

Limited evidence for probability matching as a strategy in probability learning tasks

Jessica L. Montag*

Department of Psychology, University of Illinois, Urbana-Champaign, Champaign, IL, United States

*Corresponding author: e-mail address: jmontag@illinois.edu

Contents

1. Non-human animal probability learning	4
1.1 Typical task design	4
1.2 Findings	5
1.3 Summary	13
2. Human probability learning: Adults	14
2.1 Typical task designs	14
2.2 Findings	15
2.3 Summary	21
3. Human probability learning: Children	22
3.1 Typical task design	22
3.2 Findings	23
3.3 Summary	25
4. Implications of probability learning for other domains	26
4.1 Probability learning and language development	26
4.2 A note about generalization	30
4.3 Implications for other domains	31
4.4 Implications for the process of science	32
5. Conclusions	34
Acknowledgments	35
References	35

Abstract

In probability learning experiments, a participant is typically presented with one of two alternatives to select, one of which will lead to a reward. For example, in a 70:30 task, one alternative will lead to a reward on 70% of trials while the other will yield a reward on the remaining 30% of trials. On probability learning tasks, adults are said to “probability match,” selecting each alternative with the relative frequency with which it has been reinforced. Children, on the other hand, are said to “maximize,” always guessing whichever alternative has been reinforced more often. The different patterns

between adult and child behavior are thought to have implications for language learning, especially qualitative differences in child and adult language learning skills and developmental trajectories on a range of other cognitive tasks. However, a thorough review of the literature suggests that behavioral profiles of adults and children are not as straightforward as has been claimed. Crucially, there is little empirical support for a true probability matching strategy by any participants. Differences in features of the experimental task and in meta-task knowledge contribute to variability across tasks and participants in ways that only become evident when systematically reviewing the literature. Differences in probability learning across populations may not underlie or indicate causal differences in more complex cognitive behavior, but rather may themselves be another pattern of behavior that theories of learning and development must account for.

In probability learning experiments, a participant is presented with two alternatives to select, one of which will lead to a reward. For example, in a 70:30 probability learning task, one alternative will lead to a reward on 70% (π) of trials while the other will yield reward on the remaining 30% ($1-\pi$) of trials. Across trials, these two choices are temporally independent, meaning that reinforcement on one trial does not predict reinforcement on subsequent trials. The choices are also coupled, meaning that one and only one alternative will potentially lead to a reward on any given trial. On each trial, a participant must choose one alternative, typically by pressing one of two buttons or guessing on which side a light will appear, in order to gain a reward. The strategy that would produce the highest reward would be to consistently select the alternative that is associated with the highest probability of reward on all trials. However, surprisingly, a number of studies have suggested that adults do not do this. Instead, after a period of experience with the relative probabilities of the two alternatives, adults are said to “probability match,” meaning that they select each alternative with the relative frequency with which it has been reinforced. In a 70:30 task, this would mean selecting the 70% reinforced alternative 70% of the time and the 30% alternative 30% of the time (e.g., [Grant, Hake, & Hornseth, 1951](#); [Estes & Straughan, 1954](#); [Koehler & James, 2010](#); [Neimark & Shuford, 1959](#); see [Edwards, 1956](#), [Myers, 1976](#) and [Vulkan, 2000](#) for reviews). Children, on the other hand, are said to “maximize,” meaning they always select the alternative with the higher probability of producing a reward ([Goldman & Denny, 1963](#); [Jones & Liverant, 1960](#); [Stevenson & Weir, 1959](#)). Researchers have long been interested in why participants, especially adults, seem to make non-optimal choices in probability learning tasks. If true, what might underlie this non-optimal behavior and what are the implications of behavior on probability learning asks for other aspects of cognition?

In this review, I will re-examine the evidence that adults adopt a non-optimal strategy in probability learning tasks, overviewing previous research on probability learning by nonhuman animals, as well as by human adults and children. I will then discuss how probability learning behavior has been applied to other cognitive development literatures. To briefly foreshadow the results, on many occasions probability learning behavior appears to have been misinterpreted or misstated, with consequences for how we interpret patterns of behavior on probability learning tasks and, thus, for the insights we gain into language learning processes that derive from probability learning behavior.

I will first review the non-human animal probability learning literature, which, more than any other literature, argues against simplistic explanations of developmental differences or cross-species differences on probability learning tasks. The animal literature highlights the importance of task demands that serve as better explanations for many differences in behavior. Strikingly, the conclusion from this research is that probability matching behavior, and in fact any less-rational behavior that deviates from maximizing, is likely accounted for by differences in task demands. To the extent that there is variability across species or tasks and animals appear to be doing something more like probability matching, this behavior can be attributed to task demands or experimental design, such as the clarity of the boundaries between trials, and to differences in learning abilities across species. The apparent lack of species-variant learning mechanisms (as was once proposed) provides no support for probability matching.

Next, I will review literature, primarily from the 1950s to 1970s pertaining to probability learning in adults and children. The animal literature, which emphasizes the importance of task demands, may aid in the interpretation of the human adult data, where, once again, a striking observation is the relevance of task design for data interpretation. Surprisingly, and contrary to popular belief, there is little evidence for probability matching behavior in adults (for additional discussion of this claim, see [Edwards, 1956](#), [Myers, 1976](#) and [Vulkan, 2000](#)). The departures that have been observed from maximizing behavior can be attributed to either task manipulations or learned biases that adults bring to probability learning tasks. In fact, studies that directly manipulate task design or instructions across participant groups comprise a sizable portion of the probability learning literature, and these comparative studies strongly suggest task behavior is extraordinarily affected by experimental design. The conclusion from a thorough review of this research is that trying to describe a qualitative difference

between probability matching and maximizing is not an appropriate characterization of the behavior. Instead, there are a number of different reasons people and animals will fail to maximize. Understanding the factors that push participants toward or away from true maximizing behavior, and the consequences for meta-task awareness is the true lesson of this literature.

In the final section, I address how the probability learning literature has been applied to language learning and other cognitive domains. For example, a child's ability to easily learn their native language or to regularize probabilistic language input (creolization), where adults fail to do so, is often attributed to children's tendency to maximize rather than probability match in probability learning tasks. I will discuss how the interpretation of the developmental trajectories of some behaviors can and cannot be aided by knowledge of the probability learning literature. Finally, I discuss the role of literature reviews more broadly in the practice of science, and key insights that can come from synthesizing existing literatures.



1. Non-human animal probability learning

Many findings in the animal probability learning literature are directly relevant to the study of human probability learning and its applications to other aspects of human cognition. First, an obvious lesson from the animal learning literature is the importance of task design. Second, current theories that posit that adult and child learning mechanisms are categorically different are reminiscent of many of the debates that took place in the animal literature, regarding differences in learning behavior across species and whether these cross-species differences could be caused by categorically different learning mechanisms. There is a striking overlap in the debates that took place between researchers within the non-human animal literature and within the human probability learning literature.

1.1 Typical task design

In probability learning tasks with animals, animals are typically first trained to press a button or lever, or to touch a target to receive a reward. Once this behavior has been learned, the probability learning task can begin. On a typical trial, the animal chooses one of two targets. If the animal is correct, it receives some sort of food reward. However, if the animal is incorrect, different studies provide one of three possible outcomes: guidance, correction and non-correction.

In a probability learning task with "guidance," after an incorrect response, the animal's attention is drawn to the correct target. In some

instances, the correct alternative may remain illuminated while the incorrect alternative is darkened. In other cases, the incorrect target may be removed while the correct target remains. The trial ends when the animal selects that target and is reinforced. The second method involves “correction.” In this procedure, after an incorrect response, typically nothing happens, and the animal continues to make target selections until the correct alternative is selected. As with guidance, the trial ends when the animal is reinforced. The goal of the guidance and correction procedures to ensure that the probability of reinforcement across the experiment is exactly the experimentally determined proportion. For example, in a 70:30 probability learning task, one side is indeed reinforced on 70% of trials and the other on 30% of trials. The third procedure is the “non-correction.” This is the procedure most commonly used with human participants. In this task, after an incorrect response no reinforcement is given. Whatever event marks the inter-trial interval (such as a certain elapse of time or a period of darkness) immediately occurs. Under this technique, exact probability of reinforcement may vary from the experimentally manipulated proportion. Non-correction paradigms were often considered “easier” because they tended to promote more maximizing behavior, but researchers often had reasons for either preferring to keep actual reinforcement rates consistent with the target proportions, or, in many cases, to compare and contrast behavior as a consequence of which paradigm was employed.

Another aspect that varied across studies is whether the probability learning task targets are distinguished spatially or visually. In spatial tasks, the animal must learn reinforcement probabilities associated with the spatial orientation of targets: right and left. In visual tasks, the animal must learn reinforcement probabilities associated with targets that differ visually in some way—for example, color or pattern—but that can appear on either the right or left side of the testing apparatus. The type of reinforcement procedure used (guidance, correction or non-correction) and the type of problem the animal had to learn (spatial or visual) significantly affect task difficulty and thus behavior. The experimental characteristics of the animal studies reviewed here are included in [Table 1](#).

1.2 Findings

Early probability learning researchers emphasized cross-species differences in task behavior. This led to the hypothesis, proposed by [Bitterman \(1965\)](#), that evolutionarily driven species differences led to different learning mechanisms across species. Probability learning experiments tended to show that

Table 1 Animal probability learning studies.

Author year	Animal	Ratio	Discrimination	Reinforcement type	Number of trials	sResults
Lauer and Estes (1954)	Rat	75:25	Spatial (T-maze)	Correction	56	Matched
Bitterman, Wodinsky, and Candland (1958)	Fish	70:30	Visual	Guidance, no guidance	500	Matched (guidance) Maximized (no guidance)
Bitterman et al. (1958)	Rat	70:30	Visual	Guidance	500	Maximized
Wilson and Rollin (1959)	Monkey	75:25	Spatial	No correction, correction	1920 (correction) 384 (no correction)	Maximized (more slowly with correction)
Wilson (1960)	Monkey	75:25	Spatial	Correction, guidance	1024 (correction) 768 (guidance)	Maximized (more slowly with correction)
Meyer (1960)	Monkey	100:0 to 50:50, in increments of 10 (between Ss)	Visual	No correction, then “implicit correction” (only see reward, cannot select it)	500	Maximized, even in post-experiment 55:45 reinforcement task
Behrend and Bitterman (1961)	Fish	70:30 → 0:100 → 20:80 or 40:60 → 50:50 → 100:0 or 70:30	Visual	Guidance, then no guidance (final 100:0 or 70:30 block)	500 (70:30), 420 (20:80 and 40:60) 340 (50:50) 420 (100:0 and 70:30)	Matched (guidance) maximized (no guidance)
Bullock and Bitterman (1962)	Pigeon	80:20, 70:30, 60:40	Visual	Guidance	920	Matched
Uhl (1963)	Rat	60:40, 70:30, 80:20, 90:10	Spatial	Correction, non-correction	1000	Close to maximizing, correction learned more slowly. No effect of incentive (sucrose concentration) in reward

Graf, Bullock, and Bitterman (1964)	Pigeon	70:30	Spatial, visual	Guidance or no guidance, correction or no correction	2400 (first press a center key (CK)) 1760 (no center key) but variable length blocks of reinforcement type	<i>Visual-No CK</i> : Correction and guidance: matched; non-correction: maximized <i>CK</i> : Correction: matched; guidance and Non-correction: maximized <i>Spatial</i> -(center key, correction) mostly maximized (some variation across birds)
Wilson, Oscar, and Bitterman (1964)	Monkey	70:30 → 50:50 → 30::70 → 40:60 or 60:40	Spatial	Correction	1800 (70:30) 1200 (50:50) 2000 (30:70) 1400 (40:60 or 60:40)	Matched
Longo (1964)	Cockroach	100:0 or 70:30 → 0:100 → 40:60 or 20:80 → 50:50	Spatial	Correction (Y maze, exit to escape shock)	260 (100:0 or 70:30) 140 (0:100) 160 (40:60 or 20:80) 180 (50:50)	Matched
Kirk and Bitterman (1965)	Turtle	70:30	Spatial, visual	Correction	400	Matched (visual) mixed performance (spatial)
Warren and Beck (1966)	Cat	50:50 to 100:0 by 10, within Ss	Visual	No correction, then “implicit correction” (only see reward)	500 trials each ratio	Intermediate, because individual differences and tendency to perseverate on correct alternative
Shimp (1966)	Pigeon	25:75 (red) and 75:25 (green)	Spatial	No correction	20,000	Maximized (some perseveration)
Wright (1967)	Rat	70:30	Spatial (T-maze)	Correction	110	Maximized. Faster with greater reward, no effect of drive (degree of food deprivation)
Treichler (1967)	Monkey	80:20 or 60:40	Visual	No correction	880	Maximized. Faster with 80:20, no effect of food type
Weitzman (1967)	Rat	70:30, 70:30-variant, designed to be less like the fish task	Spatial	Guidance	400	Matched (70:30) maximized (70:30-Variant)

Table 1 Animal probability learning studies.—cont'd

Author year	Animal	Ratio	Discrimination	Reinforcement type	Number of trials	sResults
Weitzman (1967)	Fish	70:30	Spatial	Guidance	400	Matched
Calfee (1968)	Rat	80:20 or 65:35	Spatial	Guidance	3680	Maximized, though some animals retained a preference to the incorrect side
Schweikert and Treichler (1969)	Cat	90:10, 80:20, 70:30 or 60:40	Visual	No correction	500, except 60:40, 1000	Maximized, more slowly than monkeys and more slowly in 70:30 and 60:40
Topping and Parker (1970)	Pigeon	70:30	Spatial	No correction	1200	Maximized. Longer delays or variable delays between selection and outcome are slower to asymptote
Shimp (1970)	Rat	70:30	Spatial	Correction	12,500	Maximized (looks like matching earlier in training)
Bitterman (1971)	Rat	100:0 → 70:30, only 70:30	Visual	Guidance	500 (switch group 340, because 160,100:0 trials)	Matched, but 70:30-switch group maintained a higher proportion of high-probability choices than 70:30 group
Williams and Albiniak (1972)	Crocodile	75:25 → 100:0 → 87.5:12.5 → 12.5:87.5	Spatial (T-maze)	Correction	120 (75:25) 72 (100:0) 72 (87.5:12.5) 144 (12.5:87.5)	Overshot matching, because a tendency to re-pick the low-frequency alternative when it had just been reinforced
Fischer (1972)	Chick	70:30	Visual, spatial (T-maze)	Correction, no correction and limited correction (finite number of choices/trial)	As many as they needed to maximize (varied by condition 100–700)	Maximized. Unlimited correction increased trials to maximize

in probability learning tasks without correction, mammals, including rats (Bitterman et al., 1958) and monkeys (Wilson, 1960; Wilson & Rollin, 1959) tended to maximize while non-mammalian species, such as fish (Bitterman et al., 1958) and pigeons (Bullock & Bitterman, 1962) tended to probability match. Thus, more “intelligent”^a species, like mammals, tended to behave optimally by maximizing, while less intelligent species behaved less optimally, and probability matched. This position is further supported by studies of decorticated rats, who probability match, like less intelligent species do (Gonzalez, Roberts, & Bitterman, 1964).

This evolution-based perspective that aimed to understand qualitative differences in learning across species encouraged researchers to investigate a wider range of species in probability learning tasks. If learning mechanisms are evolutionarily determined, then animals that are evolutionarily intermediate between the more typically studied fish and pigeons, on the one hand, and mammals such as rats and monkeys, on the other, should show learning mechanisms that yield behavior between that of the more typically studied animals. Indeed, studies with cockroaches have shown matching behavior (Longo, 1964) and studies with turtles (Kirk & Bitterman, 1965) and crocodiles (Williams & Albiniak, 1972) yielded intermediate data—neither matching nor maximizing. These studies of many different animals were generally interpreted as supporting the hypothesis that animals employ qualitatively different learning mechanisms as a consequence of evolution—specifically, the length of time since sharing a common ancestor with humans. According to this evolutionary approach, different behavior between mammals and other species was a consequence of different species employing *categorically* different learning mechanisms.

Other researchers disagreed that the species differences were best attributed to categorically different learning processes. Instead, a counter-theory argued that the differences in behavior could be attributed to different learning rates (not qualitatively different learning mechanisms) or differences in task design (Fischer, 1972; Shimp, 1966; Topping & Parker, 1970). The counter-theory posited that all animals should maximize, but the less intelligent animals might take more time to do so. On this view, all animals also should use similar learning mechanisms but might have different learning rates. If learning mechanisms are similar across species, differences in task

^a “Intelligence” was the term used in the literature at that time and was a common way to talk about differences across species. Now, because different animals might be better at different tasks, there is little agreement about what intelligence might mean in the context of cross-species differences. Today, researchers might use terms like “cognitive flexibility” or “learning rate” of a particular task to describe the specific differences between species this literature referred to.

or training procedure might better account for the pattern of maximizing and probability matching behavior observed in the literature.

Researchers therefore explicitly investigated the role of task design and number of training trials in the observed behavior across species. In all probability learning studies, animals generally began all tasks selecting each alternative at chance levels and shifted toward selecting the higher probability of reward alternative more often. A major finding was that the matching behavior, which was thought to be characteristic of some animals including pigeons, eventually came to look like maximizing when the task was prolonged (Shimp, 1966; Topping & Parker, 1970). While most studies involved only hundreds or occasionally a low-thousands number of trials, Shimp (1966) observed 20,000 trials in a variant of a typical probability matching task. In his procedure, when keys were lit red, the reinforcement probabilities of the right and left keys were 25:75 and were reversed (75:25) when the keys were green. Pigeons maximized within 20,000 trials, but had the experiment consisted of fewer trials, the behavior might have been described as probability matching. Behavior that appeared to be probability matching at first eventually came to look like maximizing as the animal gained more familiarity with the task. Again, by increasing the number of trials, Shimp (1970) also found maximizing in a correction task with rats, where maximizing was previously thought to not occur. Rats maximized in 12,500 trials in a task that had previously yielded matching with 56 trials (Lauer & Estes, 1954) or only an overshooting of matching with 1000 trials (Uhl, 1963). These findings disputed the notion of categorically different learning mechanisms across species and pointed to a feature of task design that affected behavior: number of trials. Perhaps all animals would maximize if given enough training, and studies that found probability matching simply stopped the study too soon. The finding that extending training leads to more maximizing behavior does not support the notion of separate learning mechanisms: one that asymptotes in maximizing and one in probability matching. Rather, they support a single learning mechanism that ultimately results in maximizing, plus various attributes of different animals (such as learning rate) that cause these animals to reach maximizing behavior sooner or later.

Systematic investigations also found that the type of reinforcement procedure (guidance, correction or non-correction) affected behavior. Specifically, learning rates tended to higher, and maximizing was reached more quickly, with the non-correction procedure. Though many researchers do not explicitly address the effect of task, this pattern is evident in the data.

For example, in a study with pigeons, using guidance (a harder task) and only 920 trials, [Bullock and Bitterman \(1962\)](#) found matching. In a similar task but using non-correction, [Graf et al. \(1964\)](#) and [Topping and Parker \(1970\)](#), with 1760 and 1200 trials respectively, found maximizing. As the task got harder (guidance vs. non-correction), more trials were necessary to attain maximizing. Many other findings are consistent with the idea that task substantially contributes to behavior patterns ([Bitterman et al., 1958](#); [Fischer, 1972](#); [Uhl, 1963](#); [Wilson, 1960](#); [Wilson & Rollin, 1959](#)). One possible explanation, articulated by [Fischer \(1972\)](#), is that trial boundaries are less clear in the guidance and correction procedures. As a consequence, it may be less clear to the animal what statistics they should compute or track. The clearer trial boundary of the non-correction procedure may facilitate learning by making the relative reinforcement probabilities clearer to the animal. Therefore, because reinforcement probabilities are easier to learn in non-correction tasks, animals in such tasks should be more likely to maximize.

Across many studies with non-human animals, a clear theme emerged that anything that made the probability learning task more difficult would lead to a slower learning rate and a slower shift to maximizing behavior. If probability learning is more difficult for fish and pigeons relative to the “smarter” rats and monkeys, it would make sense that the fish and pigeons need more trials to maximize, which would lead to the observed patterns of probability matching with an insufficient number of experiment trials. Further, if a guidance or correction reinforcement procedure made the probability of reinforcement more difficult to learn than a non-correction procedure, researchers should predict faster learning and more observed maximizing with the non-correction procedure. Further supporting this idea, animals participating in visual tasks (color or pattern dependencies) tended not to maximize or maximize as quickly as in spatial tasks (left/right dependencies), which is consistent with the fact that a spatial distinction is easier for an animal to learn than a visual one ([Graf et al., 1964](#); [Kirk & Bitterman, 1965](#)). A clear pattern emerges: the more difficulty a task was for an animal—either because of the task design or because of the learning rate of the animal, the more trials the animal required to attain maximizing behavior.

Task difficulty alone cannot explain all non-maximizing behavior, but in conjunction with a few other phenomena, the appearance of matching behavior seems to disappear. One well-observed phenomenon is that in many of these studies, there was a tendency for animals to perseverate.

Often after a correct response, animals tended to choose that target again (Shimp, 1966; Warren & Beck, 1966). This perseveration increased the time necessary to reach maximizing behavior and may account for some deviance from true maximizing behavior. Further, occasionally, some animals would form a preference for a particular side and continue to choose that side, or alternately gain a preference for switching and alternate between trials regardless of reinforcement proportion (Calfee, 1968). This perseveration or motor repetition behavior can explain remaining deviations from maximizing behavior after task difficulty and number of trials is considered. Crucially, these deviations from maximizing behavior seem to stem from task-specific or response-specific behaviors, not from a probability matching behavior, *per se*.


Other subtle features of the tasks used to assess probability learning behavior may account for additional deviation from maximizing. One such phenomenon is the tendency for what may appear to be matching behavior to actually be a product of attending to more local patterns, often called “momentary maximizing” (Shimp, 1966). Given outcomes on previous trials, occasionally the lower frequency alternative was, in fact, more likely to yield reinforcement when considering the subtle non-independence of trials in some paradigms. Many animals were able to learn these subtle local patterns. In deviations from typical probability matching tasks when there are indeed subtle dependencies that can be learned in order to predict with greater accuracy which alternative will yield a reward, animals seem to be able to learn them (Shimp, 1966, 1967). For example, in a complex task with variable ratio reinforcement coupled with a probability learning task (not every trial gave the animal an opportunity for reinforcement) pigeons seemed to track outcomes on prior trials and exhibited matching behavior that was a consequence of maximizing given the outcome on prior trials (Shimp, 1966). Subsequent studies with variable ratio or variable interval reinforcement with pigeons (Hinson & Staddon, 1983) or reinforcement that is contingent on the previously selected target with pigeons (Fetterman & Stubbs, 1982) and rats (Hiraoka, 1984), suggest that animals are capable of tracking outcomes on prior trials and applying that information to future behavior. Animals indeed seem to be able to perform these tasks in a very sophisticated way that optimized their reward, even when to do so they must track multiple patterns at once. Animals are clearly able to learn subtle patterns across experiment trials—whether they arise by design or incidentally.

To translate this sophisticated learning back to probability learning tasks, where local dependencies theoretically should not exist, subtle local cues nevertheless appeared in some classic probability learning designs. In many probability learning studies, targets are not determined trial-to-trial, but rather the sequence of reinforced alternatives is determined prior to the study. Experimenters ensured that blocks of a certain size maintained the target reinforcement ratio. As a consequence, in some experiments, trials are not completely independent. In some designs, a 70:30 reinforcement probability will be implemented by having 7 out of every 10 successive trials being reinforced (Weitzman, 1967) or ensure that there are never more than five high-frequency (rewarded) events in a row (Bitterman et al., 1958; Graf et al., 1964). As a consequence of these design choices, the probability of reinforcement on the low frequency target actually increases with each successive reinforcement of the high-probability target. Thus, the probability of the low frequency target, in the context of previous trials, may become the alternative more likely to lead to reward in some contexts. Animals seem to take advantage of these subtle context patterns. At least in some cases, when animals seem to be probability matching, they may instead be using a momentary maximizing strategy—indeed maximizing, but over more complex patterns than the researchers intended. The tendency of animals to momentarily maximize suggests sophisticated learning mechanisms, which should be considered when interpreting human data. This tendency to learn subtle patterns—or believe that there is a subtle pattern to learn—will be a major feature of the discussion of human probability learning behavior.

1.3 Summary

Though some researchers hypothesized qualitative species differences in behavior on probability learning tasks, subsequent findings challenged this hypothesis. Instead, there seems to be a common learning mechanism with quantitative differences in learning rate across species, and with task characteristics playing an important role in the resulting behavior. Critically, all animals seem to be using a strategy that eventually will yield maximizing. The fact that all non-human animals in these studies share a common learning mechanism on these tasks should inform our understanding of human behavior on these tasks. When the human (adult or child) behavior diverges from non-human animal behavior on these tasks, it is important to consider

whether such differences are more likely to be due to a qualitatively different learning mechanism or, instead, if they too can be explained by task characteristics or other types of knowledge that humans might bring to the task.



2. Human probability learning: Adults

A commonly held view is that in a probability learning task, adults usually probability match rather than maximize. Implicit in the assumption that behavior on a probability learning task can either be described as probability matching or maximizing is that there are two separate learning mechanisms or discrete behavioral patterns that asymptote at different levels, either 1 (maximizing) or π (probability of reward). Probability matching is thought to persist in adults despite the fact that this strategy leads to a lower rate of reward or correct responses relative to a maximizing strategy. So why might adults adopt this less optimal strategy?

What a review of the adult probability learning literature shows is that, similar to the case of non-human animal probability learning just discussed, the evidence for probability matching in human adults is far less straight-forward than it might initially seem. This is not to say that in some probability learning tasks adults do not exhibit response rates at or near π for some period of time. Rather, there is very little data to support the idea that there exists a learning mechanism or pattern of behavior that *asymptotes* at π over extended training. Instead, human behavior on probability learning tasks can better be described as a single learning mechanism or behavior that asymptotes at maximizing, along with various task effects and effects of prior knowledge that push humans toward π . These factors are described in the following sections and include experimental design, task instructions, task difficulty and various experience-based biases. Some of these factors lead to the appearance of matching-like behavior. But because the observed matching-like behavior derives from a very large list of factors, it may not make sense to treat “matching” as a categorically different learning strategy.

2.1 Typical task designs

In typically probability learning tasks (sometimes called binary choice tasks) with adults, participants predict which of two or occasionally three alternatives will occur next (Cotton & Rechtschaffen, 1958; Gardner, 1957, 1958). In some tasks, participants guess which symbol, shape or word will appear next on a card or computer screen (Edwards, 1961;

Gaissmaier & Schooler, 2008; Neimark & Shuford, 1959; Ross, 1954; Rubenstein, 1959; Singer, 1967, 1968), or will be read aloud by an experimenter (Jarvik, 1951; McCracken, Osterhout, & Voss, 1962). In other designs, participants might guess one which side a light will appear (Braveman & Fischer, 1968; Estes & Straughan, 1954; Friedman, Padilla, & Gelfand, 1964; Myers, Fort, Katz, & Suydam, 1963; Siegel & Goldstein, 1959; Suppes & Atkinson, 1960), or what color (Jones, 1961) or pattern (Hake & Hyman, 1953) of lights will turn on, or whether a light (or second light) will turn on at all (Grant et al., 1951; Grant, Hornseth, & Hake, 1950; Humphreys, 1939). Occasionally a participant might guess which lever will turn off a buzzer (Arima, 1965) or lead or a reward on a slot machine (Edwards, 1956; Goodnow & Pettigrew, 1955). There was occasionally a monetary or other award associated with correctly predicting the trial outcome (Edwards, 1956; Goodnow & Pettigrew, 1955; Peterson & Ulehla, 1965). Occasionally, studies manipulated the amount of monetary reward or risk associated with correct or incorrect guesses (Myers et al., 1963; Siegel & Goldstein, 1959; Suppes & Atkinson, 1960). However, in the majority of studies reviewed here, “being correct” is the only reward on each trial.

2.2 Findings

Contrary to the commonly-held view, there are surprisingly few studies that show true matching behavior, asymptoting at π . Though a handful of studies do indeed show true matching behavior (Neimark & Shuford, 1959), they are rare and often only show true matching behavior in a single experiment or experimental condition out of many runs. Instead, participants tend to overshoot π (for example, responding 75–85% to the 70% alternative in a 70:30 experiment (e.g., Edwards, 1961; Estes & Johns, 1958; Estes & Straughan, 1954; Friedman et al., 1964; Grant et al., 1951; Hake & Hyman, 1953; Jarvik, 1951; Myers et al., 1963; Unturbe & Corominas, 2007; see Vulkan, 2000 for an extended discussion of this pattern). In fact, this tendency for studies of adult probability learning to show behavior that overshoots true matching behavior has been a major finding of a number of reviews of adult probability learning (Edwards, 1956; Estes, 1964; Myers, 1976; Vulkan, 2000). An important initial observation about the adult probability learning literature is that what is often described as probability matching is more accurately overshooting true probability matching. For a table much like Table 1, but summarizing adult probability learning studies, see Vulkan (2000).

A number of task characteristics were consistently associated with a greater degree of overshooting probability matching. In fact, there is evidence that similar task factors affect matching versus maximizing behavior in probability in both human and non-human primates (Saldana, Claidière, Fagot, & Smith, 2020). Many studies with human participants explicitly tested the effects of experimental instructions or task design on matching versus maximizing behavior. For example, a greater monetary reward or monetary risk was consistently associated with behavior that more closely approached maximizing (Myers et al., 1963; Siegel & Goldstein, 1959; Suppes & Atkinson, 1960). I will discuss additional examples below. Of course, overshooting probability matching is not the same as maximizing. I focus this review on understanding the processes that can account for a behavior that is somewhere between 1 and π . I propose that the processes that underlie behavior can be best described as variants on maximizing, and the most fruitful approach is to understand the task and individual factors that prevent participants from exhibiting true maximizing behavior. Below I describe four factors that may prevent adults from true maximizing behavior on probability learning tasks.

One factor is experiment length. Like non-human animals, adult human participants start from a strategy of random guessing (before they understand anything about the nature of the task) and then slowly approach maximizing behavior. An experimenter sampling behavior for a short period of time (i.e., stopping before maximizing is reached) is quite likely to observe a participant's response rate close to π , or somewhere between π and 1. This tendency is visualized in Fig. 1. Stopping data collection in zone




Fig. 1 A visualization of learning trajectories in probability learning experiments.

A would yield behavior that would be described as undershooting matching, stopping data collection in zone B would yield matching behavior and stopping data collection in zone C would yield maximizing behavior. A single learning behavior, assessed at different points in time, would be categorized as qualitatively different profiles of behavior. This explanation is supported by the observation that the tendency to overshoot π becomes particularly pronounced after extended numbers of training trials (Shanks, Tunney, & McCarthy, 2002). Previous reviews of the probability learning literature have come to a similar conclusion with respect to the number of training trials. Edwards (1956), Estes (1964), Myers (1976) and Vulkan (2000) all point out that with extended training, participants tend to overshoot π . This is a point of consistency between the non-human animal and human adult literatures: given more trials, participants will approach maximizing behavior.

Another factor that may lead participants to respond at a rate between π and 1 is task difficulty. Experimenters often approach probability learning tasks with the assumption that participants will quickly learn the parameters of the task. They often overlook the extent to which a participant has to learn about the task in order to produce an optimal behavior. For example, Green, Benson, Kersten, and Schrater (2010) note that participants have to learn a number of aspects of the experimental design that are not obvious before they can choose how to behave optimally. First, participants must learn that *one alternative is more likely to lead to reinforcement than the other*. While this may seem trivial at first, consider that in a 70:30 reinforcement task, after 20 trials, about 5% of the time, the lower-probability may be reinforced at least equally often as the high-probability alternative, and nearly a quarter of the time the higher-probability alternative will be reinforced with a 12-to-8 or less extreme split. Participants may not be able to conclude that the higher-probability alternative is indeed higher probability with this data—even assuming perfect memory for past trials (Myers, 1976). Even this simple fact about the task, that one alternative will more likely lead to success than the other, may actually take upwards of 40 trials to discover, depending on reinforcement probabilities and other aspects of the task. Moreover, even when it becomes clear that one alternative is linked to a higher success rate, the participant *still* cannot be sure about the demands of the task. Participants must also learn that *the reinforcement of the two alternatives is independent*, so the “correct” alternative on one trial has no bearing on which will be correct on the next. As we will see below, this task characteristic seems particularly difficult for adults to discover. Finally, participants must learn that *the outcomes are coupled*, or that the reinforcement

probabilities of the two alternatives are complementary and add to 1. [Green et al. \(2010\)](#) emphasize that these parameters are not necessarily the most plausible in the mind of the participant at the beginning of the task, so there must be an appreciation for the length of training required for participants to realize that reinforcement of the alternatives are unequal, temporally independent and coupled.

A third factor that leads participants to respond at rates below maximizing is that experimental designs sometimes actually *violate* trial independence. In many studies, as was true in the non-human animal literature, each trial is not randomly generated independent of prior trials. Rather, a random number sequence is generated, and trial outcomes are based on that sequence, with important caveats. First, the total set of all trials must conform to the reinforcement probabilities, so the designated “correct” alternative is not independently generated for each trial. Second, in many cases, blocks of a certain number of trials must also conform to the target probabilities. For example, in many studies 8 out of every 10 or even 4 out of every 5 trials must reinforce the high-probability alternative (e.g., [Arima, 1965](#); [Jarvik, 1951](#); [Ross, 1954](#)). As noted by [Vulkan \(2000\)](#), the problem with placing these restrictions on task sequences is that sequences are no longer random. Under these circumstances, the gambler’s fallacy (believing that a long string of occurrences of one event makes the other event more likely) is no longer a fallacy, but a truth of the experimental design. After a run of a certain length, it becomes rational to choose the low-frequency alternative, which may account for some under-maximizing behavior. Similar to the animals in the previously discussed studies, adults may be “momentarily maximizing” when they fail to maximize overall.

A fourth and final factor that may lead participants away from true maximizing is the prior experience that adult participants bring to the probability learning task. This prior experience may take a number of forms, but most relevant here are experiences that adults might have had that dissuade them from believing that trials are truly independent. Adults’ experience might suggest that events are not random, but that the outcome of one trial may be informative of the outcome on the next. As [Ayton and Fischer \(2004, p. 1369\)](#) describe, “outside of gambling casinos and psychology laboratories there are very few—if any—circumstances where one can safely assume conditional independence of a succession of events.” This lack of experience is consistent with observed patterns of behavior. Adults are inaccurate at generating or even identifying random sequences ([Lopes, 1982](#)). When asked to generate random sequences, adults tend to produce

more alternations and shorter runs than would be expected in a truly random sequence (Bar-Hillel & Wagenaar, 1991; Wagenaar, 1972), and they judge sequences with more alternations than would be expected in a truly random sequence as being more random-like than actual random sequences (Ayton & Fischer, 2004; Falk & Konold, 1997). If adults have little experience with random, independent events, enter the lab assuming that events are unlikely to be independent, and do not find the experimental sequences convincing evidence of independent events, it would be difficult for adults to learn that trials are actually independent in the experimental task and thus to shift their behavior in accordance with this information.

Why might the belief that trials are not independent lead to behavior between matching and maximizing? One possible explanation is that participants are searching for the underlying structure—the pattern that predicts which alternative is “correct.” Gaissmaier and Schooler (2008) found that when there was a pattern governing which alternative was correct, participants who exhibited an exploratory “trial and error” pattern of behavior, alternating between both alternatives, tended to find the pattern more often than participants who exhibited behavior closer to maximizing. Perhaps when adults exhibit matching-like behavior, it is not because they are employing matching as an actual (irrational) strategy, but because they believe there is an underlying pattern to the sequence and are actively searching for it (Anderson, 1960; Gaissmaier & Schooler, 2008; Hake & Hyman, 1953; Wolford, Newman, Miller, & Wig, 2004). Consistent with this assessment, in a review of probability learning, Vulkan (2000) speculates an adaptive motivation behind probability matching and comes to a similar conclusion. In contrast to many assumptions that probability matching is a sub-optimal behavior that reflects flaws or shortcomings in adult learning (Koehler & James, 2009; Otto, Taylor, & Markman, 2011; West & Stanovich, 2003) or “underthinking” task demands (Koehler & James, 2010), Vulkan suggests that switching between the two alternatives can facilitate finding a pattern in the sequence. If participants believe there is a pattern to be found, maximizing may not be the best strategy to find that pattern. Such a behavior, while pushing a participant away from maximizing and toward π , is still not properly described as probability matching.

If non-maximizing facilitates a search for a pattern, then when participants can be convinced that the probabilities of reinforcement are truly random, they should show more maximizing behavior. Many studies explicitly tested this hypothesis and assessed multiple strategies for convincing

participants that trials really are independent. One simple way to convince participants that trials are independent may be to extend the number of trials in an experiment. Typically, probability learning tasks only consist of about 200 trials, but [Edwards \(1961\)](#) extended the task to 1000 trials and found that later in the experiment participants exhibited maximizing behavior. As the experiment continues and the participant still cannot find an underlying pattern, they may “give up” or become convinced there is no pattern and begin maximizing.

Another way to convince participants of the independence of trials is to vary the task design to demonstrate to the participant that the reinforced alternative is truly random. In studies where participants are explicitly shown that the reinforcement sequence is random, they tend to exhibit more maximizing behavior. Participants showed more maximizing behavior when participants rolled a fair die to determine reinforcement ([Newell & Rakow, 2007](#)), or when they watched the experimenter shuffle the cards that predicted reinforcement compared to when the cards were randomized out of their view ([Beach & Swensson, 1967](#); [Rubenstein, 1959](#)). In a clear demonstration of this phenomenon, [Peterson and Ulehla \(1965\)](#) found more maximizing behavior when participants rolled a die to determine reinforcement, than when participants guessed the color of a card in a deck, presumably because participants could be more confident that the outcome of the dice roll was truly random. These simple methodological tweaks had meaningful effects on participants’ behavior, suggesting that meta-task factors are clearly influencing participants’ choices.

In another clear demonstration of the role of task design on participant behavior, [Morse and Runquist \(1960\)](#) explicitly manipulated task design. When participants performed two tasks with *identical sequences* of reinforcement, they maximized more often in the task in which the sequence was demonstrated to be random. [Morse and Runquist \(1960\)](#) had participants perform two tasks. In the first probability learning task, participants dropped an aluminum rod onto the floor and guessed whether it could come to rest touching one of many parallel lines on the floor (this occurred about 35% of the time). The outcome (touching or not touching) was recorded and, unbeknownst to the participants, that exact sequence was used in a subsequent light prediction task, in which participants predicted whether the left or right light will turn on. More maximizing was observed in the rod-dropping task than in the light prediction task. In the rod-dropping task, it was evident to the participants that trials were random and independent, while this was not clear in the light prediction task. Again, features of the

task being employed had a clear effect on behavior, suggesting that it is not simply the reinforcement sequences, but also features of the task and participants' beliefs about the task, that affected behavior.

Simple manipulation of instructions can also encourage or discourage belief that sequences are random and promote more maximizing behavior. Like manipulating task design, different instructions might lead participants toward or away from maximizing behavior. Supporting the idea that adults enter the lab with a bias that sequences are not random, researchers found it easier to convince participants that a sequence is structured rather than random (Hyman & Jenkin, 1956). To counter this bias, many studies emphasized to participants that the sequences were random, which tended to increase maximizing behavior (Braveman & Fischer, 1968; McCracken et al., 1962; Shanks et al., 2002). Similarly, participants tended to exhibit more maximizing behavior when they were told that the task was a gambling task rather than a problem-solving task (Goodnow, 1955; Goodnow & Postman, 1955). Giving participants additional task information, such as explicitly telling participants the probabilities of the two alternatives or asking the participants to recommend a strategy to another person (Fantino & Esfandiari, 2002) or asking participants to explicitly evaluate a matching versus maximizing strategy (Koehler & James, 2010) led to more maximizing behavior. In fact, even participants who undershot maximizing were able to articulate that maximizing would be the ideal strategy to employ (Koehler & James, 2009, 2010). Adults' beliefs about the nature of the sequence they are encountering has a significant effect on behavior, and these beliefs can be altered through the experimental design and task instructions. Prior experience with the world may cause adults to approach the task with the belief that sequences are not random, but various efforts by the experimenter to convince the participant otherwise seem to lead to more maximizing behavior.

2.3 Summary

The adult probability learning literature suggests that adults do not probability match. That is, there does not seem to be evidence for a type of learning that asymptotes at the probability (π) level. Instead, they tend to maximize, or respond at rates somewhere between maximizing (1) and matching (π). When examining the factors that underlie these deviations from maximizing, we can conclude that the behavior is not well characterized as matching, given the baggage that word implies about irrational

decision processes. Instead, adults' behavior is better characterized as maximizing, but with constraints. The constraints are generally task design or individual factors, including experiences and biases that adults bring to the lab, that lead participants to believe that trials are not truly independent. The characterization of adults' probability learning behavior as maximizing, but with caveats, is important to consider when examining differences between adult and child probability learning, or other sequence learning behavior.



3. Human probability learning: Children

In this review, we saw that animals maximize in probability learning tasks, if given a sufficient number of training trials, and that variations in the species of animal or the experimental design may necessitate more or less training before this maximizing behavior emerges. We further showed that human adults' behavior is also much better characterized as maximizing plus a range of factors that lead to deviations from maximizing, such as exploratory searching for a more effective response strategy. As will be clear in the following section, characterizing the differences between adult and child performance on probability learning tasks is a difficult endeavor, because there is no clear consensus of how children behave in probability learning tasks. Evidence for consistent maximizing or probability matching behavior is generally hard to find, which makes it difficult to conclude exactly how child behavior differs from adult behavior on probability learning tasks.

Conclusions from animal and human adult performance on probability learning tasks may suggest that whatever differences exist between children and adults on probability learning task, they are likely to derive from different task strategies or learned biases that participants bring to the task, not necessarily different learning mechanisms or capabilities, as would be suggested by a dual mechanism, matching versus maximizing, account.

3.1 Typical task design

Some task designs used with children were similar to those used with adults. In some experiments, children guessed which of one or two pictures would appear on the next card in a deck (Offenbach, 1964) or which of two lights will turn on (Craig & Myers, 1963; Derks & Paclisanu, 1967). However, whereas the adult probability learning studies rarely gave participants rewards for correct responses, explicit rewards were much more common with children. Children often received a marble or a token,

which could be exchanged for coins or prizes at the end of the experiment (Goldman & Denny, 1963; Jones & Liverant, 1960; Stevenson & Weir, 1959; Sullivan & Ross, 1970; Weir, 1964; Winefield, 1980).

Another important difference between child and adult experimental designs is that experiments with children tended to consist of fewer trials than experiments with adults. This is an obvious consequence of working with younger participants. Experiments with children rarely consisted of over 100 trials, with the exception of the 200-trial design of Derks and Paclisanu (1967). Adult studies, on the other hand, rarely consisted of fewer than 100 trials, with some studies containing as many as 1000 trials (Edwards, 1961). Considering the importance of extended training in both the animal and adult probability learning literature, it is very likely that the child probability learning data is affected by the limited number of trials afforded by working with younger participants. The fact that studies with children typically employ far fewer trials is an important consideration that is often overlooked when evaluating the child data and making comparisons with adults.

3.2 Findings

There is some evidence that younger children maximize more than older children (Goldman & Denny, 1963; Jones & Liverant, 1960; Stevenson & Weir, 1959) but also some evidence that younger children *under-match* (respond at a rate less than π) more than older children (Craig & Myers, 1963). Other studies find no effect of age (Offenbach, 1964). A somewhat common effect is a U-shaped tendency to maximize, with the youngest (typically nursery school) and oldest children or adults (typically late-teens or undergraduates) exhibiting more maximizing than children between those ages. Weir (1964) first obtained U-shaped learning results in a three-choice probability learning task. The youngest children (ages 3–5) and the oldest children (age 18) tended to maximize more than did the children between those ages. Later, Derks and Paclisanu (1967) replicated this result in a more typical two-choice task, finding that nursery school aged children *overshot* matching, children between nursery school and 3rd grade *undershot* matching and children older than third grade tended to probability match. Additional support for U-shaped matching is found in Sullivan and Ross (1970) and Winefield (1980), making it a somewhat well-documented though hard to explain finding. Perhaps if it were feasible to perform more trials with children, we would begin to see more consistent patterns of behavior.

Another observation of children's probability learning behavior is that they do not tend to demonstrate experience-based biases like adults do. For example, younger children (under about 7 years) do not seem to exhibit the gambler's fallacy or expect the other alternative to be correct after a long run of one alternative (Derks & Paclisanu, 1967). In fact, children sometimes seem to alternate randomly, often regardless of events on prior trials or π , which is not a behavior that adults exhibit (Craig & Myers, 1963; Offenbach, 1964). Children's more seemingly random behavior can have a number of interpretations. Perhaps younger children are less biased about what they consider "random," reflecting different prior experience. Or perhaps children simply fail to remember events on previous trials. Consistent with the hypothesis that younger children come to the task with less bias, younger children who maximize tend to reach that asymptote faster than older children that maximize do (Jones & Liverant, 1960; Weir, 1964; Winefield, 1980). However, when there were sequence probabilities that could be learned, older children learned them faster (Goldman & Denny, 1963; Sullivan & Ross, 1970), and, in cases when patterns did not exist, older children were more likely to report attempting to find one (Craig & Myers, 1963; Offenbach, 1964). This age-related difference suggests children may not bring the same biases about the non-independence of trials to probability learning tasks as adults seem to. If, as it seems, children, particularly older children, are beginning to develop some of the biases and strategies of adults, then investigating differences in these biases may better explain age-related differences in behavior than hypothesized differences in learning mechanisms (matching vs. maximizing).

It is difficult to understand how children approach probability learning tasks and what children believe to be true about the task. For example, relative to adults, how readily do children learn the demands of a probability learning task identified by Green et al., i.e., that one alternative is more frequently reinforced and that events are independent and coupled? It is possible that if children have more difficulty remembering outcomes on previous trials, or integrating that information into a behavior plan, it could take children considerably longer than adults to learn the parameters of the task. It may take children a long time to discover that one alternative leads to a reward more often than the other, and even when they do discover this, they may not understand how to adjust their behavior to take advantage of this fact. There is some evidence that children between the ages of 4 and 11 slowly improve in their ability to select one of two alternatives that is more likely to yield a reward (Falk, Falk, & Levin, 1980), so adjusting one's

behavior to observed probabilities could be a skill that emerges with age. Simply understanding how the task works, and how to adjust one's behavior in response to this information, may itself be a challenge for children.

However, outside of the probability learning literature, there is significant evidence that even infants as young as 8 months have a sophisticated understanding of probabilities. By the end of the first year of life, children are surprised at very low-probability events, such as an experimenter withdrawing many red balls from a container with many white and few red balls (Xu & Denison, 2009; Xu & Garcia, 2008). By 15 months, infants can infer information about samples from facts about populations (Denison & Xu, 2010), even depending on the representativeness of the sample (Gweon, Tenenbaum, & Schulz, 2010). These studies suggest that children and even infants have a sophisticated understanding of probability. So perhaps observed differences between children and adults are not due to a lack of understanding of probability, but an increased difficulty learning which alternative is more often reinforced. Alternately, children may struggle to appropriately use the information they observe about relative probabilities to guide their action. From work in other domains, it is clear that task matters for assessing what children do or do not know; for example, children succeed earlier in false-belief tasks when the task imposes fewer processing demands (Scott & Baillargeon, 2017). Perhaps probability learning tasks are not ideal for assessing what children do or do not know about probability. The locus of the difficulty that children experience in probability learning tasks is unclear and many candidates, that need not be mutually exclusive, exist.

3.3 Summary

Despite the popular belief that children generally maximize in probability learning tasks, reviewing the literature reveals that there is actually limited evidence that children indeed maximize. There is evidence that children and adults may perform differently in probability learning tasks, but it is not clear why these differences exist. Some of these differences may derive from consequences of experimental design, such as the number of trials, or explicit reinforcement, which also tend to differ across the age groups.

A major finding in the adult literature that task design matters and has effects on the biases that adults bring to the tasks, which seem to predict their behavior. It is possible that an investigation of the learned biases and beliefs that participants of different ages bring to the task will end up being the best

way to describe differences in behavior across development. If children approach the task with fewer, or different, biases than adults—specifically that outcomes across trials are not independent—this may predict many of the behavioral differences between adults and children. Of course, children and adults differ in other ways too. Children may be less adept at recalling outcomes on past trials, may not understand trial boundaries as well as adults, or may find the task more difficult for a variety of reasons. It is important to consider which differences between children and adults might reflect general developmental changes, like improved memory abilities, and which differences might reflect task-specific beliefs that can emerge from experience with events and patterns in the world.



4. Implications of probability learning for other domains

The goal of this review of the probability learning literature was to first document and summarize the observed patterns of behavior, and then to understand the reasons for this observed behavior. However, the implications of this review go beyond probability learning tasks. Many other domains refer to the probability learning literature as an underlying distal explanation for observed behavior, especially in complex cognitive tasks. The logic is that the learning mechanisms and processes that underlie behavior on a simple task like probability learning behavior may also underlie other, more complex behaviors.

4.1 Probability learning and language development

An observation commonly made by researchers, parents and anyone who tends to spend a good deal of time around children, is that children learn language differently than adults. There is some empirical support for this observation. Individuals who learn language earlier in life tend to exhibit greater language proficiency (Bialystok & Hakuta, 1999; Johnson & Newport, 1991; Newport, 1990). In addition, speakers who learn a language earlier in life tend to exhibit different patterns of neural activation than speakers who learn a language later in life (Neville & Bruer, 2001). The underlying reasons for these language-learning differences between children and adults are the topic of substantial research and debate.

One potential explanation for child-adult differences in language acquisition is that there is a critical period for language learning. The idea of a language critical period was popularized by Lenneberg (1967), who

argued that language must be learned during this period for a speaker to attain a native-like competency of the language. This critical period is often, but not always, thought of as being language-specific and biologically defined.

An alternative perspective is that the observed differences in language learning in children and adults are not due to a biologically determined language learning window, but are a consequence of language learning itself. These approaches often emphasize the importance of entrenchment: that experience with the statistics of one language makes it harder learn a second language, because that involves ignoring well-practiced statistics (Elman et al., 1996; Plunkett, 1997; Seidenberg & Zevin, 2006). According to these approaches, the accumulation of learning itself is what makes subsequent language learning more difficult later in life.

A third hypothesis is that domain-general differences in learning abilities account for differences between adult and child language learning abilities. One example is the Less-is-More hypothesis (Newport, 1988, 1990), which hypothesizes that children's limited memory capacities aid language acquisition. Younger learners store smaller bits of language, perhaps because of their limited memory capacities. This smaller unit could account for differences in child and adult language learning abilities by simplifying the language to meaning mapping, which is a necessary component of language learning (Newport, 1990; Singleton & Newport, 2004). Approaches that emphasize changes in domain-general learning across the lifespan are often the approaches that refer to probability learning, or other non-language specific learning, as a means of grounding age-related changes in language learning in age-related changes in other learning tasks.

The reason that probability learning is often evoked in discussions of age-related differences in language learning is that some of the age-related language learning behavior of interest involves learning noisy or probabilistic input, much like what participants experience in a probability learning task. A great deal of literature details differences in how children and adults learn patterned input that contains some irregularity or noise. In both naturalistic and laboratory settings, adults and children tend to learn and reproduce these irregularities in different ways, and these differences in behavior are thought to be representative of fundamental differences between the ways in which children and adults learn language.

One naturalistic context in which children systematize probabilistic input is the process of creolization, in which two languages come into contact and form a new natural language that combines the two.

This creolization process is often driven by children. Sankoff and Laberge (1973) noticed that children tended to regularize the usage and stress of a future marker in Tok Pisin, a language of New Guinea. Similar observations have been made in investigations of a newly emerging sign language in Nicaragua, in which Children seem to systematize in ways that deviate from their input, such as systematically using physical space or direction of movement to encode grammatical meaning (Senghas & Coppola, 2001; Senghas, Kita, & Ozyurek, 2004). This process of regularization also occurs when children learn language from immigrant caregivers who may not be fluent speakers of the local language. One investigation found that children of immigrants in Sweden tended to generalize the irregular Swedish input they hear from their parents, but not necessarily in ways that resembled standard Swedish (Kotsinas, 1988). A similar example can be found in the case study of Simon, a child learning American Sign Language from non-natively signing parents. Simon regularized the irregular morpheme use that was present in his parents' speech, to create more consistent signs (Newport, 1999; Singleton & Newport, 2004). In real-life situations in which children and adults learn language from speakers who are sometimes inconsistent, children's behavior is characterized by a regularization of the noisy input, while adults' behavior is not.

Given that these naturalistic situations seemed to point to important differences between language learning in children and adults, they have often been approximated in the laboratory with artificial grammar learning studies. In these studies, children and adults are typically presented with sentences in an artificial language that contain a probabilistic dependency and are later asked to produce sentences in that language. The general finding is that, as in more naturalistic situations, adults tend to reproduce the inconsistencies in the input while children tend to regularize and produce patterns that did not exist in the original training set (Hudson Kam & Newport, 2005, 2009; Pitts Cochran, McDonald, & Parault, 1999). Thus, there appears to be a convergence between the creolization literature and laboratory-based investigations of artificial grammar learning. One hypothesis for why children are good at regularizing noisy input, while adults are not, is that children's regularization reflects domain-general developmental difference in learning about and responding to probabilistic information.

The hypothesis linking age-related differences in language learning to probability learning is that children's tendency to regularize noisy language input is analogous to children's tendencies to maximize in probability learning tasks. In both cases, children ignore random variation to create a

systematic pattern. Many researchers have concluded that age-related differences on probability learning tasks may help explain age-related differences in language learning (Hudson Kam & Chang, 2009; Hudson Kam & Newport, 2005, 2009; James & Koehler, 2011; Perfors & Burns, 2010; Pitts Cochran et al., 1999; Ramscar & Gitcho, 2007; Singleton & Newport, 2004; Wonnacott, 2011; Yurovsky, Boyer, Smith, & Yu, 2013). One goal of the present review is to determine if this analogy between language learning and behavior on probability learning tasks is merited.

As has been made clear from reviewing the relevant probability learning literature, differences between children and adults on probability learning tasks are not straight-forward. First, there is not even a clear consensus on what child behavior is like, given how noisy the results seem to be. Second, a number of task characteristics can systematically shift behavior, so it is hard to identify a canonical adult or child behavioral pattern independent of various experimental design choices. Even if there were a stronger consensus about differences between children and adults, it would be difficult to attribute those differences to differences in *learning* versus other kinds of factors. For example, the probability learning literature would suggest that children's maximizing behavior should be caused by not knowing to, or not being able to find a pattern in the input. Children's behavior on probabilistic artificial grammar tasks may also be a product of different strategies and biases, or something different altogether. It is unclear whether there might be commonalities in the strategic processes that underlie age-related differences in probability learning and language learning. However, what is clear from this review is that age-related differences in probability learning tasks do not readily explain age-related differences in language learning but are rather part of the question of what changes, and why, throughout cognitive development.

If qualitative differences in probability learning are not informative of age-related tendencies to generalize on language learning tasks, what might be a potential explanation? One proposed explanation is that age-related differences in memory may contribute to the observed behavior. Hudson Kam and Chang (2009) found less regularization in adults in an artificial grammar task like that of Hudson Kam and Newport (2009) when lexical retrieval was eased by presenting participants with the words of the artificial language on flashcards during the production phase of the task, suggesting a link between memory load and regularization. Likewise, Pitts Cochran et al. (1999) taught adults an artificial sign language and found that adults tended to have trouble rearranging subcomponents of complex signs, but rather tended to produce the signs holistically, as they had been learned.

However, when participants had to complete a simultaneous task to increase memory load, they were better able to recombine the learned signs. These results manipulating memory load with adults are consistent with the idea that children's limited memory abilities force them to store smaller chunks of linguistic information, which may take more time than storing larger chunks but ultimately makes the sign-to-meaning mapping easier and facilitates regularization. Future research might investigate the extent to which these memory and chunking abilities must be domain general and relate to experience-independent development, or whether they may be domain specific—related specifically to experience with language. These differences may derive from biological developmental factors and neuronal changes across the lifetime. Alternatively, these observed differences in memory might derive from greater experience with language itself. Memory in a domain is tied to expertise in a domain (e.g., [Jones & Macken, 2015](#); [Klem et al., 2015](#); [MacDonald & Christiansen, 2002](#); [Schwering & MacDonald, 2020](#); see [Simmering & Perone, 2013](#) for a discussion related to spatial memory), so there are many potential explanations for developmental trajectories of memory capacity.

Other data suggests that the relationship between memory capacity and generalization has been oversimplified. In some studies, children are only more likely than adults to over generalize probabilistic patterns in certain situations. In some tasks, children do not seem to exhibit the classic over-regularization ([Wonnacott, 2011](#)) and in fact show patterns of data in probabilistic artificial grammar learning tasks more similar to adults ([Wonnacott, 2011](#); [Wonnacott, Newport, & Tanenhaus, 2008](#)). Further, memory limitations in adults do not necessarily lead to over-regularization ([Perfors, 2012](#); [Perfors & Burns, 2010](#)), suggesting memory alone cannot explain variability in regularization behavior. Clearly the relationship between sequence learning, language, task, and memory is complicated. Future research investigating the relationship between memory, sequence learning, and various sorts of experience that an individual might bring to the laboratory may help clarify this complicated relationship and help us understand the origin of language learning differences between children and adults.

4.2 A note about generalization

A point that is often overlooked in the discussion of generalization by children is that a single behavior can be considered generalizing or matching,

depending on the units of behavior that are counted. For example, when English-speaking children learn not to over-regularize the past tense (e.g., produce *went* and *ate* rather than *goed* or *eated*), they do not always use the regular past tense in all obligatory contexts. Rather, they alternate between the correct form (*ate*) and the over-regularized form (*eated*) across situations and even from one day to the next (Marcus et al., 1992; McClelland & Patterson, 2002). While children's over regularization of irregular forms is often thought of as akin to probability maximizing, as children indeed over-regularize irregular verb forms, this is not the only interpretation of the data. Children's acquisition of the regular past tense suffix is characterized by probabilistic alternation of the correct and over-regularized form of a verb, which does not resemble maximizing, when the behavior is observed across multiple hours, days or months.

Another point to consider is that the situations approximated by artificial grammar studies, situations in which there is probabilistic variability, are quite rare in the linguistic environment. In language, variation is rarely arbitrary, and typically occurs with a difference in meaning. For example, "he" and "she" do not alternate probabilistically, but rather correlate with the gender of the referent. It is important to consider the extent to which language acquisition studies with probabilistic dependencies are representative of language acquisition. Future studies may complement existing artificial grammar studies by investigating alternations that are not probabilistic, but relate to differences in meaning, as alternations do in typical language.

Another question that merits additional attention is what noisy input actually sounds like to children. To a child receiving probabilistic linguistic input, are low-frequency productions perceived as low-frequency productions, or just nonsense? When children have little linguistic experience, perhaps these low-frequency occurrences are just interpreted as buzz or a cough might be, and do not necessarily get perceived or encoded or weighted the same as high-frequency productions do. If children are more likely than adults to discount these rare productions, it should hardly be surprising that children will exhibit more clear maximizing behavior.

4.3 Implications for other domains

The notion of probability learning is also evoked to explain the development of causal inference in children. Specifically, in a Bayesian inference framework, sometimes references to probability matching or maximizing are made to explain how children select between different hypotheses

that might explain an underlying pattern of data. In many of these tasks, children make predictions or guesses, after viewing some probabilistic string of events, about what color block is likely to emerge from a bag, or what objects make an apparatus (often called a “blicket detector”) turn on. There is evidence that children often alternate between different explanations for data, but rather than being truly random, this alternation may reflect some ordered or rational process. For example, children seem to be sensitive to the probability of different explanations for the observed data (Bonawitz & Lombrozo, 2012; Kushnir & Gopnik, 2007) and more make causal inferences consistent with the probabilities of competing explanations of the data.

Beyond simply being sensitive to the probability of competing hypotheses, some proposals for how children select, or alternate between different hypotheses evoke probability learning. Under this analogy children might “maximize” or always select the mostly likely hypothesis, or “probability match” and entertain hypotheses proportional to their probabilities. In some studies, children’s behavior seems to more closely resemble probability matching, as children select different hypotheses proportional to the relative probabilities of those hypotheses (Denison, Bonawitz, Gopnik, & Griffiths, 2013). In other studies, behavior is more like “undershooting maximizing,” where children occasionally select the less likely hypothesis, but less often than its proportionality would suggest (Kushnir & Gopnik, 2007; Sobel, Tenenbaum, & Gopnik, 2004). The goal of many of these causal inference studies is to understand the rationality (as opposed to random choice or irrationality) that may underlie children’s choice of hypothesis. Probability learning may be evoked to explain why children’s behavior is in fact more rational than it may initially appear. Thus, understanding more about probability learning, specifically whether, and under what circumstances behavior on probability learning tasks might be considered rational, is important if probability learning is evoked as an explanation for children’s behavior on causal inference tasks.

4.4 Implications for the process of science

The present review has a number of implications for the process of science, beyond the content of the review itself. Some patterns that emerge from this review may serve as broad insights in the psychological sciences as a whole.

First, a theme that emerges from the present review is the important role of task design in behavioral studies. The animal and adult literatures

most profoundly show that changing features of the task: the length of the task, what happens when an incorrect target is selected (guidance, correction and or-correction), how sequences of reinforced alternatives are selected, and various aspects of task instructions given to participants all profoundly affect behavior. An important finding of the present review, which may hold true for other domains, is that it may not be possible to identify a canonical pattern of behavior outside of these task characteristics. Patterns of behavior on a task may not be satisfactorily described without a deep understanding of the task and task demands that produced that pattern of behavior.

Second, in performing this review, it is clear that researchers occasionally miscite research, attributing claims that the authors did not make, or selectively cite research. It is only through a deeper dive into the literature that some of these citation errors become evident. For example, many papers cite [Estes \(1964\)](#), [Estes \(1976\)](#), [Myers \(1976\)](#) and [Vulkan \(2000\)](#) as providing evidence that adults probability match in probability learning studies, but these papers do not make this claim. The two Estes papers make versions of this claim, with the earlier paper making the claim of probability matching more strongly. However, both papers include substantial discussion of the experimental design dimensions that affect behavior and remind the reader that the observed behavior in probability learning tasks reflects a complex combination of different types of knowledge and learning. [Myers \(1976\)](#) and [Vulkan \(2000\)](#) do not claim that adults tend to probability match. The bulk of [Myers \(1976\)](#) considers different mathematical models of probability learning and how to best account for the observation that extended training, task instructions, and monetary reward are associated with *over-shooting* matching, not true matching. [Vulkan \(2000, p. 15\)](#) concludes that “not all hope is lost for the neoclassical approach to economics: “matching” is not robust, and some people do end up maximizing their expected utility.” This is clearly not an endorsement of probability matching in adults. Though [Vulkan \(2000\)](#) does focus on reasons that adult performance might diverge from true maximizing, it is in the context of trying to understand the biases that adults might bring to the task and the various task characteristics (not unlike those described in my review) that affect behavior. None of these papers support the strong claim that adults tend to probability match on probability learning tasks.

An example of selective citation is the high citation rate of papers that find evidence that adults probability match or children maximize. For example, [Neimark and Shuford \(1959\)](#) is extensively cited as evidence that adults probability match (even by the great William Estes), despite the fact that it

is an outlier among the adult studies and few other adult studies obtained this result. In fact, [Neimark and Shuford \(1959\)](#) has three times the citations as [Morse and Runquist \(1960\)](#), the aluminum rod study with the yoked light prediction task, which showed that how the reinforcement schedule was generated clearly affected participants' tendencies to match, overshoot matching or maximize. Authors seem to be selectively citing work consistent with claims they want to make, even if those findings are not characteristic of the literature as a whole. Put frankly, the file drawer problem is the least of our concerns if we selectively read the published literature.

On a more positive note, the present review also illustrates the richness and depth of the existing literature and how much can be gained by re-learning what was once known. Likewise, a deep reading of a related literature, for example, non-human animal probability learning, may provide enormous insight on a related literature, for example, human adult probability learning. A lot of knowledge exists in the literature, sometimes in unexpected places. Some research questions are best answered with new studies, and some are best answered, or at least guided, by revisiting old studies. Perhaps we need to remind ourselves and our students that a deep and broad understanding of the literature is an invaluable asset for even the most forward-thinking research program.



5. Conclusions

This review has yielded a number of findings. In probability learning tasks, non-human animals tend to maximize if the task consists of a sufficient number of trials. Human adults maximize under some conditions, and their failure to always maximize seems to be a consequence of attempting to find a pattern in the probability sequence, or other sources of knowledge or bias that adults bring to the task. The probability learning findings about children are less conclusive. There is some evidence that children maximize more than adults, but other studies refute that pattern. Perhaps if children do maximize more than adults, it is because they cannot, or do not know that they should, look for a pattern, or that they do not have the same expectations about the nature of the generative process of that sequence as do adults.

One important conclusion that I draw from this review is that there is no discrete learning strategy that asymptotes at the π (probability level) over training. The range of behavior on probability learning tasks, which is generally characterized by maximizing and undershooting maximizing, can better be described as a single learning mechanism that would over time

lead to maximization, plus a number of conditions and biases that impede maximizing. These conditions and biases may include experiences, beliefs or even memory limitations that the organism brings to the task. The fundamental questions of probability learning should not be “when do individuals maximize and when do they probability match” but rather, “what factors lead participants to undershoot maximization.” It is important to use terminology that accurately describes the behavior we observe, which is “maximizing” and “undershooting maximizing” to reflect the continuum of behavior that an individual may exhibit in a given probability learning task.

More generally, age-related differences in probability learning tasks are sometimes evoked to help us understand developmental trajectories in various domains. Understanding more about the factors that drive behavior in probability learning tasks themselves is important for understanding the explanatory power, or lack of explanatory power, of probability learning behavior as an explanation for patterns and trajectories of behavior in other cognitive domains.

Acknowledgments

This review was initially written as part of my preliminary exams as a graduate student. I'd like to thank my graduate advisor, Maryellen C. MacDonald, for her guidance and many helpful discussions. Look Ma! It's finally published! This research was supported by the National Institutes of Child Health and Human Development Grants T32 HD049899 and R01 HD047425; the National Science Foundation #1123788 and the Wisconsin Alumni Research Fund, all to Maryellen C. MacDonald. I'd also like to thank Jon A. Willits for helpful discussion.

References

- Anderson, N. H. (1960). Effects of first-order conditional probability in a two-choice learning situation. *Journal of Experimental Psychology*, 59, 73–93.
- Arima, J. K. (1965). Human probability learning with forced training trials and certain and uncertain outcome choice trials. *Journal of Experimental Psychology*, 70, 43–50.
- Ayton, P., & Fischer, I. (2004). The hot hand and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition*, 32, 1369–1378.
- Bar-Hillel, M., & Wagenaar, W. A. (1991). The perception of randomness. *Advances in Applied Mathematics*, 12, 428–454.
- Beach, L. R., & Swenson, R. G. (1967). Instructions about randomness and run dependency in two-choice learning. *Journal of Experimental Psychology*, 75, 279.
- Behrend, E. R., & Bitterman, M. E. (1961). Probability-matching in the fish. *The American Journal of Psychology*, 74, 542–551.
- Bialystok, E., & Hakuta, K. (1999). Confounded age: Linguistic and cognitive factors in age differences for second language acquisition. In D. Birdsong (Ed.), *Second language acquisition and the critical period hypothesis* (pp. 161–181). Mahwah, NJ: Erlbaum.
- Bitterman, M. E. (1965). Phyletic differences in learning. *American Psychologist*, 20, 396–410.

- Bitterman, M. E. (1971). Visual probability learning in the rat. *Psychonomic Science*, 22, 191–192.
- Bitterman, M. E., Wodinsky, J., & Candland, D. K. (1958). Some comparative psychology. *American Journal of Psychology*, 71, 94–110.
- Bonawitz, E. B., & Lombrozo, T. (2012). Occam's rattle: Children's use of simplicity and probability to constrain inference. *Developmental Psychology*, 48, 1156–1164.
- Braveman, N. S., & Fischer, G. J. (1968). Instructionally induced strategy and sequential information in probability learning. *Journal of Experimental Psychology*, 76, 674–676.
- Bullock, D. H., & Bitterman, M. E. (1962). Probability-matching in the pigeon. *The American Journal of Psychology*, 75, 634–639.
- Calfee, R. C. (1968). Choice behavior during long-term probabilistic reinforcement schedules. *Journal of Comparative and Physiological Psychology*, 65, 232–237.
- Cotton, J. W., & Rechtschaffen, A. (1958). Replication report: Two-and three-choice verbal-conditioning phenomena. *Journal of Experimental Psychology*, 56, 96.
- Craig, G. J., & Myers, J. L. (1963). A developmental study of sequential two-choice decision making. *Child Development*, 34, 483–493.
- Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The sampling hypothesis. *Cognition*, 126, 285–300.
- Denison, S., & Xu, F. (2010). Twelve-to 14-month-old infants can predict single-event probability with large set sizes. *Developmental Science*, 13, 798–803.
- Derks, P. L., & Paclisanu, M. I. (1967). Simple strategies in binary prediction by children and adults. *Journal of Experimental Psychology*, 73, 278–285.
- Edwards, W. (1956). Reward probability, amount, and information as determiners of sequential two-alternative decisions. *Journal of Experimental Psychology*, 52, 177–188.
- Edwards, W. (1961). Probability learning in 1000 trials. *Journal of Experimental Psychology*, 62, 385–394.
- Elman, J. L., Bates, E. L., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective of development*. Cambridge MA: MIT Press.
- Estes, W. K. (1964). Probability learning. In A. W. Melton (Ed.), *Categories of human learning*. New York: Academic Press.
- Estes, W. K. (1976). The cognitive side of probability learning. *Psychological Review*, 83, 37–64.
- Estes, W. K., & Johns, M. D. (1958). Probability learning with ambiguity in the reinforcing stimulus. *The American Journal of Psychology*, 71, 219–228.
- Estes, W. K., & Straughan, J. H. (1954). Analysis of a verbal conditioning situation in terms of statistical learning theory. *Journal of Experimental Psychology*, 47, 225–234.
- Falk, R., Falk, R., & Levin, I. (1980). A potential for learning probability in young children. *Educational Studies in Mathematics*, 11, 181–204.
- Falk, R., & Konold, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review*, 104, 301–318.
- Fantino, E., & Esfandiari, A. (2002). Probability matching: Encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology*, 56, 58–63.
- Fetterman, J., & Stubbs, D. (1982). Matching, maximizing, and the behavioral unit: Concurrent reinforcement of response sequences. *Journal of the Experimental Analysis of Behavior*, 37, 97–114.
- Fischer, G. J. (1972). Maximizing and role of correction method in probability learning in chicks. *Journal of Comparative and Physiological Psychology*, 80, 49–53.
- Friedman, M. P., Padilla, G., & Gelfand, H. (1964). The learning of choices between bets. *Journal of Mathematical Psychology*, 1, 375–385.
- Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, 109, 416–422.

- Gardner, R. A. (1957). Probability-learning with two and three choices. *American Journal of Psychology*, 70, 174–185.
- Gardner, R. A. (1958). Multiple-choice decision-behavior. *The American Journal of Psychology*, 71, 710–717.
- Goldman, A. E., & Denny, J. (1963). The ontogenesis of choice behavior in stochastic and recursive programs. *Journal of Genetic Psychology*, 102(1), 5–18.
- Gonzalez, R. C., Roberts, W. A., & Bitterman, M. E. (1964). Learning in adult rats with extensive cortical lesions made in infancy. *The American Journal of Psychology*, 77, 547–562.
- Goodnow, J. J. (1955). Response-sequences in a pair of two-choice probability situations. *American Journal of Psychology*, 68, 624–630.
- Goodnow, J. J., & Pettigrew, T. F. (1955). Effect of prior patterns of experience upon strategies and learning sets. *Journal of Experimental Psychology*, 49, 381–389.
- Goodnow, J. J., & Postman, L. (1955). Probability learning in a problem-solving situation. *Journal of Experimental Psychology*, 49, 16–22.
- Graf, V., Bullock, D. H., & Bitterman, M. S. (1964). Further experiments on probability-matching in the pigeon. *Journal of the Experimental Analysis of Behavior*, 7, 151–157.
- Grant, D. A., Hake, H. W., & Hornseth, J. P. (1951). Acquisition and extinction of verbal expectations in a situation analogous to conditioning. *Journal of Experimental Psychology*, 42, 1–5.
- Grant, D. A., Hornseth, J. P., & Hake, H. W. (1950). The influence of the intertribal interval on the Humphreys' "random reinforcement" effect during the extinction of a verbal response. *Journal of Experimental Psychology*, 40, 609–612.
- Green, C. S., Benson, C., Kersten, D., & Schrater, P. (2010). Alterations in choice behavior by manipulations of world model. *Proceedings of the National Academy of Sciences of the United States of America*, 107, 16401–16406.
- Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences of the United States of America*, 107, 9066–9071.
- Hake, H. W., & Hyman, R. (1953). Perception of the statistical structure of a random series of binary symbols. *Journal of Experimental Psychology*, 45, 64–74.
- Hinson, J. M., & Staddon, J. E. R. (1983). Hill-climbing by pigeons. *Journal of the Experimental Analysis of Behavior*, 39, 25–47.
- Hiraoka, K. (1984). Discrete-trial probability learning in rats: Effects of local contingencies of reinforcement. *Animal Learning & Behavior*, 12, 343–349.
- Hudson Kam, C. L., & Chang, A. (2009). Investigating the cause of regularization in adults: Memory constraints or learning differences? *JEP: Learning, Memory, and Cognition*, 35, 815–821.
- Hudson Kam, C. L., & Newport, E. L. (2005). Regularizing unpredictable variation: The roles of adult and child learners in language formation and change. *Language Learning and Development*, 1, 151–195.
- Hudson Kam, C. L., & Newport, E. L. (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology*, 59, 30–66.
- Humphreys, L. G. (1939). The effect of random alternation of reinforcement on the acquisition and extinction of conditioned eyelid reactions. *Journal of Experimental Psychology*, 25, 141–158.
- Hyman, R., & Jenkin, N. S. (1956). Involvement and set as determinants of behavioral stereotypy. *Psychological Reports*, 2, 131–146.
- James, G., & Koehler, D. J. (2011). Banking on a bad bet: Probability matching in risky choice is linked to expectation generation. *Psychological Science*, 6, 707–711.

- Jarvik, M. E. (1951). Probability learning and a negative recency effect in the serial anticipation of alternative symbols. *Journal of Experimental Psychology*, 41, 291–297.
- Johnson, J. S., & Newport, E. L. (1991). Critical period effects on universal properties of language: The status of subadjacency in the acquisition of a second language. *Cognition*, 39, 215–258.
- Jones, J. E. (1961). Deviations from matching behavior in probability learning. *Journal of Psychology: Interdisciplinary and Applied*, 52, 335–345.
- Jones, M. H., & Liverant, S. (1960). Effects of age differences on choice behavior. *Child Development*, 31, 673–680.
- Jones, G., & Macken, B. (2015). Questioning short-term memory and its measurement: Why digit span measures long-term associative learning. *Cognition*, 144, 1–13.
- Kirk, K. L., & Bitterman, M. E. (1965). Probability-learning by the turtle. *Science*, 148, 1484–1485.
- Klem, M., Melby-Lervåg, M., Hagtvet, B., Lyster, S. A. H., Gustafsson, J. E., & Hulme, C. (2015). Sentence repetition is a measure of children's language skills rather than working memory limitations. *Developmental Science*, 18, 146–154.
- Koehler, D. J., & James, G. (2009). Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, 113, 123–127.
- Koehler, D. J., & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition*, 38(6), 667–676.
- Kotsinas, U. B. (1988). Immigrant Children's Swedish: A new variety? *Journal of Multilingual and Multicultural Development*, 9, 129–140.
- Kushnir, T., & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology*, 43, 186–196.
- Lauer, D. W., & Estes, W. K. (1954). Observed and predicted terminal distributions of response probability under two conditions of random reinforcement. *American Psychologist*, 9, 413.
- Lenneberg, E. (1967). *Biological foundations of language*. New York: Wiley.
- Longo, N. (1964). Probability-learning and habit-reversal in the cockroach. *American Journal of Psychology*, 77, 29–34.
- Lopes, L. L. (1982). Doing the impossible: A note on induction and the experience of randomness. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 8, 626–636.
- MacDonald, M. C., & Christiansen, M. H. (2002). Reassessing working memory: Comment on just and carpenter (1992) and waters and Caplan (1996). *Psychological Review*, 109, 35–54.
- Marcus, G. F., Pinker, S., Ullman, M., Hollander, M., Rosen, T. J., Xu, F., et al. (1992). Overregularization in language acquisition. *Monographs of the Society for Research in Child Development*, 57, 1–165.
- McClelland, J. L., & Patterson, K. (2002). Rules or connections in past-tense inflections: What does the evidence rule out? *Trends in Cognitive Sciences*, 6(11), 465–472.
- McCracken, J., Osterhout, C., & Voss, J. F. (1962). Effects of instructions in probability learning. *Journal of Experimental Psychology*, 64, 267–271.
- Meyer, D. R. (1960). The effects of differential probabilities of reinforcement on discrimination learning by monkeys. *Journal of Comparative Physiological Psychology*, 53, 173–175.
- Morse, E. B., & Runquist, W. N. (1960). Probability-matching with an unscheduled random sequence. *The American Journal of Psychology*, 73, 603–607.
- Myers, J. L. (1976). Probability learning and sequence learning. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Approaches to human learning and motivation* (pp. 171–205). Erlbaum: Hillsdale, NJ.
- Myers, J. L., Fort, J. G., Katz, L., & Suydam, M. M. (1963). Differential monetary gains and losses and event probability in a two-choice situation. *Journal of Experimental Psychology*, 66, 521.
- Neimark, E. D., & Shuford, E. H. (1959). Comparison of predictions and estimates in a probability learning situation. *Journal of Experimental Psychology*, 57(5), 294–298.

- Neville, H. J., & Bruer, J. (2001). Language processing: How experience affects brain organization. In D. B. Bailey Jr., J. T. Bruer, F. J. Symons, & J. W. Lichtman (Eds.), *Critical thinking about critical periods: Perspectives from biology, psychology and education* (pp. 151–172). Baltimore, MD: Paul H. Brookes Publishing.
- Newell, B. R., & Rakow, T. (2007). The role of experience in decisions from description. *Psychonomic Bulletin & Review*, 14, 1133–1139.
- Newport, E. (1988). Constraints on learning and their role in language acquisition: Studies of the acquisition of American sign language. *Language Sciences*, 10, 147–172.
- Newport, E. L. (1990). Maturation constraints on language learning. *Cognitive Science*, 14, 11–28.
- Newport, E. L. (1999). Reduced input in the acquisition of signed languages: Contributions to the study of creolization. In M. DeGraff (Ed.), *Language creation and language change: Creolization, diachrony, and development* (pp. 161–178). Cambridge, MA: MIT Press.
- Offenbach, S. I. (1964). Studies of children's probability learning behavior: I. Effect of reward and punishment at two age levels. *Child Development*, 35, 709–716.
- Otto, A. R., Taylor, E. G., & Markman, A. B. (2011). There are at least two kinds of probability matching: Evidence from a secondary task. *Cognition*, 118, 274–279.
- Perfors, A. (2012). When do memory limitations lead to regularization? An experimental and computational investigation. *Journal of Memory and Language*, 67, 486–506.
- Perfors, A., & Burns, N. (2010). Adult language learners under cognitive load do not over-regularize like children. In S. Ohlsson, & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2524–2529). Cognitive Science Society.
- Peterson, C. R., & Uehla, Z. J. (1965). Sequential patterns and maximizing. *Journal of Experimental Psychology*, 69, 1–4.
- Pitts Cochran, B., McDonald, J., & Parault, S. (1999). Too smart for their own good: The disadvantage of a superior processing capacity for adult language learners. *Journal of Memory and Language*, 41, 30–58.
- Plunkett, K. (1997). Theories of early language acquisition. *Trends in Cognitive Sciences*, 1(4), 146–153.
- Ramsar, M. J. A., & Gitcho, N. (2007). Developmental change and the nature of learning in childhood. *Trends in Cognitive Sciences*, 11(7), 274–279.
- Ross, B. M. (1954). Probability predictions between and within series. *Journal of Psychology*, 37, 279–289.
- Rubenstein, I. (1959). Some factors in probability matching. *Journal of Experimental Psychology*, 57, 413–416.
- Saldana, C., Claidière, N., Fagot, J., & Smith, K. (2020). Rational after all: Changes in probability matching behaviour across time in humans and monkeys. In S. Denison, M. Mack, Y. Xu, & B. C. Armstrong (Eds.), *Proceedings of the 42nd annual conference of the cognitive science society* (pp. 752–758). Cognitive Science Society.
- Sankoff, G., & Laberge, Z. (1973). On the acquisition of native speakers by a language. *Kivung*, 6, 32–47.
- Schweikert, G. E., & Treichler, F. R. (1969). Visual probability learning and reversal in the cat. *Journal of Comparative and Physiological Psychology*, 67, 269–272.
- Schwering, S. C., & MacDonald, M. C. (2020). Verbal working memory as emergent from language comprehension and production. *Frontiers in Human Neuroscience*, 14, 68.
- Scott, R. M., & Baillargeon, R. (2017). Early false-belief understanding. *Trends in Cognitive Sciences*, 21, 237–240.
- Seidenberg, M. S., & Zevin, J. D. (2006). Connectionist models in developmental cognitive neuroscience: Critical periods and the paradox of success. In Y. Munakata, & M. Johnson (Eds.), *Processes of change in brain and cognitive development. Attention and performance XXI* (pp. 315–347). Oxford, UK: Oxford University Press.
- Senghas, A., & Coppola, M. (2001). Children creating language: How Nicaraguan sign language acquired a spatial grammar. *Psychological Science*, 12, 323–328.

- Senghas, A., Kita, S., & Ozyurek, A. (2004). Children creating core properties of language: Evidence from an emerging sign language in Nicaragua. *Science*, 305, 1779–1782.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15, 233–250.
- Shimp, C. P. (1966). Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, 9, 443–455.
- Shimp, C. P. (1967). Reinforcement of least-frequent sequences of choices. *Journal of the Experimental Analysis of Behavior*, 10, 57–65.
- Shimp, C. P. (1970). A within-session effect after prolonged training in probability learning by rats. *Psychonomic Science*, 18, 152–153.
- Siegel, S., & Goldstein, D. A. (1959). Decision-making behavior in a two-choice uncertain outcome situation. *Journal of Experimental Psychology*, 57, 37.
- Simmering, V. R., & Perone, S. (2013). Working memory capacity as a dynamic process. *Frontiers in Psychology*, 3, 567.
- Singer, E. (1967). Ability and the use of optimal strategies in decisions. *The American Journal of Psychology*, 80, 243–249.
- Singer, E. (1968). Type of grain preferred and strategy on a probabilistic decision-task. *The American Journal of Psychology*, 81, 36–41.
- Singleton, J. L., & Newport, E. L. (2004). When learners surpass their models: The acquisition of American sign language from inconsistent input. *Cognitive Psychology*, 49, 370–407.
- Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, 28, 303–333.
- Stevenson, H. W., & Weir, M. W. (1959). Variables affecting children's performance in a probability learning task. *Journal of Experimental Psychology*, 57, 403–412.
- Sullivan, F. J., & Ross, B. M. (1970). What is learned in probability learning. *Developmental Psychology*, 2, 58–65.
- Suppes, P., & Atkinson, R. C. (1960). *Markov learning models for multiperson interactions*. Stanford, CA: Stanford University Press.
- Topping, J. S., & Parker, B. (1970). Constant and variable delay of reinforcement effects on probability learning by pigeons. *Journal of Comparative and Physiological Psychology*, 70, 141–147.
- Treichler, F. R. (1967). Reinforcer preference effects on probability learning by monkeys. *Journal of Comparative and Physiological Psychology*, 64, 339–342.
- Uhl, C. N. (1963). Two-choice probability learning in the rat as a function of incentive, probability of reinforcement, and training procedure. *Journal of Experimental Psychology*, 66, 443–449.
- Unturbe, J., & Corominas, J. (2007). Probability matching involves rule-generating ability: A neuropsychological mechanism dealing with probabilities. *Neuropsychology*, 21, 621–630.
- Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, 14, 101–118.
- Wagenaar, W. A. (1972). Generation of random sequences by human subjects: A critical survey of the literature. *Psychological Bulletin*, 77, 65–72.
- Warren, J. M., & Beck, C. H. (1966). Visual probability learning by cats. *Journal of Comparative and Physiological Psychology*, 61, 316–318.
- Weir, M. W. (1964). Developmental changes in problem-solving strategies. *Psychological Review*, 71, 473–490.
- Weitzman, R. A. (1967). Positional matching in rats and fish. *Journal of Comparative and Physiological Psychology*, 63, 54–59.
- West, R. F., & Stanovich, K. E. (2003). Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Memory & Cognition*, 31, 243–251.

- Williams, J. T., & Albiniak, B. A. (1972). Probability learning in a crocodilian. *Psychonomic Science*, 27, 165–166.
- Wilson, W. A. (1960). Two-choice behavior of monkeys. *Journal of Experimental Psychology*, 59, 207–208.
- Wilson, W. A., Oscar, M., & Bitterman, M. E. (1964). Probability-learning in the monkey. *Quarterly Journal of Experimental Psychology*, 16, 163–165.
- Wilson, W. A., & Rollin, A. R. (1959). Two-choice behavior of rhesus monkeys in a noncontingent situation. *Journal of Experimental Psychology*, 58, 174–179.
- Winefield, A. H. (1980). Probability learning in children as a function of age and reinforcement procedure. *The Journal of Genetic Psychology: Research and Theory on Human Development*, 137, 79–90.
- Wolford, G., Newman, S. E., Miller, M. B., & Wig, G. S. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology*, 58, 221.
- Wonnacott, E. (2011). Balancing generalization and lexical conservatism: An artificial language study with child learners. *Journal of Memory and Language*, 65, 1–14.
- Wonnacott, E., Newport, E. L., & Tanenhaus, M. K. (2008). Acquiring and processing verb argument structure: Distributional learning in a miniature language. *Cognitive Psychology*, 56, 165–209.
- Wright, R. L. D. (1967). Motivational effects in probability learning. *Psychonomic Science*, 7, 329–330.
- Xu, F., & Denison, S. (2009). Statistical inference and sensitivity to sampling in 11-month-old infants. *Cognition*, 112, 97–104.
- Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences of the United States of America*, 105, 5012–5015.
- Yurovsky, D., Boyer, T. W., Smith, L. B., & Yu, C. (2013). Probabilistic cue combination: Less is more. *Developmental Science*, 16, 149–158.