success obtained by children in the first distribution predicted their subsequent success in the second distribution.

3.2 Strategies in response to the change in probability distribution

Given that the change in majority location from the pre- to post-switch probability distributions was not explicitly marked, either by instruction or by altering the visual display, there was an inevitable disruption in participants' responses immediately after the switch. Individual children displayed a range of behaviors from pre- to post-switch blocks: some children appeared to recognize and respond accordingly to the change in the probability distribution (for example see Figure 4), while others did not appear to change their strategy after the switch (for example see Figure 5).

To explore the range of pre- vs. post-shift behaviors shown by adults and children, each participant was identified as exhibiting one of four behavior patterns based on their average choice of the majority location in blocks 4 and 8 (the last block in each of the two probability distributions). First, we excluded participants who did not choose the majority location in Block 4 (Maj-1) on at least 50% of the trials because this implies they did not understand the task or were otherwise performing unreliably. Additionally, we excluded one participant who inexplicably chose the third box, which was never the majority location, on half of the trials after the switch. We refer to both of these types of participants as falling in the “other” category. Second, given Maj-1 was greater than 50%, three additional categories were defined based on how participants responded during Block 8 to the new majority location (called Maj-2). In particular we asked whether or not the participant fully or partially shifted their focus to the new majority location or whether they remained with the original majority location:

- **Switch:** Maj2 was greater than either of the other two locations (as Maj1 was in the first half).
- **Reduced:** Maj-1 remained higher than the other two locations, including Maj-2, but Maj1 dropped by at least 15% from its pre-switch level.
- **Stay:** Maj1 remained higher than the other two locations and did not drop more than 10% from its preswitch level.

Using this categorization system, the children and adults exhibited the frequency distribution shown in Figure 6. For children: Switch (19), Reduced (5), Stay (5) and Other (6); and for adults: Switch (14) and Reduced (2). A 2×4 Fisher’s exact test found that the categorization distribution across age groups was trending ($p=0.07$).

We then asked whether age was a predictor of switching. Overall, the children who showed the Switch pattern were the oldest, followed by those who showed the Reduced pattern, the Stay pattern, and finally the Other pattern (see Figure 7 below).

A one-way analysis of variance revealed a significant main effect of child age for these four categorization patterns, $F(3,31) = 4.43$, $p=0.01$, $\eta^2 = .30$. Post hoc comparisons using the Tukey HSD test indicated that the average age for the Switch group ($M=62.4$, $SD=1.8$) was significantly greater than the average age for the Other group ($M=52.1$, $SD=2.1$) ($p=0.01$). No other direct comparisons reached significance.

3.3 Effect of first half behavioral strategy and speed of recovery

Another metric of the disrupting effect of the switch in probability distributions is how quickly participants respond to the switch in the majority location. For example, a participant who spent the first half of the experiment maximizing on the majority location might be particularly quick to notice when that location suddenly became unproductive. To further explore this possibility, speed of recovery was examined by computing how rapidly participants’ responses shifted to Maj-2. We used a 5-trial moving window (i.e., a smoothing algorithm) to deal with the discrete nature of responses on our task (i.e., the fact that only 1 of 3 choices was possible on each trial). By smoothing over a small set of trials we are able to see how behavior is changing within and across groups of participants while accounting

for the high variability we would expect to find across participants and trial-to-trial. The first window is the average choice behavior for trials 1–5 after the switch, the second window for trials 2–6, and so on. For each participant, this created 10 data points over the first 14 trials. This 5-trial time window was chosen because it reduced trial-to-trial variability for time windows with fewer trials and yet provided a measure of each participant's learning curve. The first 14 trials were the focus of this learning curve because visual inspection of the data showed that participants tended to reach asymptotic performance by the 15th trial. We then averaged these smoothed functions across participants and plotted them separately for the 12 adults who were Probability Matchers in the final two blocks of the first half of the experiment and the 3 adults who were Maximizers. Because only one adult participant fell in the Low Choice category, that category was removed from further analyses so as to focus on those participants who clearly mastered the first probability distribution.

Figure 8 shows average proportion choice in the 10 temporal windows for the two behavioral pattern types (Probability Matchers and Maximizers). We fit a linear mixed model to the data using percent choice of the new majority location as the dependent variable, with both behavioral category in the first half (Match and Max) and log of the window number as independent variables, and category by log of window as the interaction term. The model demonstrated that Match participants started after the switch with a predicted accuracy intercept of 0.25. The Max participants started with a significantly higher predicted accuracy of choice of the new majority location with an intercept of 0.47, $F(1, 146) = 80.7$, $p < 0.001$. Both groups showed a positive increase in choice probability. The Match group had a positive slope of $\beta = 0.07$ and the Max group had a slope of $\beta = 0.23$. These slopes are significantly different, $F(1, 146) = 4.65$, $p < 0.05$. These analyses support the hypothesis that Maximizers more readily switch their choice behavior after a change in the probability distribution than do Matchers.

The same analysis was conducted on the child data, again only with those participants who demonstrated successful learning of the first probability distribution. The 8 children with Low Choice patterns were removed and this left 26 children: 15 Probability Matchers and 12 Maximizers (see Figure 9). The model demonstrated that Match participants started after the switch with a predicted accuracy intercept of 0.27. The Max participants started with a significantly higher predicted accuracy of choice of the new majority location with an intercept of 0.40, $F(1, 266) = 4.58$, $p < 0.05$. The Match group had a slight but not significantly negative slope of $\beta = -0.02$. The Max group had a slope of $\beta = -0.07$ which is not significantly different from the Match group, $F(1, 266) = 1.47$, $p = 0.23$. These analyses suggest that both adult and child maximizers respond to the switch in distribution more quickly than do matchers. However, the child maximizers, in contrast to adults, do not show a significantly higher slope than matchers. This suggests that this initial benefit does not appear to hold true for the children after the first few trials..

4. Discussion

The classic probability learning literature provides a method for examining how individuals make predictions about upcoming events in a probabilistic environment. In a stationary probabilistic learning environment adults tend to probability match by directly matching the

exposure probabilities observed in the environment or slightly over predicting the most common event, whereas children are more likely to maximize by focusing on the most likely option observed across trials (i.e. Gardner, 1957; Stevenson & Weir, 1959). However, the probabilities that we experience in everyday life are rarely stationary. Thus, it is possible that the developmental differences observed in probability learning tasks are at least partially reflective of how adults and children deal with non-stationarity, and not entirely about fundamental differences in probability learning per se.

4.1 Developmental differences in performance

In the current study preschool aged children and adults were exposed to such a non-stationary task. We found that tracking and responding to the change in the probabilities was not difficult for the adults, as would be expected given previous evidence that adults can track local changes in statistics (Edwards, 1961; Qian, Jaeger & Aslin, under review). Their average choice of the majority location at the end of each half of the experiment was relatively high (only 1 participant failed to either probability match or maximize responding to the majority location), and correlations between their average choice across the two halves of the experiment were trending.

In comparison to adults, children showed much more variable performance. Slightly more than half of our child participants appeared to truly respond to the change in spatial distribution of the probabilities by “Switching” strategies and now attending to the new majority location. This could be due to a difference in their underlying expectations about the world. If children expect a stationary environment, then they may be disinclined to believe that global rules may shift. Perhaps young children have a higher threshold than do adults for variability in their environment before they accept that the distribution has changed. It may take additional maturity or experience for them to discover that the variability in a task is due to global changes rather than local ones. Alternatively, it is possible that our younger participants didn’t have enough data to convince them to throw out their initial hypothesis about the learning task. With additional time and exposure, they may have gained the input needed to be certain that a change in the underlying probability distributions had occurred.

Either of these explanations is consistent with our finding that the children who demonstrated a switching behavior by changing their choice behavior in response to the change in the probability distributions were the oldest on average (mean age of 62.4 months). This is mirrored in the behavior we observed in adults. Those children who chose a strategy that was unrelated to the input by focusing entirely on a minority location were the youngest on average (mean age of 51.5 months). This significant age effect is evidence of an ongoing developmental change in strategy or general ability to track changes in statistics that begins to emerge during the preschool years. Although it did not quite reach statistical significance, those children who made the shift were also older on average than those who “stayed” (mean age of 55.1 months) by perseverating in choosing the box that was most rewarded in the first half of the study. This is in accord with findings from the reversal-shift literature showing that school aged children are more likely than preschoolers to give up on a previously learned rule (Sanders, 1971). In the present experiment, only the oldest children

were able to both learn the first rule (e.g., blue box is most common) and then later accept a new rule (e.g., red box is now the most common). As Sanders (1971) demonstrated, preschool aged children are also much more willing to abandon a previous rule when they are specifically told that a change has occurred. It is possible that these younger children, who “stayed” or merely “reduced” their responding to the previously relevant majority location (Maj-1), might have responded to the change in the majority location (Maj-2) if there had been some external cue signaling the change in the environment.

Of course they may also be a range of underlying cognitive changes that play a role in the agebased changes we see in this task but these are unlikely to be the sole explanation of what we find. For example, some of the differences we see between the adult and child participants, and within our child group may be driven by differing motivational factors. The only rewards that participants receive for a correct choice were a simple auditory feedback and the knowledge that they were learning. This is quite minimal when compared to a food reward for a foraging animal. Future work may explore whether differences in motivation are partially responsible for developmental changes in probability matching. Our goal here, however, is not to provide a mechanistic account of the behavior, but to demonstrate the changes that occur across development. Future detailed studies will be needed to sort out the relative role of attention and working memory, among other skills, on the behavior of children and adults in probability learning tasks.

4.2 Value of maximization behavior

One hypothesis about why learners might probability match, when this is clearly a sub-optimal strategy, is that probability matching allows for continual exploration of non-majority locations. If learners quickly maximized, they would never be open to the possibility that the majority location has undergone a switch. Although previous findings have demonstrated that probability matchers are more likely than maximizers to recognize a pattern when it occurs (Gaissmaier & Schooler, 2008), in the present study it was the adult maximizers who were fastest at recognizing the shift in the probability distribution (from Maj-1 to Maj-2). For the matching participants there was a slow and steady increase in choice of the new majority box over the first 14 trials after the switch. For the maximizers, however, there was an immediate change in response behavior as seen in an increase in choice followed by a plateau (see Figure 8).

The basis for this rapid switching response in maximizers is not entirely clear. One explanation for the fast recovery of the maximizers focuses on the high level of accuracy maximizers would have experienced at the end of the first half of the task. As Goodnow (1955) explains, in order for a participant to show maximization behavior, they must come into the task with a specific set of expectations or beliefs. Although maximization will lead to a higher rate of accuracy than probability matching, it can never result in perfect accuracy. Thus, in order to maximize, a participant must make a strong commitment to one option and accept a level of accuracy less than 100%. We can consider two “ideal learners” – the perfect maximizer and the perfect probability matcher. In this example, a perfect maximizer will choose the 70% location on 100% of trials. They give up on the possibility of correctly guessing on 3 out of every 10 trials in order to get that relatively high level of accuracy. By

choosing all three locations at the rate they were presented, a perfect probability matcher, on the other hand, will only be accurate on 53.5% of trials. By hoping for perfection, probability matchers actually have lower overall accuracy. Now if the switch occurs and they do not respond in any way, the perfect maximizer will drop from 70% accuracy to only 15% accuracy (because they are only selecting a minority location). A perfect probability matcher who makes no change, on the other hand, will drop from 53.5% to 23% accuracy. This occurs because they now select one of the 15% locations on 70% of trials (10.5% accuracy), and select the 70% and the other 15% location each 15% of the time (10.5% and 2.3% accuracy respectively). This is a relatively small change when compared to that experienced by the maximizers. For this reason, it is possible that a maximizer would be faster to recognize this dramatic change, as can be seen in their almost immediate high rate of choice of the new majority location. Additionally, a probability matcher who reaches this behavioral style through pattern searching may be too focused on the goal of finding the “perfect solution” to take full advantage of the available feedback.

This response to the drop in level of reward may also interact with the type of design that is used. In the choice-feedback design used in the present study, the learner is given information about the correct location of the hidden object on each trial regardless of their actual choice. This means that the participant does not need to actually sample from the minority locations in order to discover how often they are rewarded. When a response-feedback design is used rather than this ideal learning situation, how will maximizers fare? In both types of designs when the switch happens the maximizer will experience a larger drop in reward than the matcher, but in a response-feedback design they will not have been told which location they should now favor. Future work employing a response-feedback design may find that the willingness of the probability matcher to sample from the minority locations works to their benefit when detecting a shift in the probabilities.

Although we find evidence for faster recovery after the switch for adult maximizers than for adult probability matchers (a maximization benefit), this was a more mixed result in the children. Although the children who maximized did have a higher rate of choice of the new majority location in the first 5 trial window, they did not have a higher slope than the matchers. Unlike the adults where the choice behavior for the two groups clearly diverged, for children this initial benefit quickly disappeared (see figure 9). One possible explanation for this developmental difference is that to take full advantage of being a maximizer, even when feedback is available on each trial when that guess is wrong, is something that preschoolers find very difficult. Denison et al. (2013) propose the “Sampling Hypothesis” to account for this difficulty. They argue that the inconsistency demonstrated in child behavior on tasks similar to those used here may result from the fact that children are sampling from a large set of possible hypotheses about the world. Application of those hypotheses will depend on the likelihood of each. By this view, when children probability match (as 26% did in our task) it could be due to them using a complex, rational approach to generating hypotheses (i.e., implicitly juggling many competing hypotheses). It is possible that the higher rate of maximization seen in children (42% of our child participants) is based on a developmental stage in which only a single hypothesis is implicitly available in a given task context. If this single hypothesis is more entrenched in children than in adults, it may not have allowed them to have the maximization benefit seen in adults who remain open to

competing hypotheses despite having a firm grasp of the fact that there is a stable majority location.

5. Conclusion

In summary, the present study demonstrated that adults and older preschool aged children are capable of responding to an unannounced shift in outcome probabilities. For adults, this happened fastest when they were already employing a maximizing strategy. In our daily lives, however, we may be faced with multiple such shifts. Current work is exploring whether adults and children can track multiple changes. In particular, we might ask both whether adults and those older children who were able to switch would be able to retain that first distribution and swiftly return to it at a later time and whether the evidence of such shifts alters their behavioral strategies.

Acknowledgments

The work was supported by an NID research grant to RNA (HD-037082).

References

- Aslin RN and Newport EL (2012). Statistical learning: From acquiring specific items to forming general rules. *Current Directions in Psychological Science*, 21, 170–176. [PubMed: 24000273]
- Austin AC, & Newport EL (2012). When children learn more than they are taught: Regularization in child and adult learners (Unpublished doctoral dissertation). University of Rochester, Rochester NY.
- Craig GJ, & Myers JL (1963). A developmental study of sequential two-choice decision making. *Child Development*, 34(2), 483–493. [PubMed: 14023622]
- Denison S, Bonawitz E, Gopnik A, & Griffiths TL (2013). Rational variability in children's causal inferences: The Sampling Hypothesis. *Cognition*, 126, 285–300. [PubMed: 23200511]
- Derks PL, & Paclisanu MI (1967). Simple strategies in binary prediction by children and adults. *Journal of Experimental Psychology*, 73(2), 278.
- Edwards W (1961). Probability learning in 1000 trials. *Journal of Experimental Psychology*, 62(4), 385. [PubMed: 13889318]
- Estes WK & Straughan JH (1954). Analysis of a verbal conditioning situation in terms of statistical learning theory. *Journal of Experimental Psychology*, 47, 225–234. [PubMed: 13152299]
- Fiser J, & Aslin RN (2001). Unsupervised Statistical Learning of Higher-Order Spatial Structures from Visual Scenes. *Psychological Science*, 12(6), 499–504. [PubMed: 11760138]
- Fiser J, & Aslin RN (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(3), 458–467.
- Gaissmaier W, & Schooler LJ (2008). The smart potential behind probability matching. *Cognition*, 109(3), 416–422. [PubMed: 19019351]
- Gardner RA (1957). Probability-learning with two and three choices. *The American Journal of Psychology*, 70(2), 174–185. [PubMed: 13424757]
- Gardner RA (1958). Multiple-choice decision-behavior. *American Journal of Psychology*, 71, 710–717. [PubMed: 13627279]
- Goodnow JJ (1955). Determinants of Choice-Distribution in Two-Choice Situations. *The American Journal of Psychology*, 68(1), 106–116. [PubMed: 14350105]
- Gopnik A, Glymour C, Sobel DM, Schulz LE, Kushnir T, & Danks D (2004). A Theory of Causal Learning in Children: Causal Maps and Bayes Nets. *Psychological Review*, 111(1), 3–32. [PubMed: 14756583]

- Hudson Kam CL, & Newport EL (2005). Regularizing Unpredictable Variation: The Roles of Adult and Child Learners in Language Formation and Change. *Language Learning and Development*, 1(2), 151–195.
- Hudson Kam CL, & Newport EL (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology*, 59(1), 30–66. [PubMed: 19324332]
- Jeffrey W (1965). Variables Affecting Reversal-Shifts in Young-Children. *American Journal of Psychology*, 78(4), 589–595. [PubMed: 5839926]
- Kamil AC & Roitblat HL (1985). The ecology of foraging behavior: implications for animal learning. *Annual Review of Psychology*, 36, 141–169.
- Keasar T, Rashkovich R, Cohen D, & Shmida A (2002). Bees in two-armed bandit situations: foraging choices and possible decision mechanisms. *Behavioral Ecology*, 13(6), 757–765.
- Kelleher RT (1956). Discrimination learning as a function of reversal and nonreversal shifts. *Journal of Experimental Psychology*, 51(6), 379. [PubMed: 13319609]
- Kendler TS, Kendler HH, & Learnard B (1962). Mediated responses to size and brightness as a function of age. *The American Journal of Psychology*, 75(4), 571–586. [PubMed: 14031904]
- Kendler TS, Kendler HH, & Wells D (1960). Reversal and nonreversal shifts in nursery school children. *Journal of comparative and Physiological Psychology*, 53(1), 83–88. [PubMed: 14408484]
- Kendler HH, & Mayzner MS (1956). Reversal and nonreversal shifts in cardsorting tests with two or four sorting categories. *Journal of Experimental Psychology*, 51(4), 244–248. [PubMed: 13306872]
- Maye J, Werker JF, & Gerken LA (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition*, 82(3), B101–B111. [PubMed: 11747867]
- Neimark ED (1956). Effects of type of nonreinforcement and number of alternatives in two verbal conditioning situations. *Journal of Experimental Psychology*, 52(4), 209–220. [PubMed: 13367343]
- Saffran JR, Aslin RN, & Newport EL (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928. [PubMed: 8943209]
- Saffran JR, Newport EL, & Aslin RN (1996). Word segmentation: The role of distributional cues. *Journal of Memory and Language*, 35, 606–621.
- Sanders B (1971). Factors affecting reversal and nonreversal shifts in rats and children. *Journal of Comparative and Physiological Psychology*, 74(2), 192.
- Stevenson HW, & Weir MW (1959). Variables affecting children's performance in a probability learning task. *Journal of Experimental Psychology*, 57(6), 403. [PubMed: 13654658]
- Tégls E, Girotto V, Gonzalez M, & Bonatti LL (2007). Intuitions of probabilities shape expectations about the future at 12 months and beyond. *Proceedings of the National Academy of Sciences*, 104(48), 19156.
- Qian T, Jaeger TF, & Aslin RN (revision under review). Incremental implicit learning of bundles of statistical patterns. *Cognition*.
- Toro J, & Trobalón J (2005). Statistical computations over a speech stream in a rodent. *Perception & Psychophysics*, 67(5), 867–875. [PubMed: 16334058]
- Weir MW (1964). Developmental changes in problem-solving strategies. *Psychological Review*, 71(6), 473–490. [PubMed: 14218072]
- Weir MW (1972). Probability performance: Reinforcement procedure and number of alternatives. *The American Journal of Psychology*, 85(2), 261–270.
- Witting MA, & Weir MW (1971). The Role of Reinforcement Procedure in Children's Probability Learning as a Function of Age and Number of Response Alternatives. *Journal of Experimental Child Psychology*, 12, 228–239.
- Wolford G, Newman SE, Miller MB, & Wig GS (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology*, 58(4), 221–228. [PubMed: 15648726]
- Xu F, & Garcia V (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences*, 105(13), 5012–5015.

Highlights

- 4–6 year olds and adults participated in a 3-Alternative Forced Choice guessing game to find a hidden toy.
- Toy location was probabilistically determined with a 70%–15%–15% distribution.
- The distribution of events changed without warning after 80 trials.
- Adults and older children responded to the change but younger children did not.
- Suggests a developmental shift in ability to track changes in a non-stationary task.

Figure 1.
Examples of toys used in the experiment

Figure 2.
Screenshot of experiment before (image on left) and after (image on right) a choice is made.
This is identical regardless of the accuracy of the choice.

Figure 3.
Proportion of adult and child participants who fell into each of the three behavioral categories in the final 40 trials of the first and second distributions

Figure 4.
 Sample child participant (aged 62.7 months) who switched from majority choice of the 1st distribution majority location to the 2nd distribution majority location

Figure 5.
Sample child participant (aged 57.1 months) who stayed with majority choice of the 1st distribution majority location for the entire study even after the switch in majority box location

Figure 6.

Categorization of adult and child behavior in response to the change in distribution. This is based on a comparison of their level of choice of the current majority location in the final block of 20 trials both before and after the change in distribution.

Figure 7.
Average age of child participants by response to the change in distribution

Figure 8.
5-trial moving window analysis showing average choice of the new majority location choice for adults at the start of the second distribution.

Figure 9.

5-trial moving window analysis showing average choice of the new majority location choice for children at the start of the second distribution.