

Investigating how infants learn to search in the A-not-B task

Hanna Popick,¹ Melody Dye,¹ Natasha Kirkham² & Michael Ramscar¹

¹Department of Psychology,
Stanford University

²Department of Psychological Sciences
Birkbeck College, University of London

Abstract

In searching for hidden objects, infants younger than 12 months frequently commit the classic “A-not-B error,” in which they successfully search for an object in one location (A) and then fail to search for it when it is conspicuously hidden in a new location (B). The question is why they fail to make the switch and perseverate at the first location. While these errors have often been attributed to cognitive limitations or stages of neurological development, we propose that they are consistent with the early stages of learning. We present a context-learning model of “A-not-B” search, in which learning to adopt the appropriate search strategy involves attending to appropriate contextual cues. We then present the findings of an eye-tracking experiment with 9 month-olds that behaviorally supports the predictions of our learning model.

Keywords: A-not-B, Learning Theory, Computational Modeling, Causal Reasoning

Introduction

It’s Monday morning. You haven’t seen your car keys since Friday. How do you find them? In an ideal world, you might just go look where you last saw them. You’ve learned, after all, that keys don’t usually move on their own, and that what best predicts a key’s location is the conjunction of a given spot and you having last seen the key at that spot. (So you might even attempt this search pattern if that “last spot” is a very strange place for keys to be.) But what if the world is less than ideal—what if you’re not sure where you last saw the keys? You may have other memories of them—clear recollections of them hanging on the hook reserved especially for your car keys, for example—which might compete with your memories from Friday. Indeed, if you aren’t sure where you last saw them, you may still check the hook first, because you know that searching at a location where the keys are seen frequently can be a successful search strategy.

This characterization of adult “expertise” when it comes to looking for keys sheds light on the task facing a child learning how to find things in the world. A child must learn that things are likely to be at the location they were last seen, that things are also likely to be at the location they are most often seen, and that a successful search will involve weighing these considerations against what the child can remember about the last and most likely locations of an object. From this perspective, “perseverative errors,” in which a child searches for an object in a likely location rather than in the hiding location they just witnessed might be seen as a misapplication of what in other circumstances might be a completely rational strategy.

Piaget (1954) first described what are often called “A-not-B errors” in infants: 8- to 12-month infants will generally search successfully for an object in one location (A), but then fail to search for it when it is conspicuously hidden in a new location (B). Subsequent studies have confirmed that in actively searching for hidden objects, infants robustly make prototypical A-not-B errors, ignoring the most recent location of objects when they search for them after a switch (Marcovitch & Zelazo, 1999; Wellman, Cross, Bartsch, & Harris, 1986). In seeking to explain this, accounts often focus on the possibility that infants’ errors stem from problems associated with implementing a correct search, such as limited working memory and inhibitory control, or from weak memory traces for the object and hiding location (e.g. Baillargeon, Graber, Devos, & Black, 1990; Diamond, 1988; Diamond, Cruttenden, & Neiderman, 1994; Munakata, 1997; Thelen, Schöner, Scheier, & Smith, 2001).

In what follows, we explore an alternative, though complementary approach. Rather than assuming that a child already understands how to search, we consider what might be expected if a child were learning how to search. As noted above, in learning how to successfully seek out objects, children have to figure out that some things may be more likely to be where they were last seen, and others where they are most often seen, and they have to learn which strategy is most appropriate in each context.

How might children learn search strategies, and their appropriate application? One way to consider these questions is within the framework provided by formal learning theories (e.g. Rescorla & Wagner, 1972), which see learning as the process of acquiring information about the relationship between events (outcomes) in the environment, and the cues that allow them to be predicted. From this perspective, children learn to search via a process of trial and error, strengthening or weakening the value of cues depending on their predictive successes and failures. This takes time. Infants must make, and learn from, more or less successful predictions before they can “master” search, and begin to match search strategies to context. As we shall show, while to an adult, children’s search behavior may appear to be erroneous (“perseverative”) from this perspective, a child’s A-not-B “errors” can be seen as a rational, inevitable part of the process of learning to search.

The A-not-B Task

In the classic A-not-B task (Piaget, 1954), 7 to 12-month-old children search for a small object in one of two identical hiding spots. The object is first hidden in plain view of the

infant in one of the locations (A), and after a short delay, the infant usually searches successfully for the object at A. This is repeated for a few trials, after which the object is hidden at location B, and the infant is again given the opportunity to search for the object. At this point, infants often continue to reach perseveratively to the previously correct location (A).

Infants also perseverate when they learn other novel relationships between objects and locations. For example, Aguiar & Baillargeon (2000) showed infants two towels, one with a toy on it (A), and one with a toy behind it (B). The infants learned that pulling towel A enabled them to obtain and play with the toy. When the towel / toy relationships were switched, so that B now brought the toy, 7 month-old infants continued to perseveratively pull towel A rather than switch to towel B. This suggests that the perseverative response has to do with the process of learning predictive relationships between actions and outcomes, and is not solely contingent on particular properties of the objects themselves, or the specific hiding events of the task (for similar findings, see Marcovitch & Zelazo, 2006; Smith, Thelen, Titzer, & McLin, 1999).

However, infant behavior in these tasks is still, in many ways, context-dependent. For instance, while many 9 month old infants can successfully complete the towel pulling task (Aguiar & Baillargeon, 2000), they still fail the standard A-not-B task (Piaget, 1954), even though these tasks appear structurally similar. Further, Adolpho (2000) found that what an infant learns in one context does not always extend easily to another (see also Thelen, Schoner, Scheier, & Smith, 2001; Smith & Thelen, 2003). Thus, infants do not initially appear to learn abstract, generalized “search.” Rather, infant search learning is sensitive both to kind (pulling, reaching, etc) and context.

On the available evidence then, A-not-B errors cannot be attributed to motor perseveration alone: infants can make the switch when outcomes and locations change (as in Aguiar & Baillargeon, 2000). Nor do these errors result from difficulties in conceptualizing objects: infants can search successfully at the first location. Given that it seems that infants understand the task, and can switch, what needs to be explained is why they initially fail to search correctly after a switch in location. Why don't they adopt the appropriate search strategy right away?

Information Structure in Learning

Successfully searching for an object involves weighing a number of clues to its possible location: Where was it last seen? Where is it usually seen? As adults, we can assess these clues within the frame of the search task at hand. However, while it may seem patently obvious that infants should assume that “if the object is hidden at A, search A; if B, then B,” this relationship is not universal in an infant's experience: for example, people and animals will often appear and disappear at ‘random’ (out of one door, back through another); batter ‘hidden’ in the oven will reappear as a cake; and one fine summer day, mommy's ring will

vanish down the garbage disposal, never to be seen again. Given that the relationship between hiding an object at a specific location, and later finding that object in the same location is not universally warranted, it seems likely that infants will need to learn the situations in which it is appropriate, and in which it is not.

To formally illustrate our description of the role context might play in search, and how children might learn the appropriate search strategy for the A-not-B task, we simulated this learning process using the Rescorla-Wagner (1972) model, modeling the learning of cues that represented each location, and the changing strength of their predictive relationship with the object.¹ The simulations assume that in the task, an infant learns about two locations (A and B), that on each trial infants will search at the location that they most strongly associate with the object's potential location, and that the outcome of those searches will be incorporated into further learning. To reflect the fact that infants will most likely spend more time looking towards the location at which they expect objects to reappear, the saliency of the unattended location was set lower than that rate for the attended location. This allowed the model to reflect the likelihood that infants would learn less quickly about location B when they were still primarily attending to A. (This also suggests that as a consequence of their attending to A, infants are initially slow to learn about hiding events at B.) Figure 1 shows these associative strengths developing across A-not-B training trials.




Figure 1. Rescorla-Wagner simulation of cue competition between cues representing the two hiding locations and the association each has with the hidden object across two trials at location A, followed by four trials at location B. Note that although the association between the object and A decreases following the first two trials, it remains at a higher value than the association between the object and B for the first few hiding events at B.

As can be seen from Figure 1, at the outset of A-not-B learning, the relationship between a location (A) and a specific outcome (finding the object) is repeatedly reinforced. The infant learns that location A strongly predicts the reappearance of the object (in the simulation, this is illustrated by an increase in the predictive strength of location A). Then a new relationship is introduced, this time between a different location (B) and the same outcome

(finding the object). Given that the infant has no prior experience of finding objects at B, and given that location A is already a strongly learned cue to finding objects, the infant's best guess, initially, ought to be that the object will continue to reappear at A.

However, if the infant goes on to see more objects hidden at B, and being found at B, two things will happen: First, the infant will learn that location B also predicts the reappearance of the object, and the strength of this predictive relationship will strengthen over B-trials. Second, error resulting from incorrect searches at A during B-trials will weaken the relationship between A and the reappearing object, reducing A's predictive value.

This simple model makes an intriguing empirical prediction about infants in the A-not-B task: namely, that if we extend the number of B trials, infants will gradually unlearn the value of A before switching their search to B. However, this model also implies that infants will use the "most frequent location" strategy in all situations, such that if hiding were to revert to A after a number of B trials, B would need to be unlearned prior to switching back to A. Given that children do learn to search in appropriate locations, this fails to explain how young children come to learn to switch flexibly between locations and succeed at the task.

The answer to this puzzle lies in the different ways adults go about searches: Children must learn that in search, context counts. If infants are to learn to weigh an object's last location over its most frequent location, they need to learn that the conjunction of an object and its last location is the appropriate cue to where that object is most likely to be found. Accordingly, Figure 2 shows the results of an extended simulation to which this kind of "contextual" conjunctive cue has been added to the simple location cues.




Figure 2. A Rescorla-Wagner model of cue competition between two simple cues representing location alone, and two conjunctive cues using information about where the object was last hidden along with location. The model shows associations with the hidden object across two trials at A, followed by four trials at B (a fairly standard A-not-B task), and then makes predictions for how the associations would change if the task were extended with alternating trials at A and B.

Initially, there is little to distinguish the performance of the two models, because during the initial hiding events at A, the value of search at A is there is to be learned. However, after the switch in locations, the associative strengths of the simple location cues weaken relative to the conjunctive cue, because the simple cues suggesting the "search at the most frequent location" response generate error when hiding locations are switched. Because both locations are always present in the A-not-B task, on any given search, the cue value of one (successful) location will be strengthened, while the other will lose value. However, since a child will only see an object hidden in one place, the contextual cues that support searching at the place the object was hidden will prove more accurate, and will therefore strengthen relative to the simple cues over time. In this way, the infant can gradually learn to match an appropriate search strategy to the task at hand.

This process takes time; the model we present suggest that infants will need experience of making more and less successful predictions before they can learn appropriate contextual search strategies, and, importantly, unlearn their tendency to simply search in the most likely location.

Experiment


Though the standard A-not-B task generally involves infants reaching for hidden physical objects, if our hypothesis about the need for infants to unlearn inappropriate search strategies is correct, we would expect the same perseverative pattern of behavior to be apparent in all search strategies, regardless of modality of measurement (see also Diamond, 1990; Hofstadter & Reznick, 1996; Bell & Adams, 1999). Thus to examine our account of learning to search, we conducted a study of the visual search behavior of infants in the A-not-B paradigm using eye-tracking. Taking a more continuous measurement of children's visual search enabled us to examine whether children's pattern of perseverative searching after an initial switch trial simply reflected a belief that the object was in the incorrect location, or whether it was consistent with the gradual unlearning process predicted by our simulations.

Participants

32 9-month-old infants successfully completed our testing procedure (range 8 months 17 days to 9 months 17 days, median 9 months 7 days; with equal gender distribution). Data from an additional 18 infants are not reported due to poor calibration (9), fussiness during the experiment (7), and equipment failure (2). Participants were recruited from a volunteer pool, which reflects the properties of the community surrounding Stanford University.

Stimuli

The infants watched animated movies of a set of colorful keys that were accompanied by musical sound effects played at equal volumes from speakers on both sides of the screen. As part of familiarization, the keys were first shown



moving up and down and rattling in the center of the screen, and were then shown moving across the screen and disappearing into a bucket on one side. An identical bucket was also present on the other side.

Following the hiding of the keys, a pinwheel accompanied by new music appeared in the center as a distracter for three seconds, and then disappeared. For the following four seconds only the buckets were visible, while the music that accompanied the keys played to encourage searching. After this four-second search period, the keys reappeared from the same bucket into which they had disappeared, before moving back towards center screen and then moving off the top of the screen. The pinwheel animation then reappeared in the center and was displayed until the infants' attention to the center of the screen was confirmed, at which point the keys reappeared in the center to begin the next trial. Figure 3 depicts the stimuli and the sequence of events.

Figure 3 (pictured above). A screenshot of one trial of the visual search A-not-B task presented to the infants. This display is shown with two hiding events in one location, followed by four trials of hiding events in the other location.

Procedure and design

Participants sat on a caregiver's lap during testing, facing a 152cm projection screen, which was approximately 180cm from them. An Applied Science Laboratories (ASL) Model 504 corneal reflection eye tracking system collected eye movement data as infants were shown the stimulus displays. A computer script translated the gaze coordinates recorded by the system into gaze durations to regions of interest (ROI) defined around each of the hiding wells during the 4-second search period after each hiding event.

Infants were shown the key-hiding sequence six times: the keys were hidden twice in the bucket on one side of the screen, and then four times in the bucket on the other side of the screen, mimicking the sequence of a typical A-not-B task. Side of initial presentation was counterbalanced across participants.

Although the display shown to the infants was intended to mimic manual A-not-B search, it was not infant-controlled, as is often the case in manual studies. In a manual search task, the toys can continue to be hidden at location A until the infant has reached a success criterion for searching at that location, ensuring that the infant has been attending to, and learning about, the hiding events; however, in the current visual search task, the sequence shown was the same for all infants without any contingency based on where the infant looked during the search period.

Because our task did not require success at location A prior to the switch trial to B, we predicted noticeable differences between subjects depending on whether or not they actually attended to location A during A-trials. Specifically, we predicted that infants who had looked to

location A during A-trials would later continue to search for the keys there, but we did not expect children to learn about hiding events that occurred at a location to which they were not looking. Concomitantly, given that the location of the keys was the only aspect of the scene we presented to infants that varied across the trials, we did not expect to see the same pattern of unlearning in children who had not watched the hiding events.

Results

Analysis confirmed that infants varied in how much they looked towards location A during the initial hiding at A trials, with 17 infants (accurately) looking more at location A, and another group of 15 infants looking more at location B. Looking within the defined ROIs was considered "searching" behavior. An analysis of variance (ANOVA) comparing the searching patterns between these two groups of infants revealed the predicted difference in the patterns of infants' looking across the study, $F(2,90)=34.597$, $p<.001$. Accordingly, the children were separated for remaining analyses: an 'attenders' group of children who looked more to A during the initial search trials, and a 'non-attenders' group who looked more to B during the initial search trials, even though the keys were hidden at location A.

A further omnibus ANOVA, including attending status as a variable, revealed an overall 'side' x 'time' interaction, $F(1,92)=2.622$, $p=.022$, and a 'side' x 'attending status' interaction, $F(1,92)=5.435$, $p<.001$ (Figure 4). These results revealed an overall change in where the infants were looking during the search period across trials, showing that this change was driven by the attenders, who searched first at A and then slowly changed their locus of search over time, as more hiding events occurred at B. Importantly, the non-attenders did not change their searching behavior throughout the study.




Figure 4. A plot of the difference in looking time to A and B across the six trials of the A-not-B task (first two at A, and then four at B) for the two groups, those who searched more at A during trials at that location than those who didn't.

An analysis of the visual search of the attenders revealed a change in looking bias across the trials, with decreasing looking to location A, $F(1,49)=14.057$, $p<.001$. Despite this trend however, there was still a main effect of side in the

study, $F(1,49)=29.468$, $p<.001$, with significantly more looking to A ($M=2038$ ms) than B ($M=1192$ ms), $t(1,50)=4.611$, $p<.001$ (despite the fact that there were twice as many hiding events at B than at A over the course of the study). This finding is noteworthy because it is consistent with the perseverative trend seen in the typical A-not-B task with manual search.

We hypothesized that infants' perseverative search at A would gradually lessen as the cues that predict A were weakened following hiding events at B, but that the rate of learning to search at B would be negligible until the value of the cues directing attention to A were unlearned. Individual differences in rate of the incremental trend away from looking to A in the attenders provided evidence of precisely this pattern of unlearning, with a regression showing that the extent of the searching bias A events "predicted" the extent of bias on early B-trials, $p=.018$, a relationship that was not significant for non-attenders (see Figure 5 for a comparative plot of search behavior at each location for the attenders).



Figure 5. For the attenders, looking time (in ms) to each of the two locations across the six trials of the task.

The non-attenders, who did not learn about location A or the hiding events that occurred there, were not expected to behave in the same way the attenders were. These infants simply showed a main effect of side across the six trials, $F(1,42)=10.979$, $p=.002$, which resulted from more overall looking to location B, $t(1,43)=7.282$, $p<.001$. Given that the non-attending infants initially looked to the side where nothing was happening, and that later hiding events did take place (meaning that B became more interesting as the experiment progressed), it is perhaps unsurprising that the non-attenders did not change their looking bias over the course of the trials in different locations, $F(1,42)=.378$, ns.; moreover, given that the non-attenders do not change their bias over time, it suggests that pattern shown by the attenders is not simply a result of their regressing to the mean (a possible concern, because groups were split based on early search behavior), but rather that these data reflect different patterns of learning over time in the two groups.

Discussion

Children who initially learned about an object hidden at one location continued to search visually at that location even after the object was hidden in a new location, but showed an incremental shift in their search behavior away from the initial location and towards the new location. This pattern of data is consistent with the idea that children have to learn to search, and that unlearning, as a consequence of prediction error, is a key part of that process (Ramscar et al., 2010; Ramscar & Dye, 2009; Rescorla & Wagner, 1972). If infants initially learn that hidden objects are to be found at A, they will perseverate in that response until the balance of evidence (learning) favors the prediction that the objects' will be found at B. The correlation between the attenders' bias during A trials and the early B trials—but not the later B trials—further supports the idea that the initial bias towards A must be unlearned, and that this will happen only as more hiding / appearance events are shown at location B (see also Diedrich, Thelen, Smith, & Corbetta, 2000). This gradual change in looking preference over time is consistent with our hypothesis that search is something children have to learn, and that success or failure at different kinds of search may be, to a degree, a reflection of experience.

While there is much to explain with regards to the development of children's ability to search—and not least how the learning of conjunctive cues over extended trials might impact performance on a modified A-not-B task—we believe that there is insight to be gained from seeing infants' behavior in the A-not-B task in terms of learning to search, and the patterns of behavior that accompany such learning, rather than simply as a failure to search correctly. Not only does this approach offer some insight into the often puzzling behavior of children as they learn to search, but we believe that the combination of eye-tracking and computational modeling methods used in the current study offer a helpful formal framework for other work in this area.

Acknowledgements

This material is based upon work supported by NSF Grant Nos. 0547775 and 0624345 to MR. Many thanks to Daniel Richardson for programming Experiment 1.

References

- Adolph, K. (2000). Specificity of learning: Why infants fall over a veritable cliff. *Psychological Science*, 11, 290-295.
- Aguiar, A., & Baillargeon, R. (2000). Perseveration and problem solving in infancy. In H. Reese (Ed.), *Advances in child development and behavior* (Vol. 27, pp. 135-180). New York: Academic Press.
- Ahmed, A., & Ruffman, T. (1998). Why do infants make A-not-B errors in a search task, yet show memory for the location of hidden objects in a nonsearch task. *Developmental Psychology*, 34(3), 441-453.
- Baillargeon, R., Graber, M., Devos, J., & Black, J. (1990). Why do young infants fail to search for hidden objects? *Cognition*, 36(3), 255-284.
- Barlow H. (2001). Redundancy reduction revisited, *Network*, 12, 241-253

Bell, M. A., & Adams, S. E. (1999). Comparable performance on looking and reaching versions of the A-not-B task at 8 months of age. *Infant Behavior and Development*, 22(2), 221–235.

Diamond, A. (1988). Abilities and Neural Mechanisms Underlying AB Performance. *Child Development*, 59(2), 523-527.

Diamond, A. (1990). The development and neural bases of memory functions as indexed by the AB and delayed response tasks in human infants and infant monkeys. In A. Diamond (Ed.), *The development and neural bases of higher cognitive functions* (pp. 267-309). New York: New York Academy of Sciences Press.

Diamond, A., Cruttenden, L., & Neiderman, D. (1994). AB with multiple wells: 1. Why are multiple wells sometimes easier than two wells? 2. Memory or memory+ inhibition. *Developmental Psychology*, 30(2), 192–205.

Diedrich, F. J., Thelen, E., Smith, L. B., & Corbetta, D. (2000). Motor memory is a factor in infant perseverative errors. *Developmental Science*, 3(4), 479-494.

Hofstadter, M., & Reznick, J. S. (1996). Response Modality Affects Human Infant Delayed-Response Performance. *Child Development*, 67(2), 646-658.

Johnson, M. H., Posner, M. I., Rothbart, M. K. (1994). Facilitation of saccades toward a covertly attended location in early infancy. *Psychological Science*, 5, 90-93.

Marcovitch, S., & Zelazo, P. D. (1999). The A-Not-B Error: Results from a Logistic Meta-Analysis. *Child Development*, 70(6), 1297-1313.

Marcovitch, S., & Zelazo, P. D. (2006). The influence of number of A-trials on 2-year-olds' behavior in two A-not-B type search tasks: A test of the hierarchical competing systems model. *Journal of Cognition and Development*, 7, 477-501.

Munakata, Y. (1997). Perseverative reaching in infancy: The roles of hidden toys and motor history in the AB task. *Infant Behavior and Development*, 20(3), 405–416.

Piaget, J. (1954). *The Construction of Reality in the Child* (M. Cook, Trans.). New York: Basic Books, Inc.

Ramscar, M., & Dye, M. (2009) Expectation and negative evidence in language learning: the curious absence of mouses in adult speech. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, Amsterdam, Netherlands.

Ramscar, M., Yarlett, D., Dye, M., Denny, K. & Thorpe, K. (2010). The Feature-Label-Order Effect in Symbolic Learning. *Cognitive Science*, 34(7).

Rescorla, R. A., & Wagner, A. R. (1972). A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical Conditioning II: Current Research and Theory* (pp. 64-99). New York: Appleton-Century-Crofts.

Rescorla, R. A. (1988) Pavlovian conditioning: It's not what you think it is. *American Psychologist*, 43(3), 151-160.

Smith, L. B., Thelen, E., Titzer, R., & McLin, D. (1999). Knowing in the context of acting: the task dynamics of the A-not-B error. *Psychol Rev*, 106(2), 235-260.

Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, 24(01), 1-34.

Topal, J., Gergely, G., Miklosi, A., Erohegyi, A., & Csibra, G. (2008). Infants' perseverative search errors are induced by pragmatic misinterpretation. *Science*, 321, 1831-1834.

Wellman, H. M., Cross, D., Bartsch, K., & Harris, P. L. (1986). Infant Search and Object Permanence: A Meta-Analysis of the A-Not-B Error. *Monographs of the Society for Research in Child Development*, 51(3).

ⁱ The effects of learning were simulated using the **Rescorla-Wagner model**, a widely used learning rule that has been applied to numerous learning effects in animals and humans, and for which there is strong neurobiological evidence (Waelti, Dickinson & Schultz, 2001; Schultz, 2006). The Rescorla-Wagner model simulates changes in the associative strengths between individual cues and an outcome as the result of discrete learning trials. If the presence of a cue or outcome X at time t is defined as present(X, t), and its absence as absent(X, t), then the predictive value V of a cue i for an outcome O after a learning event at time t + 1 can be stated as:

$$V_i^{t+1} = V_i^t + \Delta V_i^t,$$

while the change (Δ) in the predictive value of i after t can be defined as:

$$\Delta V_i^t = \begin{cases} 0 & \text{if ABSENT}(C_i, t) \\ \alpha_i \beta_1 \left(\lambda - \sum_{\text{PRESENT}(C_j, t)} V_j \right) & \text{if PRESENT}(C_j, t) \ \& \ \text{PRESENT}(O, t) \\ \alpha_i \beta_2 \left(0 - \sum_{\text{PRESENT}(C_j, t)} V_j \right) & \text{if PRESENT}(C_j, t) \ \& \ \text{ABSENT}(O, t) \end{cases}$$

Learning is thus governed by a **discrepancy function** where λ_j is the value of the predicted event (the maximum amount of associative strength that an event j can support) and V_j is the predictive value for j given the set of cues present at time t.

In trials in which there is **positive evidence** – i.e., in which expected outcomes do occur – the Rescorla-Wagner learning rule produces a negatively accelerated learning curve (the result of events being better predicted, which reduces the discrepancy between what is expected and what is observed) and asymptotic learning over repeated trials (as events become fully predicted). Conceptually, this happens because the model embodies the idea that the function of learning is to align our expectations with reality, and the better that alignment becomes over time, the less we need to learn.

In trials in which there is **negative evidence** – i.e., in which an expected outcome fails to occur – λ_j (the expected outcome) takes a value of zero because it didn't occur. In such cases, the discrepancy function ($\lambda_j - V_j$) produces a negative value, resulting in a reduction in the associative strength between the cues present on that trial and the absent outcome j. Conceptually these prediction errors can be thought of as violations of expectation that allow the model to learn from the negative evidence.

The total amount of predictive (cue) value any given outcome can support in learning is finite. (Informally, we can think of this as capturing the idea that if predictive confidence keeps rising, it must eventually reach a point of certainty.) As a result, cues compete with one another for relevance, and this produces learning patterns that often differ greatly from those that would arise by simply recording the correlations between cues and outcomes (i.e., simply tracking base rates – a common misconstrual of learning; Rescorla, 1988).

Notably, the amount of learning that occurs on any given trial is determined by two factors: the overall learning rate β_j (where $0 \leq \beta_j \leq 1$), and the individual saliency of cues, denoted by a parameter α_i (where $0 \leq \alpha_i \leq 1$). These parameters establish the rate at which the discrepancy between λ_j and V_{TOTAL} reduces. In the simulation we conducted, we set $\lambda = 100\%$ for a location when the keys are visible there, or 0% when not visible, $\alpha_i=0.05$ for attended stimuli and 0.075 for unattended stimuli, and $\beta_j=0.5$. The model is similar to, and consistent with, models of representational strength of latent versus active memory traces (e.g., Munakata, 1998).