

Too much of a good thing: How novelty biases and vocabulary influence known and novel referent selection in 18-month-old children and associative learning models

Sarah C. Kucker

Department of Psychology, The University of Wisconsin Oshkosh; The DeLTA Center

Bob McMurray

Department of Psychological and Brain Science, The University of Iowa; The DeLTA Center

Larissa K. Samuelson

School of Psychology, The University of East Anglia; The DeLTA Center

#### Author Note

This research is supported by R01HD045713 from the Eunice Kennedy Shriver National Institute of Child Health and Human Development to LKS. The content is solely the responsibility of the authors. We wish to thank the members of the Language and Categorization Lab, John Spencer, Prahlad Gupta, and Karla McGregor for helpful discussions.

This research was part of the first author's doctoral dissertation submitted to the University of Iowa. Sarah Kucker is now at the University of Wisconsin Oshkosh, Department of Psychology. Address correspondences to Sarah C. Kucker, the University of Wisconsin Oshkosh, Department of Psychology, 800 Algoma Blvd, Oshkosh, WI 54901 email: [kuckers@uwosh.edu](mailto:kuckers@uwosh.edu)

Keywords: word learning, novelty, referent selection

## Abstract

Identifying the referent of novel words is a complex process that young children do with relative ease. When given multiple objects along with a novel word, children select the most novel item, sometimes retaining the word-referent link. Prior work is inconsistent, however, on the role of object novelty. Two experiments examine 18-month-old children's performance on referent selection and retention with novel and known words. The results reveal a pervasive novelty bias on referent selection with both known and novel names and, across individual children, a negative correlation between attention to novelty and retention of new word-referent links. A computational model examines possible sources of the bias, suggesting novelty supports in-the-moment behavior but not retention. Together, results suggest that when lexical knowledge is weak, attention to novelty drives behavior, but alone does not sustain learning. Importantly, the results demonstrate that word learning may be driven, in part, by low-level perceptual processes.

Keywords: word learning, novelty, referent selection

**Too much of a good thing: How novelty biases and vocabulary influence known and novel referent selection in 18-month-old children and associative learning models**

Word learning is a difficult problem that young children appear to solve with ease. Toddlers can successfully get their new *banjo*, point to the *pangolin* in a book, or choose between a familiar food and the new *falafel* despite minimal experience with these words and concepts. Over the past half century, hundreds of studies with infants through adults have investigated the processes driving this behavior, commonly referred to as disambiguation, fast-mapping, or referent selection. Studies have examined attentional biases (e.g. Mather & Plunkett, 2012; Samuelson & Smith, 1998, 2000), novelty<sup>1</sup> (e.g., N3C; novel-name name-less category principal; Mervis & Bertrand, 1994), lexical knowledge (Halberda, 2003; Markman & Wachtel, 1988), and social/pragmatic cues (Akhtar, Carpenter, & Tomasello, 1996; Diesendruck & Markson, 2001). By many accounts, word learning requires an interaction of processes (Hollich et al., 2000). For example, it is well known that children can select a novel item over a known one when given a novel label (so-called fast-mapping). Yet even this simple choice has been hypothesized to rely both on knowledge of the known foils as well as recognition that the target item and word are novel (Golinkoff, Mervis, & Hirsh-Pasek, 1994; Markman & Wachtel, 1988; Mather, 2013a; Mervis & Bertrand, 1994).

Attempts to understand this trade-off in word learning are complicated by two factors. First, the relative novelty and familiarity of stimuli change continuously as exposure to words and objects increase and knowledge of their lexical mappings expands. Words and objects are perceived as novel when they are *relatively* less familiar than known words and objects. That is, early word learning draws on a child's lexical knowledge (or lack of), but also on their reaction

---

<sup>1</sup> Here and elsewhere, we use novelty to refer to the visual familiarity with the object (regardless of name).

to novelty. Second, children's responses to novel and familiar stimuli change over learning (Kidd, Piantadosi, & Aslin, 2012), and development (Fantz, 1964; Wetherford & Cohen, 1973). Thus, to understand the role of attention to novelty in referent selection and word learning, we must understand how the presence of a novel word or object impacts both novel *and* familiar word and object recognition.

The studies and model simulations presented here extend recent work on 24-month-old children's referent selection and retention to 18-month-olds, with a focus on children's attention to novelty. Even this modest age difference is enough to reveal surprising changes in children's responses to known words in the context of novel objects, and suggests a new possible source of children's bias towards novelty in early vocabulary development.

### **Novelty biases over development**

Children preferentially attend to novel over familiar stimuli as early as 7-months gestation (Sandman, Wadhwa, Hetrick, Porto, & Peeke, 1997) and throughout infancy (Fantz, 1964). By one year, novelty preferences are complex, influenced by exposure time (Roder, Bushnell, & Sasseville, 2000; Wiebe et al., 2006), competition, (Hunter, Ross, & Ames, 1982; Rose, Gottfried, Melloy-Carminar, & Bridger, 1982), and modality/stimulus characteristics (Cohen, DeLoache, & Rissman, 1975; Slater, Morison, & Rose, 1984), suggesting increasing sophistication in how children process and attend to novelty.

Novelty can also bias children's attention in ways relevant for word learning. Infants' responses to novel objects are influenced by auditory input (Bahrick, Lickliter, & Flom, 2004; Lewkowicz, 1996) and habituation to novel objects slows when a novel label is added (Mather, Schafer, & Houston-Price, 2011). Likewise, novel labels (but not known ones) drive 10-month-olds to preferentially look to novel over familiar objects (Mather & Plunkett, 2010); this referent

selection behavior becomes more consistent by 17-months (Halberda, 2003).

At 24-months, changes in visual novelty have been shown to alter children's referent selection and retention. Horst, Samuelson, Kucker, and McMurray (2011) showed that 24-month-old children's ability to select a novel object in response to a novel name was, to some extent, driven by attraction to novelty. They used a pre-exposure period to manipulate the visual familiarity of some of the objects presented in a referent selection task. When two-year-olds heard a novel word in the context of two pre-familiarized (though name-unknown) novel objects, and one never-seen novel object, children selected the never-seen object even though none of the objects had known names (see also Mather & Plunkett, 2012). This suggests a unique role for object novelty during the first two years (Mather, 2013b), though it remains likely that novelty interacts with lexical knowledge.

### **The growing influence of lexical knowledge**

Children's ability to identify familiar words and objects also changes with increasing exposure to them over development. One-year-old infants look at toys longer in the context of their known labels (Baldwin & Markman, 1989), but as children age (Bion, Borovsky, & Fernald, 2013) and their word knowledge increases (Borovsky, Ellis, Evans, & Elman, 2015), real-time word recognition improves. Over the second year, children access familiar words increasingly efficiently (Fernald, Pinto, Swingley, Weinberg, & McRoberts, 1998; Fernald, Perfors, & Marchman, 2006). This suggests there is a growing ability to rapidly build activation for familiar words and their referents (and see Rigler et al., 2015; Sekerina & Brooks, 2007 for change through adolescence).

In addition to changes in familiar word and object recognition, there is also evidence that knowledge about the items used as foils (in disambiguation tasks) can affect attention to a novel

target. For instance, Halberda (2006) demonstrated that 14-month-old infants look primarily to name-known items both before and after hearing a novel word, suggesting that known objects can draw gaze, even when a heard word is novel. Older 17-month-old infants also look first to name-known competitors, but eventually settle on the novel referent (Halberda, 2006). Borovsky et al. (2015) found that at 24-months, referent selection performance for novel items from well-known semantic categories is higher than those from weakly known categories. Similarly, a child's level of familiar word knowledge predicts novel referent selection in both toddlers and preschool-age children (Grassmann, Schulze, & Tomasello, 2015; Merriman & Schuster, 1991). Thus, novelty likely plays an increasingly weaker role as children age and their selection processes begin to rely on name-knowledge to a greater extent (Nazzi & Bertoncini, 2003).

### **Disentangling novelty and knowledge**

The literature clearly demonstrates an influence of both object novelty and name-knowledge on how children respond to novel words (Golinkoff et al., 1994; Mather, 2013b), and the role of both factors in real-time referent selection also shifts over development (Nazzi & Bertoncini, 2003). Even as children's ability to recognize familiar words develops, the factors driving their ability to rapidly identify the referents of *novel* words in disambiguation or fast-mapping type situations also changes. However, the exact nature of these changes is unclear. What is particularly less well understood is the impact of object novelty on recognition of *familiar* words early in vocabulary acquisition.

Computational models have lent insight to this issue as they can control and measure the strength of word-referent links. One simple model of children's ability to select the referent of a word posits a form of competition in which candidate referents vie for attention after a novel word is heard (McMurray, Horst, & Samuelson, 2012; see also MacWhinney, 1987). Under this

framing, it is possible that either novelty or knowledge could lead to the selection of a given object as the winner. More robust knowledge of a word-referent link could allow children to access the names of familiar referents more efficiently, which would also allow them to be more easily dismissed as referents of the novel name (c.f., Halberda, 2006). Conversely, when a novel word fails to activate a known referent, attention to novelty may fill the gap. Such models, then, propose novelty (lack of familiarity) is an inherent component to word learning. This is similar to prior theories, such as N3C, that propose either that individual objects/words selectively draw attention because of their novelty.

More recently, modeling work, such as that by McMurray and colleagues (2012), has suggested the possibility that attraction to novelty might be related to the state of the child's entire lexical system (rather than individual associations). Under this framing, an individual's lexicon consists of a series of connections or associations between words and referents and learning is the process of adjusting those connections or associations. Over time, then, correct connections between words and referents are strengthened and incorrect, spurious, connections are weakened resulting in the overall network becoming more refined. In later states of learning/development activation of a word will efficiently lead to activation of the associated referent. However, in the early stages of development, the lexical network will be unrefined, with many weak connections and with incorrect connections that need to be removed. In this state, activation of a word will not efficiently lead to activation of a single associated referent, and may instead lead to several possible referents being weakly activated, or to the activation of seemingly unrelated referents via spurious connections that have not yet been pruned. Such an unrefined lexicon may not be well organized enough to drive consistent responses (even to familiar words), leading toward behavior that is biased toward novelty.

These theoretical ideas fit with prior empirical work on novelty that has emphasized what a child selects in response to a *novel* word. However, such empirical work does little to disentangle the basis for selection because in such contexts both lexical knowledge and the novelty of the object support the same choice. Consider a child who is presented with a well-known *cup* and a novel object and asked to select the *cheem*. The child can correctly choose the novel object because she already knows the label for *cup* and does not yet have a label for the novel item (nor a referent for *cheem*). Or, the child can correctly choose the novel object because it is exciting and new, thus grabbing attention. Thus, novel selection trials do little to differentiate what is driving the child's choice. Consider instead if the child were asked to select the cup. Here, if she relies on lexical knowledge, she will correctly select the *cup*. However, if she relies on novelty, she will incorrectly select the novel item. As this example illustrates, an investigation of responding to both familiar and novel words in the same task may help isolate the role of novelty, and identify how it plays a role throughout lexical processing.

Relevant prior work exists, but differences in tasks and inconsistent findings obscure a clear conclusion. Samuelson and colleagues (Horst & Samuelson, 2008; Kucker & Samuelson, 2012; see also Bion et al., 2013) tested responses to novel words in the presence of both novel and known objects. As typically observed, with a novel word, 24-month-old children avoided familiar items and selected the novel object. These studies, however also included trials testing the response to known words, using the same array of novel and known items. On these *known* referent selection trials, children were able to *avoid* the novel, name-unknown item. Thus, by 24-months children are often able to overcome the pull to novelty.

A handful of prior referent selection studies with younger children have also contained requests for both familiar/known and novel objects across trials. Bion and colleagues' (2013)



looking task with 18-month-old children included known referent selection trials, but the results of these trials were not reported separately. White and Morgan (2008) found correct preference for the known-name objects in the context of novel foils with 19-month-old children. However, this preference also existed in an initial no-name portion of the trial, making it hard to determine if it was an overall object preference or driven by lexical knowledge. In a canonical study, Mervis and Bertrand (1994) examined the N3C principle in 14- to 17-month-old children but children were corrected and praised after their selection, potentially biasing later trials. Thus, it is not clear from the extant literature that 18-month-old children can reliably overcome a bias towards novelty to select the referent of a familiar word. Consequently, it is not clear that good performance on novel word selection trials is not simply an attraction to novelty.

Further, work with even younger children has shown that novelty may disrupt known word processing - Mather and Plunkett (2010) found that though 10-month-old infants did eventually settle on the correct known referent by the end of a preferential looking trial, they displayed a slight initial preference for novel objects across all trials (novel, known, and non-label control). Likewise, in Halberda (2003), 14- to 17-month-old children seemingly maintained a slight preference for novel objects (Halberda, 2003 Fig. 1) but no details on mean accuracy or trial-by-trial performance is given. Taken together, prior results suggest novelty influences responses to both known and novel words during the second year, although differing experimental procedures and a blending of results across trials types obscures firm conclusions.

In summary, by 24-months, children appear to behave in a way that is consistent with the relevant information in the task: they accurately respond to known, familiar words despite the presence of a novel object (Bion et al., 2013; Horst & Samuelson, 2008; Kucker & Samuelson, 2012). In line with this, models of 24-month-old word learning can simulate behavior without

any additional cues toward novelty or knowledge (e.g. McMurray et al., 2012). However, given changes in novelty biases and the strength of internal representations by which familiar words are recognized, responses to familiar words may differ earlier in development. Thus, the current study uses referent selection with name-known and novel objects in the same forced-choice task in order to examine the roles of visual novelty and lexical knowledge in younger children and a formal model. We examine 18-month-olds—a critical age at which children’s vocabulary is rapidly growing (Fenson et al., 1994), but known word knowledge still relatively fragile (Fernald et al., 1998) and thus less likely to out-compete novel objects. If referent selection largely derives from novelty, children should choose the novel object regardless of name, even when a familiar word is the target. However, if lexical knowledge is the driver, children should respond appropriately to both types of trials. This may also require some rethinking of how computational models of word learning treat referent selection early in development.

Of course, selecting a referent is only part of word learning. Recent work suggests children do not always retain word-referent links from referent selection (Horst & Samuelson, 2008); rather retention is not reliable until 30 months of age (Bion et al., 2013; Spiegel & Halberda, 2011). Thus, long-term learning does not occur instantly, but develops over a longer timescale (Carey & Bartlett, 1978; McMurray et al., 2012; Swingley, 2010). Prior work further shows that retention is influenced by several factors, including age (Bion et al., 2013), repetition (Goodman, McDonough, & Brown, 1998; Axelsson & Horst, 2014), praise (Mervis & Bertrand, 1994), and number of competitors (Horst, Scott, & Pollard, 2010). This suggests that though initial exposure to a word is only a small step in learning, the context in which a word is initially encountered may have cascading effects on learning (see Kucker, McMurray, & Samuelson, 2016; McMurray, Horst, & Samuelson, 2012). That is, the role of novelty in referent selection

during novel word mapping may also impact retention. Thus, experimental tests should include both novel and known labels, and examine mapping *and retention* of each.

Finally, computational models of word learning have not fully explored the space of possible mechanisms by which novelty may play a role in referent selection. If an isolable effect of novelty on referent selection is found, it is important to examine the mechanisms that might achieve such results in a model. In this regard, the model of McMurray et al. (2012) offers a useful start as it treats real-time decision and attentional processes (where novelty most likely has a role) as distinct from long-term learning. It has also already been shown to model both referent selection (both familiar and novel words) and retention. However, even within this model there are multiple possible loci for novelty. Here we examine two: 1) an attentional bias derived from a system biased toward individual novel or unknown items, or 2) a system-based bias derived from an unrefined lexicon that is too weak to support use of prior knowledge and thus, displays a global preference for unknown word-referent links.

### **Experiments**

Experiment 1 examines referent selection with known and novel words as well as retention of novel word-referent links in 18-month-old children. On referent selection (RS) trials, we examined how children respond to a known word/object in the context of a novel object (*known RS*) and to a novel word/object in the context of name-known objects (*novel RS*, classic fast-mapping or disambiguation). After a short delay, we tested retention of novel word-referent links. This procedure directly compares name-unknown/unfamiliar and name-known/familiar trials, and also assesses retention. We address three questions.

First, we asked whether the presence of a novel foil item disrupts familiar word recognition in these younger children via known RS trials. Children should perform well, given

they know the words. However, an enhanced role for visual novelty early in development could induce more competition from novel foils, reducing accuracy even for known words, especially early in vocabulary growth. Second, we ask whether 18-month-olds reliably select a novel object in response to a novel word via novel RS trials. We would expect both object novelty and the lack of a known label to support selection of the novel object, consistent with similar work in looking-based tasks (e.g. Bion et al., 2013; Halberda, 2003) and in tasks in which children were corrected and praised for their selections (Mervis & Bertrand, 1994). Third, we ask whether 18-month-olds demonstrate retention of these new mappings. While there are not strong reasons to suspect that 18-month-olds will retain in a task where 24-month-olds did not, this test provides an important complement to prior work.

Experiment 2 adds a manipulation of item familiarity (Horst et al., 2011; Kucker & Samuelson, 2012) to ask if the bias to novelty is tied to specific items (in which case it should be reduced) or reflects a broader developmental state (which would not be affected by short exposure to individual items). Combined analyses then examine how variation in retention among individual children is related to performance during the learning phase. This complements prior work in examining what aspects of the initial selection impact retention. Finally, our computational simulations examine two possible underlying mechanisms that may support a bias toward novelty. We examine both an exogenous bias specifically toward novel visual stimuli, and an endogenous bias that is product of an early under-developed lexical system.

## Experiment 1

### Methods

**Participants.** Thirty-two, 18-month-old children (16 females,  $M = 18:24$ ; range 17:22-

19:25) from primarily middle-class families in a small Midwestern town participated. They had a mean productive vocabulary of 77.5 words (range 1-375). Data for an additional eight children were dropped due to a failure to complete the majority of trials by refusing to participate, crying, or being noncompliant (6), parent/sibling interference (1) or an incomplete vocabulary questionnaire (1). Children received a small prize for participating.

**Stimuli.** For each child, three known items were randomly selected from a pool of 16 items whose labels are known by most 18-month-olds (Fig. 1, left). Parents confirmed that their child knew the names of these items prior to the task, and items were replaced as necessary. Eight unknown items were novel toys with which 18-month-olds are not likely familiar (Fig. 1, right). These items are similar to items rated by adults to have low familiarity (< 34% familiarity rating) in an online database (Horst & Hout, 2016) with the exception of two: a ball-and-cup toy and a “worm ball” that are more commonly known by adults. However, because familiarity to adults and 18-month-olds are not equivalent, and what is novel can vary between individuals, parents confirmed the novelty of all the unfamiliar items prior to the task with familiar items being replaced to ensure that children had not previously been exposed to any novel or similar items. Four novel words were randomly selected from a pool of eight phonologically legal CVC words that had no known referents (*cheem, dite, fode, lorp, pabe, stad, roke, yok*).

-----Insert Figure 1 about here -----

**Procedure and design.** The child was seated across a table from the experimenter in a booster seat next to their parents or on their parents’ lap. Parents completed the MacArthur-Bates Communicative Development Inventory: Words and Sentences (MCDI-WS; Fenson et al., 1994)

during the session and were instructed to avoid interacting with their child, offering minimal encouragement only if needed.

**Warm-up/comprehension.** The procedure began with three warm-up/comprehension trials. Three name-known items were presented on a white tray in a horizontal row, equally spaced. The experimenter arranged the items out of sight, then, while maintaining eye contact, placed the tray within sight, but out of reach of the child. After three seconds, the experimenter prompted the child to retrieve one name-known item by name (e.g. “Can you get the *shoe*?”) and pushed the tray forward. The child was re-prompted if necessary and praised and/or corrected as needed. Each object was the target only once. Target location and trial order were randomized.

**Referent selection.** Eight referent selection trials immediately followed warm-up. On these trials, children saw three objects – two name-known and one novel – and were asked to retrieve an object by name. The name-known items were the same three from warm-up. This was to ensure that children were familiar with the known items and corresponded to prior work (e.g. Kucker & Samuelson, 2012). The novel item on each trial was randomly selected from the pool. On half the trials, children were asked for a name-known item (“Can you get the *shoe*?”). The same item was never requested on two consecutive known RS trials, and no item was requested more than twice. On the other four novel RS trials, children heard a novel word (“Can you get the *roke*?”). Across trials, novel items were present only once and object locations and trial order were randomized. Children were re-prompted up to three times if needed, but without correction or praise.

**Break.** There was a five-minute break immediately following referent selection during which children played in an adjacent waiting room or colored quietly in the experiment room. No experimental stimuli were present and no novel names were heard.

**Retention.** A single warm-up trial using the same three known items followed the break. Two retention trials immediately followed. Children saw two novel objects that had previously been targets on novel RS trials along with a third novel object that had been presented as a foil during a known RS trial. Thus, all three had been seen an equal number of times and two had potentially been mapped to a novel word. Items and prompts were not repeated across trials.

**Preference trials.** Two no-label preference trials with items from the task followed retention. During these, three items were presented: 1) one unnamed novel foil item from a known RS trial (which was not used on retention); 2) a novel target item from a novel RS trial (which was used as a foil during retention); and 3) one novel item that was the target during both novel RS and retention. Thus, the amount of the child's prior exposure to each varied, as did their prior selection and use as a target. Items were placed on the tray as before and children were asked "Can you get one?" These trials provided a measure of children's preference for novelty without labels.

### **Coding**

Naïve coders indicated children's final selections off-line via video recordings. Data from 11 random subjects (61.1%) were re-coded for reliability. Inter-coder agreement was nearly 100%. One discrepancy was settled via discussion with a third coder.

### **Results and discussion**

On the warm-up/comprehension trials, children ultimately selected the target with 94.8% accuracy ( $SD=.15$ ; Fig. 2), significantly above chance (33%),  $t(31) = 23.4, p < .0001, d = 4.14$ . This is not surprising given that children were corrected and praised on these trials. Very few children, however, required praise or correction and children easily selected the correct known items.

In contrast, children performed surprisingly poorly on the known RS trials, with only 30.7% (SD=.31) accuracy, a level no different from chance,  $t(31) = -.42, p = .68, d = -.07$ . This is startling given that 1) the target words are known by most 18-month-old children, 2) parents confirmed their child's knowledge of these words, and 3) children accurately selected these same referents on the preceding warm-up/comprehension trials. Rather than selecting the target, children selected the *novel foil* 70.3% of the time (SD=.31),  $t(31) = 6.70, p < .0001, d = 1.19$ . Thus, children's knowledge of the known words was overridden by a novelty bias. This novelty bias was maintained throughout referent selection with no differences in on the first two versus last two known RS trials ( $M_{\text{first}} = 30.2\%$ , SD=.38;  $M_{\text{last}} = 31.3\%$ , SD=.37), between groups  $t(30) = .59, p = .56, d = .11$ .

On novel RS trials, children selected the novel (name-unknown) target on 77.9% of trials (SD=.27), significantly above chance,  $t(31) = 9.40, p < .0001, d = 1.66$ . This is higher than known RS performance,  $t(31) = 5.64, p < .0001, d = 1.63$ , and close to prior work (e.g., Markman & Wachtel, 1988). Given the dominant attention to novelty on known RS trials, this excellent performance is not surprising.

Retention of a new word-referent link after one exposure is difficult for children under 30-months (Bion et al., 2013; Horst & Samuelson, 2008). However, if novelty acts as a goal-driven learning strategy, then heightened attention to novelty might translate to better learning and subsequently better retention. The data do not support this. As a group, participants failed to retain the new word-referent links, selecting the target at 32.8% accuracy (SD=.36), no different from chance,  $t(28) = -.04, p = .97, d = -.01$ .<sup>2</sup> Thus, even with a strong bias to choose the most

---

<sup>2</sup> Only trials for children who mapped the target item during the respective novel RS were included. This is consistent with prior work (Horst & Samuelson, 2008), and ensures children had initial associations to recall. Including the trials from the three children who did not respond accurately on novel RS does not affect the results (34.38%, SD=.35),  $t(31) = .225, p = .82, d = .04$ .



novel object during referent selection, 18-month-olds failed to retain these new word-referent links. Instead, children chose the unnamed foil items at above chance levels 48.4% of the time ( $SD=.35$ ),  $t(31)=2.51$ ,  $p=.0008$ ,  $d=.44$ , and the other named item below chance 17.2% of the time ( $SD=.24$ ),  $t(31)=-3.71$ ,  $p=.02$ ,  $d=-.66$ . This suggests novelty may not be a key for word *learning*, but rather a bias affecting in-the-moment behavior.

-----Insert Figure 2 about here -----

Finally, we examined the preference trials. These suggested that children maintain preferences for novelty. Children chose the novel item (unnamed items seen once as a known RS foil) on 54.69% of trials ( $SD=.37$ ), significantly above chance (33%),  $t(31) = 2.57$ ,  $p = .02$ ,  $d = .59$ . This confirms children are biased to select to the most novel, name-unknown object regardless of prompt, even outside of word learning contexts.

These analyses confirm the strong pull of novelty on 18-month-old's referent selection. When novel words are heard, this leads to good (though perhaps misleading) performance. However, when familiar words are heard, this leads to near-chance performance. Children also failed to retain the novel word-referent links, suggesting the novelty bias did not strengthen the association between newly encountered words and objects. This pattern of performance – selection of the novel object on both name-known and novel name trials—stands in contrast to prior findings and conflicts with prior theoretical accounts. Accurate novel RS performance has been interpreted as a logically-based use of prior knowledge to disambiguate the referent of a novel word (Golinkoff et al., 1994; Mervis & Bertrand, 1994). When this good performance on novel RS is contrasted with their poor known RS performance, however, 18-month-old's

behavior seems less likely a result of information-laden strategies and more likely a result of low-level biases.

Nevertheless, novelty is a relative concept; a single item may be perceived as very novel in comparison with well-known items, but less novel in the context of other unknown items. Likewise, an item that has been seen previously is less novel than one which has never been seen. Thus, it is possible children's selection of the novel objects on name-known trials was based on a fleeting property of the novel object - that it was simply much more unfamiliar than the contrasting known items. It is difficult to diminish the familiarity of the known items (but see Samuelson, Kucker, & Spencer, 2017), but prior work has shown that the effect of novel items changes when children are given even minimal time to explore them prior to a task (Kucker & Samuelson, 2012). Thus, we may be able to reduce 18-month-old children's novelty bias in this task with brief pre-familiarization to the novel items, thereby making them more familiar. Such a familiarization period could reduce the novelty bias in referent selection, improving performance and help 18-month-old infants reach above-chance performance. If 18-month-old children continue to perform poorly on Known RS, the familiarization period would then provide evidence for a child-based pervasive novelty bias instead of an object-specific mechanism driving behavior. These possibilities are tested in Experiment 2.

## **Experiment 2**

Children in Experiment 1 overwhelmingly chose the novel object even when the task was to ignore it in favor of the familiar name-known object. This is surprising given that prior work with similar tasks showing 24-month-old children consistently select the correct referent for familiar words in forced choice selection (Horst & Samuelson, 2008; Kucker & Samuelson, 2012), and looking-based tasks (with even younger children; Bion et al., 2013; Mather &

Plunkett, 2010; White & Morgan, 2008). One explanation for the novelty bias observed here is that the items were so perceptually engaging that children had a hard time disengaging from them during the referent selection trials. One way to reduce both the relative novelty of the unknown items is to familiarize children to the novel objects. Prior work has shown that very short periods of familiarization can not only change attention to an item (Fantz, 1964), but can dramatically change performance in referent selection tasks. Kucker and Samuelson (2012) gave 24-month-old children one minute of familiarization time with the novel objects prior to a referent selection task that pitted two name-known objects against a novel object. Contrary to the findings of Horst and Samuelson (2008), in Kucker and Samuelson (2012) 24-month-old children demonstrated significant retention of the novel word-referent links that included these pre-familiarized objects. Thus, it is possible that such familiarization would reduce the contrast between known and novel items on each trial, decreasing the novelty of unknown target items relative to name-known items (see also Horst et al. 2011), thereby increasing performance on the known word trials and potentially boosting retention (Kucker & Samuelson, 2012).

However, it is also possible that a specific object novelty bias is not the cause of the poor performance with familiar words. Rather, the 18-month-olds' novelty bias may be a marker of a relatively unrefined (immature or less organized) lexical system along the lines suggested by McMurray et al. (2012). Because the network of word/object associations is still unrefined, these young children's knowledge of the *known* competitor items may be strong enough to be brought to bear in the task. If children have a hard time acting in-the-moment based on their weak knowledge of the names for "known" items, other factors like novelty may determine the ultimate response (Mather, 2013b; Mervis & Bertrand, 1994). In this case, then, a brief pre-exposure to novel items will not change the strength of known word-object associations, leading

to continued poor performance.

## Methods

**Participants.** A new group of twenty-seven 18-month-old children from the same community participated (12 females;  $M = 18:15$ ; range = 17:12-19:90). They had a mean productive vocabulary of 69.22 words (range = 4-301, median = 26 words), which was not significantly different from the vocabulary of children in Experiment 1,  $t(57) = .40$ ,  $p = .69$ ,  $d = .03$ . Data for two additional children were dropped due to a failure to complete the majority of trials by refusing to participate, crying, or being noncompliant (2).

**Stimuli.** Stimuli were identical to Experiment 1.

**Procedure and design.** The procedure was identical to Experiment 1 but with the addition of a familiarization period. Prior to warm-up, children were given the eight novel objects, in two sets of four, to explore silently. The experimenter did not label items and interacted minimally. If a child did not engage with an item, the experimenter drew the child's attention to it by picking it up and saying "Look!" Once the child had explored each item, and the familiarization time had passed, or when the child became fussy and disengaged, the experimenter removed all the items and continued as in Experiment 1.

Prior studies with pre-familiarization periods (Horst et al., 2011; Kucker & Samuelson, 2012) used one minute as a standard length. However, because these children were younger and because we thought the novelty bias from Experiment 1 would require more time to diminish, we began by giving children five minutes of familiarization ( $n = 9$ , 3 female). However, most children did not engage with the stimuli for this long - only one child used the entire five minutes and children averaged 2-3 minutes before disengaging from the task, suggesting object novelty was no longer sustaining attention. To ensure children remained engaged with the task, and to

keep the time consistent with prior work, the remaining 18 children (9 female) had a shortened familiarization period of roughly one minute per set, in line with the prior studies (Horst et al., 2011; Kucker & Samuelson, 2012).

### **Coding**

Naïve coders timed each child's familiarization period and coded children's final selections off-line via video recordings of the entire session. Choice data from 18 subjects (66.6%) were re-coded for reliability purposes. Inter-coder agreement was 100%.

### **Results and Discussion**

We first examined exposure time during the familiarization period. The nine children with the longer familiarization time averaged 250.8 seconds (range 185.4-356.9) of exposure to the first set of objects, and a marginally shorter 224.0 seconds (172.8-302.5) to the second,  $t(8) = 2.17, p = .06, d = .73$ . The 18 children with a shorter familiarization time averaged 118.6 seconds (range 64.2 seconds – 228.2 seconds) for the first set and 104.9 seconds (54.4 seconds -154.9 seconds) to the second,  $t(17) = 1.00, p = .33, d = 0.24$ . This suggests that the shorter familiarization time was appropriate in order to ensure both sets were treated equally.

We next asked how familiarization time correlated with performance on referent selection. Pearson's correlations between familiarization time and performance were near 0;  $r(25) = 0.0084, p = .65$  for warm up/comprehension;  $r(25) = .006, p = .69$  for known RS;  $r(25) = .0005, p = .91$  for novel RS; and  $r(24) = .007, p = .69^3$  for retention. Thus, for the rest of the analyses, all 27 children are analyzed together.

As in Experiment 1, children performed well during the warm-up/comprehension trials, eventually selecting known targets with 97.50% accuracy (SD=.09) after correction and praise,

---

<sup>3</sup> Only trials mapped on novel RS were included. If all children were included, the results were similar;  $r^2(25) = .016, p = .52$ .

significantly above chance  $t(26) = 12.56, p < .0001, d = -1.24$ . Generally, children did not need much correction or praise with the majority of children engaging with the stimuli correctly on the first warm-up/comprehension trial. However, on known RS trials, children were only 42.60% accurate ( $SD=.33$ ), at chance,  $t(26) = 1.51, p = .14, d = .29$ , and not different from Experiment 1,  $t(57) = 1.42, p = .16, d = .37$ . As in Experiment 1, children instead selected the novel foil object the majority of the time (53.7%,  $SD=.34$ ), above chance,  $t(26) = 3.19, p = .002, d = 1.25$ . Given the pre-familiarization trials, this suggests known RS performance is not a response to the unfamiliarity of these specific items.

On the novel RS trials, children performed well, selecting the target 76.9% of the time ( $SD=.23$ ),  $t(26) = 9.94, p < .0001, d = 1.91$ . This is not different from Experiment 1,  $t(57) = .84, p = .41, d = .90$ , suggesting that pre-familiarization had no influence on a child's ability to attend selectively to the novel target when asked.

In addition, pre-familiarization did not affect retention. Children selected the target on 40.4% of trials ( $SD=.42$ ), not different from chance,  $t(25) = .89, p = .38, d = .17$ , and not different from Experiment 1,  $t(53) = .72, p = .48, d = .08$ . Thus, relative to Experiment 1, at no point during either referent selection or retention did pre-familiarization alter children's performance<sup>4</sup>. However, children here no longer preferentially selected the unnamed foil item instead of the target, selecting it only 42.6% of the time ( $SD=.41$ ), no different from chance,  $t(27)=1.22, p=.23, d=.23$ . Children selected the other named item below chance levels, 16.7% of the time ( $SD=.24$ ),  $t(27)=3.53, p=.0016, d=-.68$ . Thus, pre-familiarization appears to have altered a child's attraction to specific unnamed foil items.

Similarly, pre-familiarization significantly altered the preference trials. Children no

---

<sup>4</sup> Like Experiment 1, analysis was done across children who initially mapped on novel RS trials. If the one child who failed to select the target is included, the results hold, 40.7% ( $SD=.37$ ),  $t(26) = 1.09, p = .28, d = .21$ .

longer demonstrated a strong preference for the most novel item, instead choosing all items at chance levels: unnamed known RS foil 40.7% ( $SD=.42$ ),  $t(26) = .96$ ,  $p = .34$ ,  $d = .19$ ; items only named during novel RS 33.3% ( $SD=.34$ ),  $t(26) = .05$ ,  $p = .96$ ,  $d = .01$ ; and items named during both RS and retention 25.9% ( $SD=.29$ ),  $t(26) = -1.3$ ,  $p = .22$ ,  $d = -.24$ .

As a whole, this suggests that item-level novelty is not what is driving poor performance with the familiar words. Even after to-be-learned objects have been familiarized to the point that children no longer show a novelty bias, we still observed marked decrements in familiar word performance when these objects were present as a foil. This suggests that the bias to choose novel objects on these trials may not reflect a reaction to the novelty of a single item. A remaining hypothesis then is that an overall unrefined network of word-referent links across the lexicon may make it hard to rely on word knowledge, and leave the system resorting to novelty during this word learning task. We test this idea through a combined analyses and computational model.

### **Combined Analysis**

Across both experiments, children varied in the novelty bias exhibited during RS. Some selected the novel item on 100% of trials (failing on all known RS trials), whereas a few were 100% accurate during known RS trials. Because children at this age show wide variability in language learning abilities (Fernald & Marchman, 2012) and differences in attention during initial exposure can predict differences in learning (Smith & Yu, 2013; Yu & Smith, 2011), we examined how retention is predicted by these individual differences in performance on the two types of referent selection trials. Critically, if novelty seeking during novel RS trials is crucial for learning, performance on novel RS trials should predict retention. If on the other hand, the relevant factor is the general level of organization in the lexical system we may predict that

familiar trial performance would be a better predictor, either alone, or along with novel RS performance.

To examine these issues, each child's accuracy on known RS and novel RS trials was computed. This was entered into a logistic mixed model predicting retention from both factors. Data from both experiments was pooled, resulting in 94 data points. The fixed factors were novel and known RS performance, experiment (1 and 2), and the interactions. Vocabulary (z-scored) was a covariate. The dependent variable was trial-by-trial retention accuracy. A random intercept of participant was used; adding a random intercept for word did not improve model fit, and there were no within-participant variables to use as random slopes.

The analysis<sup>5</sup> showed no relationship between accuracy on novel RS trials and retention ( $\beta = -.26$ ,  $SE = .28$ ,  $Z = .93$ ,  $p = .35$ ). However, there was a highly significant relationship between known RS performance and retention ( $\beta = .51$ ,  $SE = .20$ ,  $Z = 2.55$ ,  $p = .01$ ): children who performed better on the Known RS trials (overcoming the novelty bias) tended to retain better (see Table 1). The nine children who showed retention on both trials averaged 63.9% correct in known RS; the 26 who retained one of two were 40.3% correct on known RS; and the 24 children who were incorrect on both retention trials averaged 23% correct on known RS. However, there was no effect of overall vocabulary size ( $\beta = .11$ ,  $SE = .25$ ,  $Z = .45$ ,  $p = .65$ )<sup>6</sup> or experiment ( $\beta = .43$ ,  $SE = .89$ ,  $Z = .48$ ,  $p = .63$ ), and experiment did not interact with known RS ( $\beta = -.14$ ,  $SE = .38$ ,  $Z = .37$ ,  $p = .71$ ) or novel RS ( $\beta = -.36$ ,  $SE = .57$ ,  $Z = .63$ ,  $p = .53$ ). The lack of a correlation with vocabulary may be because known RS and vocabulary were correlated ( $r = .28$ ,  $p = .03$ ), making it more difficult to

---

<sup>5</sup> This analysis was repeated using all retention trials (not just those corresponding to items for which accurate responding was shown on novel RS trials). There was a continued significant effect of known RS ( $\beta = .467$ ,  $p = .007$ ), and no other significant effects.

<sup>6</sup> Vocabulary remained non-significant even if the number of object names (sections 2-10 of the MCDI) in a child's vocabulary (rather than their total vocabulary size) were used.



see a unique effect of vocabulary. Alternatively, it may be that overall vocabulary size (number of words produced) is not equivalent to the strength of individual word knowledge. That is, two children who produce the word “cup” may still respond to cups in the world differently. Thus, this leads to a null result for overall vocabulary, but a significant correlation between trials tapping specific word knowledge (known RS). Similarly, all children were reported to comprehend (and typically produce) labels for all know items used, thus parent-report of individual word knowledge did not vary across children. Moreover, though it is reliable and valid as a parent-report measure of vocabulary (Fenson et al., 1994), the MCDI is an imperfect measure of individual word knowledge. A parent indicating that a child produces a specific word does not guarantee the child would recognize that word in a new context nor that they would have robust lexical entries for that word. Likewise, the fact that two parents indicate that their children have the same size vocabulary does not mean those children will necessarily process words in the same way, especially in ambiguous tasks like that used here.

-----Insert Table 1 about here -----

Overall, this pooled analysis suggests that a novelty bias does not facilitate known object selection, and may impair retention of novel word-referent links. Only the children who overcame this bias on known RS trials demonstrated retention. Importantly, no relationship was found between *novel* RS and retention (of the same items). Rather, *known* RS predicts retention. Though nearly all children performed near ceiling on novel RS, only those who selected correctly when prompted with a *known* word (known RS) showed evidence of retaining the *novel* word-object links. This suggests that at 18-months-old, it is the balance between novelty biases and known word knowledge that predicts long-term retention rather than performance in-the-

moment of exposure to a new word. However, it is difficult with the discrete selection responses made in this task to differentiate whether word learning is driven by an object-specific novelty preference or a more generally immature system that, because of weak lexical connections, is more likely to attend to novelty. We explore these possible mechanisms further in our computational simulations.

### **Computational investigations**

The striking finding from Experiments 1 and 2 is that 18-month-olds have a pervasive novelty bias which can override their ability to select a known referent. This could stem from the objects—their relative novelty attracts attention. Alternatively, it could stem from overall weaker word-referent links within the lexical system such that knowledge cannot drive behavior in-the-moment. This could more broadly affect performance on both familiar and novel trials. While Experiment 2 suggests that it is not the latter, the former is difficult to directly capture empirically. Thus, we examined both possibilities in a computational model to understand where the novelty bias might come from and how it might play out over the real-time dynamics of referent selection and the long-term dynamics of learning.

This required a model that could simulate referent selection and retention with both familiar and novel words. We adopted the dynamic associative architecture of McMurray, Samuelson, and colleagues (McMurray, Zhao, Kucker, & Samuelson, 2013; McMurray et al., 2012). This model captures initial in-the-moment referent selection as well as slowly acquired word-referent links. The model begins by assuming any word could map to any object. It then eliminates incorrect links and strengthens correct ones using simple associative learning sensitive to co-occurrence. The model captures real-time referent selection (known and novel) as a process of dynamic competition among available referents. Competition plays out over the associative

network; thus, learning shapes in-the-moment referent selection.

McMurray et al.'s (2012) model did not originally capture visual novelty. Thus, we integrated novelty into it in two ways. First, we treated novelty as external to the lexical system arising from highly salient objects by boosting activation for novel objects. Second, we treated novelty as an endogenous part of the lexical system, deriving from immature word-referent links.

### **Architecture**

A complete computational description of this model is provided in the appendix. Here we give a conceptual overview. This model consists of banks of localist auditory and visual inputs, indicating which word was heard and which object(s) were present. These are connected via an intermediate bank of lexical nodes (Fig. 3). During processing, a small amount of lateral inhibition is applied to the input layers and activity within each layer is then normalized to sum to 1.0, simulating competition between items. Activation of the input layers then spreads to the lexical layer. A larger amount of lateral inhibition is applied to this layer (simulating competition among potential referents, forcing the model to choose one word), and again, activation is normalized and sent back down to the input layers. Cycling continues until activation stops changing, simulating competition among candidate interpretations of the event. After competition and settling, the visual unit with the highest activation is the model's choice.

-----Insert Figure 3 about here -----

To implement learning, association weights start with small random values. The model is then exposed to a series of words with a random number of visual competitors. Learning is applied at every cycle: weights are adjusted with a Hebbian learning rule strengthening the

connections between co-active units, and weakening connections with only one node active. Learning occurs throughout processing. Each input layer had 40 localist nodes, each representing a word or object. This number was selected to make computation time reasonable while still allowing for an adequate representation of known and novel words in the task (see McMurray, Horst, & Samuelson, 2012 for discussion). Thirty nodes corresponded to objects and words in the model's vocabulary, 10 were held out as novel objects. The lexical layer had 500 units (McMurray et al., 2012).

A simulation began by teaching the model a small number of words, akin to the vocabulary of an 18-month-old. On each learning trial, we activated a single "known" auditory unit and a random number of visual units ( $M=15$ ). Visual units could be both known and novel. The model settled and associations were updated. A single epoch consisted of learning/training trials for each of the 30 words and then testing for each word. Testing was done with a 30AFC production task (one name-known visual unit with all 30 known auditory units) and a 3AFC comprehension test (one known auditory unit with three name-known visual units). This assessed the model's known word vocabulary prior to testing referent selection and retention for novel words.

The model was periodically tested on referent selection (RS) and retention using a procedure similar to Experiment 1. For RS trials, a subset of three "known" items the model had correctly chosen during the 3AFC comprehension test was selected. On each of the 10 RS trials, two known items and one of the 10 untrained objects were active. The active auditory unit corresponded to either a known or novel word. Retention trials tested five of the novel words from RS, along with their corresponding visual unit. Foil items on retention trials included one additional item from the Novel RS trials, and one completely novel item. This differed slightly

from the children for whom the third item was not completely novel (it had been seen but not named on the known RS trials). RS and retention were tested every 500 epochs. This assessed the effect of novelty on referent selection early in lexical development (e.g. an 18-month-old), and also checked that the novelty manipulation did not disrupt performance later in development (e.g. 24-months).

### **Simulation 1: Object novelty**

The first simulations implemented novelty as an item-specific attentional bias. On each trial, objects were given a temporary boost of activation that was a function of the cumulative amount of prior exposure to that object.

$$n_x^{(t+1)} = n_x^{(t)} - \delta \cdot v_x^{(t)} \cdot (n_x^{(t)} - 1) \quad (1)$$

Consequently, for the trained words this boost decreased over the course of exposure (since their objects had been seen many times). Novelty boost ( $n_x^{(t)}$ ) to visual unit  $x$ , at time  $t$  is given by Equation 1. The novelty boost ( $n_x$ ) at the next time step is equal to its current novelty boost minus some proportion ( $\delta$ ) of that unit's current activation ( $v_x$ ). This value is multiplied by ( $n_x - 1$ ) so the ultimate value of  $n_x$  asymptotes at 0. This novelty boost was added to the initial activation of visual units, prior to competition. Novelty boost was adjusted at the beginning of the trial. The boost was applied and updated continually across all simulation phases, including within testing phases to simulate real-world situations in which testing is itself a learning phase. However, after a testing phase was complete, weights were reset to their pre-testing levels to allow for continual unbiased testing at later points.

The change in novelty boost over the simulation is a function of the units' activation, not frequency. Since activation within an input layer is normalized to 1.0, the initial activations in the input (prior to the novelty bias) depend on the number of competitors – if only two objects

are active, they will each be initially active at 0.5; if five are active they are each active at 0.2 (see Fig. 4). Further, this boost does not violate the overall spirit of the model in which lexical behavior operates as a product of simple, domain-general mechanisms, as this object boost can be seen as something akin to a Hebbian bias weight.

-----Insert Figure 4 about here -----

We compared simulations with and without this novelty boost (see Appendix for parameter values). For the novelty and baseline models, 10 simulations were trained on 30 word-referent links. With no clear benchmarks to determine when the models performed like an 18-month-old, we tested the models throughout acquisition and made inferences about early and late development on the basis of performance across tasks.

Results are depicted in Figure 5. Baseline simulations (top panel) capture the original findings of Horst and Samuelson (2008; replicating McMurray et al., 2012). Novel RS starts and remains very high throughout development, well above chance at 100% for all epochs. This is consistent with excellent novel RS performance of 18-month-old children. On known RS selection, baseline simulations show an early period of reduced performance on known RS followed by near ceiling performance later. This is *unlike* 18-month-old children in Experiments 1 and 2, who show chance-level performance.

However, the addition of the novelty bias yields data more like children (Fig. 5, bottom panel). Performance on known RS, early in development is significantly lower than baseline performance. As the model develops, the novelty bias systematically diminishes (due to continual exposure to novel items) and known RS increases, mimicking prior work with 24- and

30-month-old infants (Bion et al., 2013; Horst & Samuelson, 2008). Also like children, novel RS performance starts and remains high throughout development.

Finally, at retention, the baseline simulations showed near chance performance across development. When a novelty boost is added to the model, retention remained near chance, nearly identical to Experiments 1 and 2. The lack of eventual retention could be seen as a possible limitation. However, prior instantiations of this model that have included more epochs do demonstrate an increase in retention as vocabulary size/age increases (McMurray, Horst, & Samuelson, 2012, Figure 14C), but with a slightly different parameter space (e.g. different levels of exposure to items during training).

-----Insert Figure 5 about here -----

These simulations validate the notion that an object-based novelty bias can give rise to the pattern of familiar word recognition observed behaviorally: strong performance when only known-name objects are present, but weak performance with a single novel foil (Fig. 6). Secondly, the fact that good novel RS was observed in both the baseline and the novelty boost models suggests that disambiguation does not entirely derive from object-based novelty biases.

-----Insert Figure 6 about here -----

Intriguingly, even the baseline model shows reduced performance on known RS trials (though not at chance like 18-month-olds). This suggests there may be information *within* the lexical system that tracks novelty and could play a similar role. Simulation 2 used this insight to explore a potential endogenous, lexically-based, novelty bias.

**Simulation 2: Endogenous novelty**

Simulation 1 suggests the possibility that the novelty bias might be a marker for deeper developmental differences in the lexical system. The depressed performance on known RS early in training is reminiscent of behavioral performance on those trials. One potential cause is that activation for visual units is based in part on feedback from the lexical units. Early in training, associations are random and high – most words are connected to most objects. By the end of training, associations only link known words and their referents, so feedback reflects “lexical knowledge”. However, as McMurray et al. (2012, Figure 12) show, this is not the case for novel objects whose associations remain relatively unpruned throughout training. This is because weights are only pruned when a visual unit is active *and* the corresponding lexical unit is inactive. Thus, early in learning, novel objects are likely to have several high random connections to familiar words, which may boost their activation when familiar words are presented. Critically, this mechanism does not rely on bottom-up activation for novel objects; rather the pull to novelty comes from the fact that the relatively immature lexical system has not sufficiently pruned these connections.

This could explain the small reduction in familiar word performance observed early in the baseline simulation, but why was it small? One possibility is that the weights were initialized to rather small values (between 0 and 0.25); consequently, feedback across unpruned connections does not offer much boost to visual units. To check this, we increased the range of the random weights to 0 – 0.75, and eliminated the external novelty boost. Preliminary simulations suggested that this change affected the speed and quality of learning (i.e. vocabulary levels remained low longer) so other parameters (such as the level of ambiguity during learning) were also adjusted to ensure optimal model performance (see Appendix).



Results are shown in Figure 7. These simulations, though not identical, show strong novel RS performance (100%) and weaker retention, but poor early performance on known RS.

However, these simulations do have some limitations, notably they also show slower vocabulary growth (the model “produced” only 5/30 words). There are several reasons for slower vocabulary growth—more in-the-moment ambiguity makes it harder to learn, and more randomness in initial weights means the model needs more time to acquire a robust associative network. This means Simulation 2 does not perfectly map onto children whose RS and retention performance tends to increase as overall vocabulary size increases. Overall, however, this simulation suggests that novelty effects on familiar referent selection can also emerge endogenously from within the lexical system. In fact, an exploration of the model space suggested a range of other ways to achieve a similar profile: for instance, increasing the number of visual competitors during training (the degree of referential ambiguity) can depress familiar word performance without degrading novel RS performance (McMurray et al., 2012).

Most importantly, Simulation 2 demonstrates how novelty can arise *within* the lexical system, deriving from unpruned associations left over from initialization. As novel objects are never seen, their initial weights are not pruned, and these can give rise to a strong bias to attend to them even as the referent for a known word (see also Kidd et al., 2012 for a similar example with stimuli complexity). Computationally, this is a way to enable new word-object mappings in a model that is not growing new connections. However, it may also have some neurophysiological basis as it is well known that there is substantial neural pruning through early childhood (e.g., Huttenlocher & Dabholkar, 1997). In addition, these unpruned connections could serve to functionally bias the early lexical network towards novel objects—often the right response when the child knows very few words (Mervis & Bertrand, 1994, and see McMurray et

al., 2012 for further discussion of this issue).

-----Insert Figure 7 about here -----

### **Discussion of simulations**

The empirical results from Experiments 1 and 2 suggest that 18-month-olds' behavior in referent selection tasks with is driven by object novelty. Simulation 1 confirms that this could derive from attraction to novel objects. However, Simulation 2 also shows the viability of endogenous mechanisms—an immature lexical system in which unpruned connections allow activation from familiar words to spread to novel ones. While the resulting behavior – selection of the most novel object – is certainly consistent with extant ideas of endogenous novelty biases, the simulations provide insight into a new possible basis for this—an immature lexical system with high associations between novel objects and *familiar* words. Indeed, this possibility parallels the results of the combined analysis; children who demonstrated stronger lexical knowledge (overcoming the novelty bias on known-name trials) were the ones to demonstrate retention of novel word-referent links.

### **General Discussion**

Across experiments and simulations the same story emerged: 18-month-old children showed a strong bias towards novel objects during referent selection. They even chose novel foils over familiar targets when requested with a known name. Furthermore, Experiment 2 shows the pull to novelty remains even when we diminished the novelty of the foils via familiarization. Together these experiments suggest that whatever lexical knowledge 18-month-olds bring to this task is overridden by novelty. Attention to novelty in these children created seemingly good

performance when the novel object was the target (novel RS), but poor performance when the target was known (known RS). These young children also demonstrated chance performance on retention of these novel words. Additional analysis counterintuitively revealed no relationship between novel RS performance and retention, but a strong link between known RS and retention (for those same novel words). This suggests the critical predictor of learning is not the accuracy of initial referent selection, but rather the ability of knowledge for contrasting familiar words to outweigh novelty in these challenging naming situations (see also Smith & Yu, 2013). That is, accurate initial novel selection is not necessarily indicative of long-term learning (see also Fitneva & Christiansen, 2011). In addition, the results suggest that behavior of young children during word learning situations may be driven by simple, low-level perceptual biases and not high-level intelligent learning systems.

The computational models show the novelty bias could derive from two possible mechanisms: 1) an exogenous attentional boost to novel objects, or, 2) an internally-generated push towards a novel object caused by an unrefined lexicon (e.g. poorly organized associative network that includes weak knowledge (correct word/referent associations) and many unpruned spurious connections). Both models yielded patterns of referent selection similar to 18-month-olds: novel RS near ceiling early in development, but known RS near chance. And both implementations are reasonable: novel objects are likely more salient, but the idea that novelty bias may simply be a remnant of unpruned initial connections is parsimonious. Critically, however, both forms of the bias were implemented with very simple model of associative learning and competitive decision making. Thus, again, the results suggest that the source of the novelty effect may be low level, and not necessarily geared to the task of learning words.

The simulations also point towards a possible explanation for the difference between our

finding with 18-month-old children and the prior finding that 24-month-old children correctly select known objects. The model assumes that all items could be associated with labels, but critically, they differ in the strength of their connections in part as a function of pruning. Thus, the critical difference between 18- and 24-month olds could be in the strength of [unnecessary] connections between known words and novel objects. This fits with Grassmann and colleagues (2015) who found that children's ability to exclude a known object as a referent for a novel word was best predicted by a child's production not comprehension. Together, this work demonstrates that referent selection and retention could be based on a dynamic balance between the pull of novelty and the coherence of the mappings between words and referents in the developing vocabulary (Samuelson et al., 2017).

Overall, the current study augments prior investigations into the processing underlying language acquisition in young word learners and suggests a more nuanced view of mechanisms that might support word learning biases. Some prior work implicates high-level reasoning and specific word learning constraints (e.g. mutual exclusivity, N3C) to explain children's rapid word learning abilities. We cannot rule this out for older children who overcome their novelty biases (e.g. Horst & Samuelson, 2008), or for some children in this study who may have operated based on name knowledge. That said, our results are in line with other studies that show the pervasive influence of novelty in word learning (Mather, 2013b; Mervis & Bertrand, 1994) in that the young vocabulary learners here chose the novel objects even when the task was to ignore them (see also control trials in Frank, Sugarman, Horowitz, Lewis, & Yurovsky, 2016). This challenges the idea that children are using novelty in a goal-directed way to advance their linguistic knowledge (Horst et al., 2011). The same bias that appears intelligent when used in response to a novel word looks quite the opposite in response to a familiar word. Importantly,

however, the mechanisms of this bias need not be mutually exclusive. The differences between children—some biased towards novelty and others able to correctly select name-known targets and retain new novel word links – suggest individual children’s behavior may have been driven by multiple factors. Likewise, because the strength of individual word-referent links in an individual child’s vocabulary are likely to differ, we should also expect variability across trials within children.

To understand word learning, we must account for how novelty drives behavior in-the-moment of mapping but also how it, in conjunction with lexical knowledge, may, or may not, influence learning. Our results support arguments that word learning operates over both the shorter timescale of in-the-moment referent selection processes and longer learning timescales of retention and vocabulary building (Bion et al., 2013; McMurray et al., 2012; Kucker et al., 2016). These studies demonstrate that behavior during referent selection may be a product of novelty biases in the-moment, but novelty may not cascade to affect long-term learning. Importantly, our work also complements prior arguments unifying the processes that support novel referent selection and known word comprehension (c.f., McMurray et al., 2012) by suggesting that one unintelligent mechanism driving behavior during referent selection with known words is not likely different from the mechanism operating when selecting the referent of novel words.

### References

- Akhtar, N., Carpenter, M., & Tomasello, M. (1996). The role of discourse novelty in early word learning. *Child Development, 67*(2), 635. <https://doi.org/10.2307/1131837>
- Axelsson, E. L., & Horst, J. S. (2014). Contextual repetition facilitates word learning via fast mapping. *Acta Psychologica, 152*, 95–99. <https://doi.org/10.1016/j.actpsy.2014.08.002>
- Bahrnick, L. E., Lickliter, R., & Flom, R. (2004). Intersensory redundancy guides the development of selective attention, perception, and cognition in infancy. *Current Directions in Psychological Science, 13*(3), 99–102.
- Baldwin, D. A., & Markman, E. M. (1989). Establishing word object relations: A first step. *Child Development, 60*, 381–398.
- Bion, R. A. H., Borovsky, A., & Fernald, A. (2013). Fast mapping, slow learning: Disambiguation of novel word–object mappings in relation to vocabulary learning at 18, 24, and 30months. *Cognition, 126*(1), 39–53.  
<https://doi.org/10.1016/j.cognition.2012.08.008>
- Borovsky, A., Ellis, E. M., Evans, J. L., & Elman, J. L. (2015). Lexical leverage: category knowledge boosts real-time novel word recognition in 2-year-olds. *Developmental Science, 19*(6), 918–932. <https://doi.org/10.1111/desc.12343>
- Carey, S., & Bartlett, E. (1978). Acquiring a single new word. *Papers and Reports on Child Language Development, 15*, 17–29.
- Cohen, L. B., DeLoache, J. S., & Rissman, M. W. (1975). The effect of stimulus complexity on infant visual attention and habituation. *Child Development, 46*(3), 611–617.  
<https://doi.org/10.2307/1128557>

- Diesendruck, G., & Markson, L. (2001). Children's avoidance of lexical overlap: A pragmatic account. *Developmental Psychology, 37*(5), 630–641. <https://doi.org/10.1037//0012-1649.37.5.630>
- Fantz, R. L. (1964). Visual experience in infants. decreased attention to familiar patterns relative to novel ones. *Science, 146*(3644), 668–670.
- Fenson, L., Philip, P. S., Reznick, J. ., Bates, E., Thal, D. J., & Pethick, S. J. (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development, 59*(5), i–185.
- Fernald, A., & Marchman, V. A. (2012). Individual differences in lexical processing at 18 months predict vocabulary growth in typically developing and late-talking toddlers. *Child Development, 83*(1), 203–222. <https://doi.org/10.1111/j.1467-8624.2011.01692.x>
- Fernald, A., Perfors, A., & Marchman, V. A. (2006). Picking up speed in understanding: Speech processing efficiency and vocabulary growth across the 2nd year. *Developmental Psychology, 42*(1), 98–116. <https://doi.org/10.1037/0012-1649.42.1.98>
- Fernald, A., Pinto, J. P., Swingle, D., Weinberg, A., & McRoberts, G. W. (1998). Rapid gains in speed of verbal processing by infants in the 2nd year. *Psychological Science, 9*(3), 228–231.
- Fitneva, S. A., & Christiansen, M. H. (2011). Looking in the wrong direction correlates with more accurate word learning. *Cognitive Science, 35*(2), 367–380. <https://doi.org/10.1111/j.1551-6709.2010.01156.x>
- Frank, M. C., Sugarman, E., Horowitz, A. C., Lewis, M. L., & Yurovsky, D. (2016). Using tablets to collect data from young children. *Journal of Cognition and Development, 17*(1), 1–17. <https://doi.org/10.1080/15248372.2015.1061528>

- Golinkoff, R., Mervis, C. B., & Hirsh-Pasek, K. (1994). Early object labels: The case for a developmental lexical principles framework. *Journal of Child Language*, *21*(1), 125–155.
- Goodman, J. C., McDonough, L., & Brown, N. B. (1998). The role of semantic context and memory in the acquisition of novel nouns. *Child Development*, *69*(5), 1330–1344.
- Grassmann, S., Schulze, C., & Tomasello, M. (2015). Children’s level of word knowledge predicts their exclusion of familiar objects as referents of novel words. *Frontiers in Psychology*, *6*, 1–8. <https://doi.org/10.3389/fpsyg.2015.01200>
- Halberda, J. (2003). The development of a word-learning strategy. *Cognition*, *87*(1), B23–B34. [https://doi.org/10.1016/S0010-0277\(02\)00186-5](https://doi.org/10.1016/S0010-0277(02)00186-5)
- Halberda, J. (2006). Is this a dax which I see before me? Use of the logical argument disjunctive syllogism supports word-learning in children and adults. *Cognitive Psychology*, *53*(4), 310–344. <https://doi.org/10.1016/j.cogpsych.2006.04.003>
- Hollich, G. J., Hirsh-Pasek, K., Golinkoff, R. M., Brand, R. J., Brown, E., Chung, H. L., ... Bloom, L. (2000). Breaking the language barrier: An emergentist coalition model for the origins of word learning. *Monographs of the Society for Research in Child Development*, *65*(3), i–135. <http://www.jstor.org/stable/3181533>
- Horst, J.S., & Hout, M.C. (2012). The Novel Object and Unusual Name (NOUN) Database: A collection of novel images for use in experimental research. *Behavioral Research Methods*, *48*(4), 1393-1409. doi: 10.3758/s13428-015-0647-3
- Horst, J. S., & Samuelson, L. K. (2008). Fast mapping but poor retention by 24-month-old infants. *Infancy*, *13*(2), 128–157. <https://doi.org/10.1080/15250000701795598>



- Horst, J. S., Samuelson, L. K., Kucker, S. C., & McMurray, B. (2011). What's new? Children prefer novelty in referent selection. *Cognition, 118*(2), 234–244.  
<https://doi.org/10.1016/j.cognition.2010.10.015>
- Horst, J.S., Scott, E.J., & Pollard, J.A. (2010). The role of competition in word learning via referent selection. *Developmental Science, 13*(5), 706-713. doi: 10.1111/j.1467-7687.2009.00926.x.
- Hunter, M. A., Ross, H., & Ames, E. (1982). Preferences for familiar or novel toys: Effects of familiarization time in 1-year-olds. *Developmental Psychology, 18*(4), 519–529.
- Huttenlocher, P.R., & Dabholkar, A.S. (1997). Regional differences in synaptogenesis in human cerebral cortex. *Journal of Comparative Neurology, 387*(2), 167-178. doi: 10.1002/(SICI)1096-9861(19971020)387:2<167::AID-CNE1>3.0.CO;2-Z
- Kidd, C., Piantadosi, S., & Aslin, R. N. (2012). The Goldilocks effect: human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS ONE, 7*(1), e36399.
- Kucker, S. C., McMurray, B., & Samuelson, L. K. (2016). Slowing down fast mapping: Redefining the dynamics of word learning. *Child Development Perspectives, 9*(2), 74–78.  
<https://doi.org/10.1111/cdep.12110>
- Kucker, S. C., & Samuelson, L. K. (2012). The first slow step: Differential effects of object and word-form familiarization on retention of fast-mapped words. *Infancy, 17*(3), 295–323.  
<https://doi.org/10.1111/j.1532-7078.2011.00081.x>
- Lewkowicz, D. J. (1996). Perception of auditory–visual temporal synchrony in human infants. *Journal of Experimental Psychology: Human Perception and Performance, 22*(5), 1094.

- Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology*, *20*(2), 121–157.
- Mather, E. (2013a). Bootstrapping the early lexicon: How do children use old knowledge to create new meanings? *Frontiers in Psychology*, *4*.  
<https://doi.org/10.3389/fpsyg.2013.00096>
- Mather, E. (2013b). Novelty, attention, and challenges for developmental psychology. *Frontiers in Psychology*, *4*, 1–4. <https://doi.org/10.3389/fpsyg.2013.00491>
- Mather, E., & Plunkett, K. (2010). Novel labels support 10-month-olds' attention to novel objects. *Journal of Experimental Child Psychology*, *105*(3), 232–242.  
<https://doi.org/10.1016/j.jecp.2009.11.004>
- Mather, E., & Plunkett, K. (2012). The role of novelty in early word learning. *Cognitive Science*, *36*(7), 1157–1177. <https://doi.org/10.1111/j.1551-6709.2012.01239.x>
- Mather, E., Schafer, G., & Houston-Price, C. (2011). The impact of novel labels on visual processing during infancy: Novel labels and visual processing. *British Journal of Developmental Psychology*, *29*(4), 783–805. <https://doi.org/10.1348/2044-835X.002008>
- McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review*, *119*(4), 831–877. <https://doi.org/10.1037/a0029872>
- McMurray, B., Zhao, L., Kucker, S. C., & Samuelson, L. K. (2013). Pushing the envelope of associative learning: Internal representations and dynamic competition transform association into development. In L. Gogate & G. Hollich (Eds.). In *Theoretical and computational models of word learning: Trends in psychology and artificial intelligence* (pp. 49–80). Hershey, PA: IGI Global.

- Merriman, W. E., & Schuster, J. M. (1991). Young children's disambiguation of object name reference. *Child Development, 62*(6), 1288–1301.
- Mervis, C. B., & Bertrand, J. (1994). Acquisition of the novel name–nameless category (N3C) principle. *Child Development, 65*(6), 1646–1662.
- Nazzi, T., & Bertoncini, J. (2003). Before and after the vocabulary spurt: two modes of word acquisition? *Developmental Science, 6*(2), 136–142.
- Rigler, H., Farris-Trimble, A., Greiner, L., Walker, J., Tomblin, J. B., & McMurray, B. (2015). The slow developmental time course of real-time spoken word recognition. *Developmental Psychology, 51*(12), 1690–1703. <https://doi.org/10.1037/dev0000044>
- Roder, B. J., Bushnell, E. W., & Sasseville, A. M. (2000). Infants' preferences for familiarity and novelty during the course of visual processing. *Infancy, 1*(4), 491–507.
- Rose, S., Gottfried, A., Melloy-Carminar, P., & Bridger, W. H. (1982). Familiarity and novelty preferences in infant recognition memory: Implications for information processing. *Developmental Psychology, 18*(5), 704–713.
- Samuelson, L. K., Kucker, S. C., & Spencer, J. P. (2017). Moving word learning to a novel space: A dynamic systems view of referent selection and retention. *Cognitive Science, 41*(S1), 52–72.
- Samuelson, L. K., & Smith, L. B. (1998). Memory and attention make smart word learning: An alternative account of Akhtar, Carpenter, and Tomasello. *Child Development, 69*(1), 94–104.
- Samuelson, L. K., & Smith, L. B. (2000). Children's attention to rigid and deformable shape in naming and nonnaming tasks. *Child Development, 71*(6), 1555–1570.

- Sandman, C., Wadhwa, P., Hetrick, W., Porto, M., & Peeke, H. V. (1997). Human fetal heart rate dishabituation between 30 and 32 weeks gestation. *Child Development, 68*(6), 1031–1040.
- Sekerina, I. A., & Brooks, P. J. (2007). Eye movements during spoken word recognition in Russian children. *Journal of Experimental Child Psychology, 98*(1), 20–45.  
<https://doi.org/10.1016/j.jecp.2007.04.005>
- Slater, A., Morison, V., & Rose, D. (1984). New-born infants' perception of similarities and differences between two-and three-dimensional stimuli. *British Journal of Developmental Psychology, 2*(4), 287–294.
- Smith, L. B., & Yu, C. (2013). Visual attention is not enough: Individual differences in statistical word-referent learning in infants. *Language Learning and Development, 9*(1), 25–49.  
<https://doi.org/10.1080/15475441.2012.707104>
- Spiegel, C., & Halberda, J. (2011). Rapid fast-mapping abilities in 2-year-olds. *Journal of Experimental Child Psychology, 109*(1), 132–140.  
<https://doi.org/10.1016/j.jecp.2010.10.013>
- Swingley, D. (2010). Fast mapping and slow mapping in children's word learning. *Language Learning and Development, 6*(3), 179–183.  
<https://doi.org/10.1080/15475441.2010.484412>
- Wetherford, M.J., & Cohen, L.B. (1973). Developmental changes in infant visual preferences for novelty and familiarity. *Child Development, 44*(3), 416-424.  
<http://www.jstor.org/stable/1127994>

White, K. S., & Morgan, J. L. (2008). Sub-segmental detail in early lexical representations.

*Journal of Memory and Language*, 59(1), 114–132.

<https://doi.org/10.1016/j.jml.2008.03.001>

Wiebe, S. A., Cheatham, C. L., Lukowski, A. F., Haight, J. C., Muehleck, A. J., & Bauer, P. J.

(2006). Infants' ERP responses to novel and familiar stimuli change over time:

Implications for novelty detection and memory. *Infancy*, 9(1), 21–44.

Yu, C., & Smith, L.B. (2011). What you learn is what you see: using eye movements to study

infant cross-situational word learning. *Developmental Science*, 14(2), 165-180.

[10.1111/j.1467-7687.2010.00958.x](https://doi.org/10.1111/j.1467-7687.2010.00958.x)

**Appendix.** Details and parameters of the model

In this section we describe the computational details of the model. Parameter settings used for all of the free parameters in the below equations are provided in Table A1.

On any given trial, the activation for the single auditory unit (corresponding to the word that was heard) is set to 1.0. Similarly activation of each of the present visual objects is set to 1.0. Auditory and visual activations then undergo a small amount of inhibition, and normalized so that activity of all of the units sums to 1.0 using Equation A1.

$$v_x^{(t+1)} = \frac{(v_x^{(t+1)})^\tau}{\sum_{z \in V} (v_z^{(t+1)})^\tau} \quad (\text{A1})$$

Here,  $v_x^{(t)}$  refers to activation of visual unit  $x$  at time  $t$  (the same formula is also used for the auditory units); and  $\tau$  is a temperature parameter controlling the amount of inhibition (usually set to a small value like 1.05).

Next activation spreads from both input layers to the lexical layer (Equations A2 & A3).

$$\Delta l_x = (\sum_{z \in A} w_{xz} a_z + \sum_{z \in V} u_{xz} v_z) \quad (\text{A2})$$

$$l_x^{(t+1)} = l_x^{(t)} + \tau_{ff} \cdot \Delta l_x^{(t)} \quad (\text{A3})$$

Here,  $l_x$  refers to lexical unit  $x$ ,  $a_z$  is the auditory unit  $z$ , and  $w_{xz}$ , is the associative weight between  $a_z$  and  $l_x$ . Similarly,  $v_z$  refers to visual unit  $z$ , and  $u_{xz}$  is its connection to  $l_x$ . In Equation S3, the current lexical units' activation is a function of its prior activation, and  $\Delta l_x$  multiplied by a temperature parameter,  $\tau$ . Next activation in the lexical layer is squared and normalized as in Equation A4.

$$l_x^{(t+1)} = \frac{(l_x^{(t+1)})^2}{\sum_{z \in L} (l_z^{(t+1)})^2} \quad (\text{A4})$$

This quadratic normalization scheme causes the most active unit to become more active, and units with less activation to become less so. If run over several iterations, it gradually leads to a winner take all pattern of activation. Next, activation then feeds back to the input layers (Equations A5 and A6)

$$\Delta a_y = a_y \cdot \sum_{z \in L} l_z w_{yz} \quad (\text{A5})$$

$$a_y^{(t+1)} = a_y^{(t)} + \tau_{fb} \cdot \Delta a_y^{(t)} \quad (\text{A6})$$

Note that because the sum of the current lexical to auditory weights is multiplied by the current activation of the auditory weights, the resulting changes in activation only account for currently active units. Thus, non-active auditory units do not get activated by top-down processes. These layers are then normalized with Equation S1, and the updated activation of the input layers is fed back up to the lexical layer. This cycle continues until the RMS change in the lexical layer between cycles is below a very small threshold. The resulting visual node with the highest activation is then taken to be the model's choice in response to the original auditory prompt.

Weights are altered during each cycle of processing using Equations S7 and S8.

$$\begin{aligned} \Delta w_{xy} = & a_x l_y (1 - w_{xy}) \\ & - .5 \cdot (1 - a_x) \cdot l_y w_{xy} \\ & - .5 \cdot a_x (1 - l_y) \cdot w_{xy} \end{aligned} \quad (\text{A7})$$

$$w_{xy}^{(t+1)} = w_{xy}^{(t)} + \eta \cdot \Delta w_{xy}^{(t)} \quad (\text{A8})$$

The learning rate,  $\eta$ , controls how much change there is on each cycle, and is usually very small. Under this learning scheme, a connection is updated if both of its units are active (A7, 1<sup>st</sup> line), and it decays if the input is active but the output is not (2<sup>nd</sup> line) or vice versa (3<sup>rd</sup> line). If neither unit is active, no change occurs. Positive change in weights is multiplied by  $(1-w_{xy})$  so that as the weight approaches 1.0, change gets progressively smaller. Negative changes in

weights are multiplied by the current weight so that as the weight approaches 0, change similarly gets progressively smaller.

Table A1. Model Parameters

Parameter	symbol	computational role	Baseline Rep.	Sim1:Exon.NB	Sim2:Endog.NB
Input units		number of nodes in each input layer	40	40	40
Known words/objects		total number of words to acquire during vocabulary acquisition	30	25	30
Novel words/objects		unnamed objects present a small amount of time during vocab acquisition	10	10	10
Supernovel words/objects		novel objects held out until referent selection	0	5	5
Lexical units		number of nodes in lexical/decision layer	500	500	500
Weight size		initial random weights of associations	.25	.25	0 – 0.7
Learning rate $\eta$		rate at which weights increase on each cycle	.0005	.0005	.0005
Decay		rate at which spurious connections die away relative to the learning rate	.5	.5	.5
Referential ambiguity		average number of competitors during vocabulary acquisition	.5	.5	.70
Novelseen		amount of time novel items are present as foils during vocabulary acquisition	.01	.01	.01
Feed-forward temperature	$\tau_{ff}$	input to lexical layer	.01	.01	.01
Feedback temperature	$\tau_{fb}$	lexical to input layer	2.0	2.0	2.0
Novelty bias starting value	$n_{xy}$	boost to un-seen items	n/a	2.5	n/a
Novelty bias decay rate	$\delta$	rate at which novelty boost decays at each exposure	n/a	.01	n/a