Cognitive Development 48 (2018) 105–116

Contents lists available at ScienceDirect

Cognitive Development

journal homepage: www.elsevier.com/locate/cogdev

Probability learning in an uncertain world: How children adjust to changing contingencies[☆]

Sarah J. Starling[a,*], Patricia A. Reeder[b], Richard N. Aslin[c]

[a] Department of Social Sciences, Psychology Program, DeSales University, Center Valley, PA, 18034, United States
[b] Department of Psychological Science, Gustavus Adolphus College, St. Peter, MN, 56082, United States
[c] Haskins Laboratories, New Haven, CT, 06511, United States

ARTICLE INFO

Keywords:
Probability learning
Non-stationary environment
Maximization
Predicting events

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.

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 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

☆ The work was supported by an NID research grant to RNA (HD-037082).
* Corresponding author at: DeSales University, 2755 Station Avenue, Center Valley, PA, 18034, United States.
E-mail address: Sarah.Starling@desales.edu (S.J. Starling).

https://doi.org/10.1016/j.cogdev.2018.06.008
Received 17 August 2016; Received in revised form 19 October 2017; Accepted 25 June 2018
0885-2014/

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 (Téglas, Girotto, Gonzalez, & Bonatti, 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 non-stationary 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