Accounting for the build-up of proactive interference across lists in a list length paradigm
reveals a dominance of item-noise in recognition memory

Julian Fox

University of Melbourne

Simon Dennis

University of Melbourne

Adam F. Osth

University of Melbourne

Address correspondence to:

Julian Fox (E-mail: jfox1@student.unimelb.edu.au)

Author Note

Data, model code, and supplementary materials from this article can be found on our Open Science Framework (OSF) page (<https://osf.io/7qxdm/>). This work was supported by a grant from the Australian Research Council, ARC DE170100106, awarded to Adam Osth.

Abstract

There has been a longstanding debate concerning whether interference in recognition memory is attributable to other items on the study list (i.e., item-noise) or to prior memories (i.e., context-noise and background-noise). Recently, Osth and Dennis (2015) devised a global matching model that could estimate the magnitude of each interference contribution and they found that context-noise and background-noise were dominant in recognition. In the present investigation, data from a list length experiment were analysed using variants of the Osth, Jansson, Dennis and Heathcote (2018) model, that integrates the memory retrieval components of the Osth and Dennis (2015) model with the diffusion decision model (Ratcliff, 1978) to jointly account for choice probabilities and RT distributions. The standard version of the model, like existing recognition models, treated each condition as if no proactive interference had accumulated over the session. A more comprehensive version of the model allowed both study and test items from prior conditions to contribute proactive interference (PI) to future conditions. While the standard model estimated a dominance of background-noise, the PI model estimated a dominance of item-noise, reversing the conclusions made by Osth and Dennis (2015). Along with list length, the experimental design provided a measure of the test position effect (TPE). While the standard model attributed the TPE to context drift, the PI model attributed the TPE to both context drift and increases in item-noise.

Keywords: list length effect; recognition memory; proactive interference; global matching model; diffusion decision model; test position effect

Accounting for the build-up of proactive interference across lists in a list length paradigm reveals a dominance of item-noise in recognition memory

For many decades, global matching models have been used to account for the existence of various phenomena in recognition memory. Global matching models of memory specify that retrieval involves a global match between retrieval cues and the contents of memory (Clark & Gronlund, 1996). In recognition memory, retrieval cues consist of item cues (i.e., the test items themselves) and context cues (i.e., the experimental context that corresponds to the study list in question). These cues are combined into a single memory probe and globally matched with all items in memory, generating a degree of memory strength that forms the basis of recognition decisions. In traditional approaches, memory strength is compared to a criterion to make a decision, but in some more recent approaches (e.g., Nosofsky, Little, Donkin & Fific, 2011; Osth, Jansson, Dennis & Heathcote, 2018) memory strength is used to drive an evidence accumulation process to produce choice and response latency.

The vast majority of global matching models make the assumption that only items from the current study list contribute to the global match, and as such, they do not account for the influence of memories from prior lists. However, in memory research there have been some notable cases where considering the entire set of items learned in an experiment and the order in which they were presented led to a shift in theoretical conclusions. A classic example stems from two studies that explored the Brown-Peterson paradigm (Brown, 1958; Peterson & Peterson, 1959). Brown (1958) initially found that the rate of forgetting stimuli could be well predicted by the length of the decay interval between study and test, leading to the conclusion that forgetting was caused by the passage of time. However, when Keppel and Underwood (1962) evaluated forgetting curves for each individual item over six trials, they found that the rate of forgetting was not constant, but was steeper for each successive trial. In fact, performance for the first study item was almost completely unaffected by the decay

interval. They concluded that the reduction in performance for the study items was caused by proactive interference from the study items themselves and not by the passage of time. This re-analysis was notable in that it reversed the conclusions of Brown (1958) – from the notion of time causing forgetting to interference causing forgetting – simply by considering the order of trials, and not just the set of trials as a whole. Recently, a similar shift in conclusions has occurred in the recognition literature, as common findings in the list length paradigm have been challenged by restricting analyses to particular conditions in an experiment (Brandt, Zaiser & Schnuerch, 2019).

The List Length Effect and Global Matching Models

The list length effect (LLE) refers to a phenomenon whereby recognition performance is poorer for items from long lists than items from short lists (Strong, 1912). The LLE was one of the earliest findings in the memory literature (Strong, 1912) and has been replicated several times throughout a century of research (Gillund & Shiffrin, 1984; Gronlund & Elam, 1994; Ohrt & Gronlund, 1999; Ratcliff, Clark & Shiffrin, 1990).

Global matching models, such as the search of associative memory model (SAM; Gillund & Shiffrin, 1984), and the MINERVA 2 model (Hintzman, 1988) have been able to account for the LLE in recognition by assuming that each item in memory contributes additional variance to the strength of the global match. This assumption leads to the prediction of an LLE as more study items (i.e., a longer list) will lead to greater variance in the global match. This additional variance increases the overlap between the signal and noise distributions and thus leads to less accurate recognition judgments. Since models like SAM and MINERVA 2 predict that interference in the global match stems from the items on the current study list other than the target, they are usually referred to as *item-noise* models.

In recent times, however, global matching models like the bind-cue-decide model of episodic memory (BCDMEM; Dennis & Humphreys, 2001) have been developed that predict

a null LLE in recognition. These models distinguish themselves from the aforementioned item-noise models in two key ways. Firstly, they assume that the representations of different items in memory are sufficiently non-overlapping, such that the match of the retrieval cues to a study item (other than the target) is minimal or non-existent and produces no interference. This assumption leads to the prediction of a null LLE as the additional study items from longer lists will not contribute variance to the global match, and as such, will not decrease the signal-to-noise ratio or recognition accuracy. Secondly, these models assume that the set of items involved in a global match is not restricted to items from the study list in question, but also includes pre-experimental memories. As a consequence, the strongest matches will be matches between the test item and memories corresponding to prior instances in which the same item has been witnessed. Naturally, test items that are associated with many prior contexts, such as high frequency (HF) words, will lead to greater interference than items associated with fewer prior contexts, such as low frequency (LF) words. Since these models predict that the majority of interference in a global match stems from prior contexts in which the target item was witnessed, and not the other list items themselves, they are usually referred to as *context-noise* models.

In addition to item-noise and context-noise, some models predict that there is an influence of *background-noise* in recognition decisions (e.g., Murdock & Kahana, 1993a, 1993b; Osth & Dennis, 2015). Background-noise is interference stemming from matches to items other than the target in pre-experimental contexts. Like context-noise models, models that assume a predominance of background-noise will predict that additional study items contribute minimal variance to the global match, and therefore that there is a minimal LLE in recognition.

It is important to emphasize that both context-noise and background-noise are classifiable as *proactive interference*, given that they refer to interference generated by items prior to the

current study list. While the terms context-noise and background-noise conventionally refer to interference from pre-experimental sources, they can also be used to encompass interference that is generated by items from prior study lists within the experiment. Likewise, while the term proactive interference traditionally refers to interference generated by prior list items within an experiment, it can also be used to encompass pre-experimental sources of interference. A summary of these noise contributions and other key terms is provided in Table 1 below.

Table 1.

Definitions of key terms used throughout this article

Term	Definition
List length effect (LLE)	Finding that recognition performance is worse for long study lists than short study lists.
Test position effect (TPE)	Finding that recognition performance decreases over test trials.
Global matching model	Mechanistic model framework for recognition memory. Retrieval cues are matched towards all memories simultaneously; the sum of all the similarities produces a degree of memory strength that is used to make a decision.
Diffusion decision model (DDM)	Evidence accumulation model of decision-making. Evidence accumulates noisily until either a top boundary ('YES' response) or bottom boundary ('NO' response) are reached and the respective response is made.
Item-noise	Interference generated by items from the current study list other than the probe item.
Proactive interference (PI)	<p>Interference generated by items prior to the current study list. Can be divided into context-noise and background-noise.</p> <p>Context-Noise. Interference generated by occurrences of the probe item prior to the current study list (either from within the experiment, or prior to the experiment taking place).</p> <p>Background-Noise. Interference generated by occurrences of items not from the current study list (either from within the experiment, or prior to the experiment taking place).</p>
Context drift	Change in the current context representation in response to events, such as successive study trials or test trials.

Evidence Against the Item-Noise Account of Memory

Given the robustness of the LLE throughout decades of research, models of memory have tended to favor the item-noise account of recognition. However, Dennis, Lee and Kinnell (2008) suggested that previously observed LLEs may have been due to confounding factors in the standard list length paradigm (see Figure 1). Three of the potential confounds they highlighted were: retention interval, attention, and contextual reinstatement.

The first of these alludes to the fact that retention intervals following short and long study lists are often not equated. Since recognition testing sometimes begins immediately after the presentation of a study list, there is often a greater retention interval for items from a long list than for items from a short list. This is a potential confound since the longer retention interval for the long list items may result in weaker memory strength. This confound can be controlled by having a filler task precede the study list (i.e., a proactive design) or follow the study list (i.e., a retroactive design), thus ensuring that test begins at the same time for both short and long lists.

The second confound highlighted by Dennis et al. (2008) was decreases in attention throughout the study list. They argued that attention is particularly likely to decrease throughout the presentation of a long study list and that this has a detrimental effect on items in the later positions of the list. This is a potential confound since recognition performance may appear lower for long lists simply due to a subset of those items being learned under poor conditions, as opposed to there being more interference among the items. This confound can be controlled by testing participants with words from the same positions across both lists. Since all items from a short list effectively appear at the start of their list, the tested long list items must also appear at the start of their list. While maintaining equal retention intervals, this method can only be achieved with a retroactive design. Put simply, a proactive design can allow for the control of either retention interval or attention, but not both, while a

retroactive design allows for both to be controlled simultaneously. This distinction has been supported empirically as proactive designs have been shown to pronounce the LLE (Bowles & Glanzer, 1983; Ohrt & Gronlund, 1999; Underwood, 1978), while retroactive designs have been shown to nullify the effect (Dennis et al., 2008; Kinnell & Dennis, 2011; Murnane & Shiffrin, 1991; Schulman, 1974).

A third confound outlined by Dennis et al. (2008) is ‘contextual reinstatement’. Throughout the completion of a list length experiment, it can be assumed that a participant’s contextual state will fluctuate. In a retroactive design, the current context following a long study list will be associated with the items near the end of the list, while the current context following a short list will be associated with the filler task. The end-of-list context following a long list is likely to be considered a sufficient context cue for memory retrieval during testing. However, the ‘filler task’ context following a short list will not be, and participants will instead have to reinstate the context associated with the beginning of the study list. The lack of contextual reinstatement following a long list is a disadvantage since the end-of-list context may not be sufficient to retrieve the critical items from the beginning of the study list. This confound can be controlled by enforcing a delay period between study and test in the long list condition, and between the filler task and test in the short list condition, thus ensuring that both lists lead to the reinstatement of the study list context.

Once Dennis et al. (2008) controlled for all these confounding factors (see Figure 1), they found a null LLE. They used this finding to argue against the item-noise account of recognition, and to instead support the context-noise account of recognition, as context-noise models do not predict the effect. Note that while null LLEs have been reported in other studies using similar controls (Dennis & Humphreys, 2001; Kinnell & Dennis, 2011, 2012; Schulman, 1974), recent research has suggested that the presence of the LLE is also affected by stimulus class (Kinnell & Dennis, 2012). For instance, while word stimuli may lead to a

null LLE, small LLEs have been observed for face stimuli.

Figure 1. Schematics of a standard list length design (top) and a list length design that controls for the confounds highlighted in Dennis et al. (2008) (bottom)

Further evidence against the item-noise account of memory, however, was established by Osth and Dennis (2015). In their work, they proposed a variant of the matrix model (Humphreys, Bain & Pike, 1989) that involves bindings between item and context vectors. The parameters of the model are the similarities between the cued item/context vectors and the ones stored in memory, which allows for the quantification of each interference contribution (i.e., item-noise, context-noise and background-noise). When fit to ten recognition memory datasets with diverse stimuli classes and experimental manipulations (e.g., list-length, list-strength and word frequency), the model's parameter estimates revealed that most of the interference stemmed from pre-experimental sources (i.e., pre-experimental context-noise and background-noise), and not from the study list items themselves (i.e., item-noise), particularly for datasets using word and scene stimuli. Background-noise, in particular, was estimated to be the dominant source of interference, with the exception of HF words that led to roughly equal estimates for context-noise and background-noise. It is important to note, however, that both context- and background-noise relate to memories formed prior to the experiment. In its standard form, the Osth and Dennis (2015) model does

not account for prior lists that are part of the same experiment. Regardless, the conclusions of Osth and Dennis (2015) were reinforced when the model was expanded to account for source memory data (Osth, Fox, McKague, Heathcote & Dennis, 2018), and when the model was integrated with the diffusion decision model (DDM; Ratcliff, 1978) to account for response times (RTs) and various effects (e.g., LLE, test position effect) simultaneously (Osth, Jansson et al., 2018).

Evidence For the Item-Noise Account of Memory

The conclusions of Osth and Dennis (2015), however, have been challenged by recent work from Brandt et al. (2019) that demonstrates how the LLE is dependent on condition order. Throughout three experiments and with various classes of photos used as stimuli (e.g., landscapes, flowers and faces), participants would perform recognition on one short and one long list in a random order. When performance was averaged across multiple cycles, a null LLE was observed. However, when analysis was restricted to the first study-test cycle in the experiments, an LLE was observed. They argued list length *does* impact recognition performance, but that after an initial list, the learning of subsequent lists produces an increase in proactive interference that equates the total interference in memory between groups (i.e., those currently studying a short list or a long list), resulting in a null LLE.

To clarify this point, consider that in the first study-test cycle of a list length experiment, participants will commit either a short or long list of items to memory (e.g., 20 or 80 items), but that in the second study-test cycle, participants will often study the opposite list type. The consequence of this is that by the test phase of the second cycle, each participant would have committed an equal number of items to memory (i.e., 100 items). If it is assumed that participants create an experiment-wide context cue at retrieval that does not fully differentiate between the current list and prior lists, then there will be roughly equivalent amounts of interference between groups (i.e., those currently studying a short list or a long list) and a null

LLE will be observed. In other words, the accumulation of proactive interference from prior lists may be masking an underlying list length effect throughout the experiment.

Brandt et al. (2019) were able to support these assumptions by fitting a modified version of the MINERVA 2 model to their data. While the basic MINERVA 2 model assumes that only items from the current study list are represented in memory, Brandt et al. highlighted that this relies on the assumption that the episodic context of each list is distinguishable from the contexts of other lists (Hintzman, 1988). They argued that the contexts of separate lists may be quite similar, due to them all involving item presentations on a computer screen, and consequently that it is plausible for episodes prior to the study list to be represented in memory. In other words, they argued that the context cue used at retrieval is unable to fully isolate the current study list. These assumptions were incorporated into MINERVA 2 by allowing additional traces in memory for each extra-list item, but also by associating the extra-list items with a lower learning parameter than the current list items to reflect the fact that some of their features may have been forgotten over time. Note that the incorporation of extra-list items into MINERVA 2 is comparable to the assumption of background-noise in other memory models (e.g., Murdock & Kahana, 1993a, 1993b; Osth & Dennis, 2015). The modified version of the MINERVA 2 model was able to capture the diminished LLE found in the 2nd study-test cycle of each experiment due to the fact that each extra-list item contributes additional variance to the target and lure distributions. Since MINERVA 2 explains the LLE as resulting from a greater increase in distributional variance for long lists compared to short lists, increases in variance from the extra-list items serve to mask the relative differences in variance between the short and long lists, and consequently reduce the LLE.

Overall, by simply considering condition order effects, Brandt et al. were able to establish evidence for the LLE in recognition, while suggesting that the item-noise account of recognition may be more tenable than previously believed.

The Present Study

The first aim of the present investigation will be to replicate the results of Brandt et al. (2019) using word stimuli and a greater number of lists. This first aim is important as Brandt et al.'s design only involved picture stimuli and two study-test cycles; it is unclear whether their results will be replicated using different stimuli and across a larger number of lists. Recall that Kinnell and Dennis (2012) observed a null LLE for word stimuli but a small LLE for face stimuli; it could be the case that Brandt et al. observed large LLEs due to their usage of picture stimuli. Our second aim will be to investigate the results with a computational model that will be able to measure the sources of interference (i.e., the contributions of item-noise, context-noise and background-noise). This second aim is of particular importance since the conclusions of Osth and Dennis (2015) – that pre-experimental context-noise and background-noise are the primary sources of interference in recognition – may not hold when a computational model is made to account for each individual study list, and not simply all conditions as a whole (as it most commonly done). Specifically, it is necessary to determine whether estimates of item-noise will increase once the model accounts for the influence of proactive interference across conditions. While it is rare for inter-list differences to be investigated using memory models, especially with regards to recognition memory, notable exceptions in other memory tasks are Lohnas, Polyn and Kahana (2015) who modelled build-up of proactive interference across lists in free recall tasks, and Mewhort, Shabang and Franklin (2018) who did the same for both free recall and cued recall.

A limitation of Brandt et al. (2019) was that they did not consider how RTs or decision-related factors may have influenced their findings. As such, the results of the present study will be analysed using the Osth, Jansson et al. (2018) global matching model, that combines the memory retrieval components of the Osth and Dennis (2015) model with the DDM (Ratcliff, 1978), an evidence accumulation model of decision making. The advantage of

using the Osth, Jansson et al. model is that it is able to account for both choice probabilities and RT distributions. By allowing the model to be constrained by RTs, it can be revealed whether changes in performance across cycles are the result of memory strength (e.g., interference and/or learning) or decision-based factors (i.e., changes in response bias and/or the speed-accuracy trade-off). In contrast, any memory model lacking a decision-making component – such as the model of Brandt et al. (2019) – will exclusively attribute changes in performance to memory-related factors. The distinction between memory- and decision-based factors is non-trivial; for instance, Donkin and Nosofsky (2012) found in a short-term memory task that increases in RT with increasing list length were due entirely to higher speed-accuracy thresholds. In the context of the present investigation, it could be the case that the Brandt et al. (2019) findings are not attributable to a build-up of proactive interference (as they claim), but rather to changes in attention or strictness in responding. Support for the latter can be found in Osth, Jansson et al. (2018) who showed that the speed-accuracy threshold changes across trials in a list (n.b., we define trial as the presentation of a single stimulus within a list, not the list as a whole). It is not unreasonable to consider that similar changes are occurring *between* lists (not just within them) and that this could help explain the changes in the LLE across lists that Brandt et al. reported.

In addition to accounting for changes in decision-based factors across lists, our model will also account for changes in learning rate across lists. Justification for this comes from the model of Mewhort et al. (2018) that accounts for changes across trials and lists by assuming a build-up of proactive interference. In Mewhort et al.'s model, the learning rate is inversely proportional to the similarity between the retrieval cues and the contents of memory (prediction-error learning; Farrell & Lewandowsky, 2002). Since a build-up of observed items over lists increases the chance of a particular item being more similar to the contents of memory, the model predicts that the learning rate should decrease over lists. It could

therefore be the case that the Brandt et al. (2019) findings are partly attributable to decreases in learnings across lists, and not the aforementioned changes in memory- and decision-based factors.

Although the present experiment largely mirrors that of Brandt et al. (2019), a number of amendments have been made to the design in addition to the use of word stimuli and a greater number of lists. Firstly, participants will not be tested with targets and lures exclusively, but will also be tested with an additional item type which we will refer to as 'studied lures'. These items will originate from the study list prior to the current one, and will not have been used as targets in any test list. Participants will be instructed to treat these items as lures and to respond to them as 'NEW'. An example of a studied lure would be a word (e.g., 'cat') that is presented once during the study phase of the 8th study-test cycle, not presented at test in the 8th cycle, but is later presented at test in the 9th cycle. The inclusion of these studied lures will be crucial throughout the modelling exercise as they will provide considerable constraint in determining the extent to which items from prior lists match with the current context cues. Despite evidence in the list discrimination literature suggesting that participants can distinguish between list contexts (Criss & Shiffrin, 2005; Hintzman, Caulton & Levitin, 1998), the MINERVA 2 model of Brandt et al. (2019) specifies that the context cues at retrieval do not isolate the current study list. The model would therefore predict that there is a similar match to studied lures as there is to targets, and consequently that the false alarm rate (FAR) for studied lures should be comparable to the hit rate (HR).

The experimental design will also involve a lengthy test list that participants complete at the end of each session, thus providing a measure for the test position effect (TPE). The TPE refers to the phenomenon in which memory performance decreases during the course of testing (Criss, Malmberg & Shiffrin, 2011). In addition to providing additional constraint in the modelling exercise, the inclusion of the long test lists will allow participant estimates of

the TPE to be calculated and subsequently correlated with other measures, such as item-noise, changes in context throughout the course of testing (i.e., context drift), and changes in the speed-accuracy trade-off over trials. Using a comprehensive global matching model, Osth, Jansson et al. (2018) found a high correlation between the TPE and context drift, suggesting that context drift was the primary cause of the TPE. Such comparisons are theoretically pertinent as item-noise models specify that the TPE is caused by a build-up of interference from additional item presentations (Criss et al., 2011; Gillund & Shiffrin, 1984; Shiffrin & Steyvers, 1997) and therefore predict that the TPE should be correlated with measures of item-noise.

The data from this experiment can be found on our Open Science Framework (OSF) page (<https://osf.io/7qxgm/>).

Method

Participants

Prior to data collection, we determined that 50 participants would be sufficient for the purposes of the present study. This figure was chosen as it maintains consistency with key prior studies (e.g., Dennis et al., 2008; Kinnell & Dennis, 2012). We also allowed for the testing of additional participants in any instances where a prior participant could not complete both test sessions (e.g., they displayed near-chance accuracy in the first session and were not invited back). Note that our target of 50 participants was not determined through a formal power analysis. Large degrees of statistical power were not required as our main inferences of interest were derived from computational modelling, and not data analysis.

The participants were $N=51$ students aged 18 or older at the University of Melbourne. Age and gender details were not collected for each participant as they were not relevant to our study. Participants were recruited through the Melbourne School of Psychological

Sciences Reimbursed Participant System (RPS). The participants were reimbursed \$15 per session by participating in this research. Participants were only recruited for the experiment if they were native English speakers (either monolingual or bilingual). One participant displayed overall discrimination performance in the first session that was very close to chance ($d' = 0.204$). Consequently, they were not invited to participate in the second session and their first-session data was removed from the analysis. The data of another participant was removed from the analysis as they expressed that they had misinterpreted the experimental instructions after testing.

Design

The experiment initially took place over 4 sessions, the purpose of which was to provide a strong within-subjects analysis of the data. It was expected that performance would be comparable across sessions since it is unlikely that proactive interference from a prior session would impact future sessions. However, it was found that performance decreased drastically in the 3rd and 4th sessions (a formal analysis of this is provided in the Results section). Consequently, the length of the experiment was reduced to 2 sessions.

In each of the sessions, there were 13 cycles of a recognition task, each involving a study phase and test phase. The first 12 cycles of each session were comprised of 6 pairs of cycles, each involving 1 short study list of 20 words (10 HF, 10 LF) and 1 long study list of 80 words (40 HF, 40 LF). The first session would begin with either the short or long study list and the second session would begin with the opposite list length. The ordering of the short and long study lists in the remaining pairs of cycles in each session was randomized (see Figure 2, row A). The purpose of having a short and long list in each pair of cycles was to ensure that the build-up of within-experiment interference (i.e., proactive interference and item-noise) would be unequal between participants in all the odd-numbered cycles, but equal in all the even-numbered cycles, thus preserving the integrity of Brandt et al.'s (2019) 2-cycle design over

the lengthy course of the experiment. The logic behind this was that if Brandt et al.'s explanation for the appearance and disappearance of the LLE from cycle 1 to cycle 2 was accurate, then there would be an LLE across all odd-numbered cycles and a null LLE across all even-numbered cycles.

The first cycle of each session, regardless of the length of the study list, involved a test phase of 16 targets (8 HF, 8 LF) and 16 lures (8 HF, 8 LF). For cycles 2-12 in each session, the test phases were comprised of 16 targets (8 HF, 8 LF), 12 lures (6 HF, 6 LF) and 4 studied-lures (2 HF, 2 LF). The studied lures consisted of words from the study list immediately prior to the current list that had not already been used as a target. Both targets and studied-lures were always selected from the first block of 20 words in each study list so to ensure that attention was controlled for between short and long lists.

The 13th and final cycle of each session involved a long study list of 80 words (40 HF, 40 LF) and a long test of 80 targets (40 HF, 40 LF) and 80 lures (40 HF, 40 LF). The retention interval between study and test was manipulated so that the length between study and test was equivalent for both short and long lists (see Figure 2, row B).

Figure 2. Schematic of the 13 study lists that comprised each testing session (row A) and descriptions of the structure and stimuli used for the study and test phases of the short lists, long lists, and final list (row B). Note that in row A, half the participants would have the labels for each session reversed, such that the first session begins with a long list and the second session begins with a short list

Materials

The experimental data was collected through a computer program using the Python programming language with the Pygame package. The word list of the experiment was initially derived from the MRC Psycholinguistic Database. It was specified that the length of the words be limited to 3-10 letters and that all plurals or derivational variants of words be excluded. This preliminary word list was then analysed through N-Watch (Davis, 2005) where the CELEX frequencies for each word were generated. From N-Watch, two new

wordlists were created: a list of HF words ($N=1946$) with CELEX frequencies ranging between 25-2500 occurrences per million, and a list of LF words ($N=2326$) with CELEX frequencies ranging between .5-3 occurrences per million.

Procedure

The overall length of each session was roughly 45 minutes. Throughout the study phase of each of the experimental cycles, participants were presented with a series of individual words for 1000 milliseconds each, separated by an interstimulus interval (ISI) of 250 milliseconds, with each word appearing in white font in the centre of the screen. Upon seeing each word, participants were required to make a “pleasantness” rating, in that they pressed the “i” key if they thought the word was pleasant, or the “k” key if they thought it was unpleasant. If the study phase involved a short list, participants would be presented with a card game which served as a distractor task. The card game required participants to respond to sequences of cards with rules such as, “Hit the SPACEBAR if you see two cards with the same SUIT in a row”. The card game lasted 75 seconds. This was to ensure that the retention interval was equated between short and long lists.

In the test phase of each cycle, participants were presented with a series of words in white font in the centre of the computer screen. For each word, participants were instructed to press the “1” key if they thought the word was OLD, or the “0” key if they thought it was NEW. Participants were also instructed to respond with NEW if they believed the word had been seen on a study list other than the current one (i.e., a studied-lure). Responses to each test item that were faster than 0.2 seconds resulted in a ‘TOO FAST’ warning. Responses to items that were slower than 7s resulted in a ‘TOO SLOW’ warning, upon which the current trial would terminate and the next word would appear. At the end of the first session, participants scheduled their next session at their own convenience. The second session could take place from 1-10 days after the first session.

Analysis was performed on the data via a series of Bayesian ANOVAs and t-tests. These analyses were performed using the statistics program JASP (Version 0.7.5.6).

Results

For all of the analyses, the first test trial of each list was excluded. This was due to the fact that the initial test trials were frequently associated with abnormally long RTs. Additionally, trials with RTs below 200 ms or above 2500 ms were excluded. Overall, 3.40% of trials were excluded. All analyses were conducted using Bayesian ANOVAs and t-tests due to their ability to provide evidence for null effects (Wagenmakers, 2007). Bayes factors were calculated using the JZS method (Rouder, Morey, Speckman & Province, 2012) with default prior scales. Each of the calculated Bayes factors (BF_{10}) represent the ratio of the null model likelihood to the alternative model likelihood, such that $BF_{10} > 1.000$ constitutes evidence in favor of the alternative model, and $BF_{10} < 1.000$ constitutes evidence in favor of the null model. It is assumed that Bayes factors larger than 3.000 or lower than 0.333 provide significant evidence in favor of the alternative or null models respectively, whereas Bayes factors within the range of 0.333 to 3.000 provide marginal evidence. Furthermore, it is assumed that Bayes factors above 10.000 or below 0.100 constitute strong evidence in favor of a model.

From the initial 10 participants who completed four sessions of testing, we found marginal evidence that discriminability (d') in the first session ($M = 1.988$, $SD = 0.759$) was the same as the second session ($M = 1.860$, $SD = 0.774$, $BF_{10} = 0.520$), indeterminate evidence that discriminability in the first session was the same or different in the third session ($M = 1.724$, $SD = 0.910$, $BF_{10} = 1.048$), and significant evidence that discriminability in the first session was higher than the fourth session ($M = 1.508$, $SD = 0.967$, $BF_{10} = 3.310$). Since the potential/observed declines in performance in the third/fourth sessions would confound the ability to perform a within-subjects analysis, all subsequent participants completed only

two sessions, and the third/fourth session data from the initial 10 participants were removed from the final dataset.

In the final dataset, it was found that overall discriminability was slightly higher in the first session ($M = 1.955$, $SD = 0.672$) than the second session ($M = 1.880$, $SD = 0.734$). However, there was marginal evidence in favor of the null effect ($BF_{10} = 0.354$). Consequently, the remainder of the analyses in this section will be conducted on data from both sessions simultaneously. Throughout the entirety of both sessions, it was also revealed that discriminability was significantly higher for LF words ($M = 2.199$, $SD = 0.812$) than HF words ($M = 1.651$, $SD = 0.611$, $BF_{10} > 1,000.000$).

Across the first 12 cycles of the experiment, it was found that the mean discriminability for short lists ($M = 2.059$, $SD = 0.769$) was significantly higher than that for long lists ($M = 1.767$, $SD = 0.605$, $BF_{10} > 1,000.000$). These results are indicative of a small, yet highly significant, LLE throughout the data. This finding is consistent with previous research that did not control for contextual reinstatement (Dennis et al., 2008).

When analysing the first cycle of the experiment in isolation, it was also found that mean discriminability for short lists ($M = 2.632$, $SD = 0.792$) was significantly higher than long lists ($M = 2.221$, $SD = 0.738$, $BF_{10} = 16.979$). However, when the second study-test cycle was analysed, it was found that discriminability was lower in short lists ($M = 1.872$, $SD = 0.832$) than long lists ($M = 2.103$, $SD = 0.681$). A one-tailed Bayesian t-test revealed strong evidence in favor of the null effect ($BF_{10} = 0.092$). The appearance of a positive LLE in the first cycle and its disappearance in the subsequent cycle is consistent with the results of Brandt et al. (2019). As such, we demonstrate that the results of Brandt et al. can be extended to word stimuli.

Extending beyond the two-cycle scope of Brandt et al. (2019), we analysed the LLE across all odd- and even-numbered cycle of the experiment. As described earlier, the amount

of within-experiment interference (i.e., proactive interference and item-noise) is equated between groups (those currently studying a short list or a long list) on the even-numbered cycles but not in the odd-numbered cycles. If Brandt et al.'s explanation for the null LLE is accurate (i.e., it is caused by an equation of within-experiment interference between groups), then there should be an LLE in the odd-numbered cycles, but a null LLE in the even-numbered cycles. It was found that short list discriminability ($M = 2.085$, $SD = 0.866$) was significantly higher than long list discriminability ($M = 1.686$, $SD = 0.795$) across all odd-numbered cycles of the experiment; not just the first cycle ($BF_{10} > 1,000.000$). In contrast, performance for short lists ($M = 1.812$, $SD = 0.762$) was not significantly different than that for long lists ($M = 1.805$, $SD = 0.777$) across the even-numbered cycles ($BF_{10} = 0.120$). These results provide further support for the findings of Brandt et al. (2019) since LLEs were observed in all the cycles where participants had been exposed to an unequal total of items throughout the experiment (i.e., within-experiment interference was unequal), and null LLEs were observed in all the cycles where participants had been exposed to the same total of items (i.e., within-experiment interference was equal). This data pattern can be observed in Figure 7 (presented in the Modelling section).

Regarding the choice proportions for each item type, it was observed that HR ($M = 0.694$, $SD = 0.153$) was significantly higher than the FAR for studied lures (FAR_{SL} ; $M = 0.437$, $SD = 0.157$) across all cycles ($BF_{10} > 1,000.000$). Furthermore, it was found that FAR_{SL} was significantly higher than the FAR for lures ($M = 0.111$, $SD = 0.085$) across all cycles ($BF_{10} > 1,000.000$). These results indicate that there is a degree of context match to items from previous study lists during retrieval, but that participants are reasonably capable of distinguishing the current list context from the previous list context.

Regarding the HR and FAR for items in the final test list, it was found that the HR was significantly higher in the first block of 40 trials ($M = 0.642$, $SD = 0.174$) than in the last

block of 40 trials ($M = 0.453$, $SD = 0.220$, $BF_{10} > 1,000.000$). Conversely the difference in FAR between the first block ($M = 0.134$, $SD = 0.112$) and last block ($M = 0.109$, $SD = 0.131$) was minimal ($BF_{10} = 0.876$). Naturally, it was found that overall discriminability was significantly higher in the first block ($M = 1.597$, $SD = 0.732$) than in the last block ($M = 1.266$, $SD = 0.667$, $BF_{10} = 182.937$). These findings are consistent with the TPE typically observed in prior research (Criss et al., 2011).

Modelling

While one aim of the present investigation was to replicate the results of Brandt et al. (2019), the primary aim was to determine how accounting for their findings would impact the conclusions from a computational model. The data from the present experiment was investigated using several different variants of the Osth, Jansson, Dennis and Heathcote (2018) global matching model. As mentioned earlier, the Osth, Jansson et al. (2018) model is comprised of a front-end memory model (Osth & Dennis, 2015) and a back-end DDM (Ratcliff, 1978), thus allowing for the prediction of both choice probabilities and RT distributions. Formal descriptions of both the Osth & Dennis (2015) model and DDM are provided in the following sections. Following this, descriptions of each model variant will be provided. The model code we used throughout this study can be found on our OSF page (<https://osf.io/7qxgm/>).

The Osth and Dennis (2015) Model

The primary constituents of the Osth and Dennis (2015) model are items (I) and contexts (C). In the model, items and contexts are represented as vectors, and bindings between them are represented as outer products. Each outer product can be seen to signify a particular learned experience; for example, a word (item) that was seen on a study list (context). For a given set of items on a study list L , the stored contents of memory are represented by a matrix

M that corresponds to the sum of each outer product:

$$M = \sum_{t \in L} r(I_t \otimes C_s) \quad (1)$$

where r is a scalar learning rate parameter, the subscript t corresponds to each item on the study list, and the subscript s corresponds to the study list context. Note that the r scalar can reflect phenomena such as the number of item repetitions or individual ability.

When an item is presented at test, the memory strength s is derived by matching the retrieval cues against the memory matrix M :

$$s = (I'_t \otimes C'_s) \cdot M \quad (2)$$

where the dashes indicate that the retrieval cues may not be identical to the vectors stored in memory. While many models of recognition define the contents of the stored vectors and retrieval cues, the Osth and Dennis (2015) model employs an approximate analytic solution that allows the similarities between the stored vectors and retrieval cues to be parameterized directly (as normal distributions). The model's analytic derivations and the equations for the means and variances of the target and lure distributions are outlined in Appendix A.

Each item in the memory matrix M can be classified into four categories depending on whether they match or mismatch to the item and context cues at retrieval. The four categories are: self-match, item-noise, context-noise, and background-noise (refer to Figure 3 for a visual demonstration of each category). Note that in the model, only pre-experimental contributions of context-noise and background-noise are accounted for.

Figure 3. Schematic of the Osth and Dennis (2015) model. A self-match occurs when a probe – consisting of an item cue (e.g., orange) and context cue (i.e., study list) – is globally matched to memory and there is a match on both the item and context dimensions. Item-noise occurs when there is a match on the context dimension but a mismatch on the item dimension. Context-noise occurs when there is a match on the item dimension but a mismatch on the context dimension. Background-noise occurs when there is a mismatch on both dimensions

The overall influence of each noise type can be expressed as the product of the item match/mismatch strength and context match/mismatch strength. The means and variances for the target and lure distributions are as follows (refer to Appendix A for the full derivations):

$$\begin{aligned}\mu_{old} &= r\mu_{tt}\omega \\ \mu_{new} &= 0\end{aligned}\tag{3}$$

$$\begin{aligned}\sigma_{old}^2 &= r^2(\omega^2\sigma_{tt}^2) + && \text{Self-Match} \\ &+ r^2(l-1)(\omega^2\sigma_{ti}^2) + && \text{Item-Noise} \\ &+ (\mu_{tt}^2\rho + \sigma_{tt}^2\rho) + && \text{Context-Noise} \\ &\beta && \text{Background-Noise}\end{aligned}\tag{4}$$

$$\begin{aligned}\sigma_{new}^2 &= r^2l(\omega^2\sigma_{ti}^2) + && \text{Item-Noise} \\ &+ (\mu_{tt}^2\rho + \sigma_{tt}^2\rho) + && \text{Context-Noise} \\ &\beta && \text{Background-Noise}\end{aligned}\tag{5}$$

In the model, a self-match occurs when a cued target item in conjunction with the cued study list context are matched to a representation of the same item and context pairing in memory. In such cases, there is a match between the cued item vector and stored item vector (i.e., an item match). The strength of the item match is drawn from a normal distribution with mean μ_{tt} (conventionally fixed at 1) and variance σ_{tt}^2 , both model parameters. The item match variability factor σ_{tt}^2 can be seen to reflect encoding variability of the cued item features. In addition to the item match, there is also a match between the cued context vector and stored context vector (i.e., a context match), the strength of which is represented by the model's ω parameter. The value of ω decreases with test position due to the influence of context drift (described below). The overall strength of the self-match is derived by multiplying the item and context matches together, and further multiplying by the learning rate parameter r (first row of Equation 3 and Equation 4).

Item-noise occurs when the same retrieval cues are matched to a study list item (other than the cued item) that is bound to a representation of the study list context in memory. While there is a match between the context vectors, represented by ω , there is a mismatch between the cued item and stored item (i.e., an item mismatch). The strength of the item mismatch is drawn from a normal distribution with mean 0 and variance σ_{ti}^2 , the latter of which is a model parameter. For a single item, the similarity to the retrieval cues is calculated by multiplying the context match and item mismatch together. The overall contribution of item-noise is calculated by summing the similarities of each study list item (other than the target) to the retrieval cues (context match * item mismatch), and multiplying the sum by the learning rate parameter r (second row of Equation 4, first row of Equation 5). As such, a greater number of items in the study list will lead to a greater contribution of item-noise. Furthermore, if an item elicits a lower context match (e.g., the item is from a prior study list), then the interference contribution from that item will also be lower. Additionally, it should be

noted that increases in the contribution of item-noise (i.e., higher values of item mismatch variability σ_{ti}^2) directly increase the influence of proactive interference, and increase the magnitudes of the LLE and TPE.

Context-noise occurs when the retrieval cues are matched to a representation of the same item in memory that is bound to any context prior to the study list context. Note that ‘prior contexts’ here refer to pre-experimental contexts, and not prior study lists within the experiment. Similar to the self-match, context-noise involves a match between the item vectors, the strength of which is drawn from a normal distribution with mean μ_{tt} (fixed to 1) and variance σ_{tt}^2 . However, there is a mismatch between the cued context and stored context (i.e., a context mismatch). The context mismatch is drawn from a normal distribution with mean 0 and variance ρ , another model parameter. As HF words are associated with a greater number of prior contexts than LF words, the ρ parameter is allowed to vary between word frequency classes. The overall contribution of context-noise is equal to the product of the item match and context mismatch (third row of Equation 4, second row of Equation 5).

Finally, background-noise occurs when the retrieval cues are matched with an item other than the target that is bound to a pre-experimental context in memory. In such cases, there is an item mismatch and context mismatch between the cued and stored vectors. The overall contribution of background-noise, or the product of the item mismatch and context mismatch, is represented by the β parameter (fourth row of Equation 4, third row of Equation 5). Following convention, each model employed in this work had the β parameter fixed at .05 (Osth, Jansson et al., 2018).

The resulting memory strengths from the global match are transformed into likelihood ratio distributions (expressions for the means and variances can be found in Appendix A). The primary purpose of this transformation is to capture the mirror effect. The mirror effect is a finding where manipulations that increase performance have opposite effects on the HR and

FAR (Glanzer & Adams, 1985). After the likelihood ratio transformation, a drift criterion c is applied to the log likelihood ratio distributions, reflecting variance in bias across participants. Each model in this work allowed c to vary between short and long lists.

Context Drift

To implement changes in performance through the course of testing, the model includes context drift. The term ‘context drift’ can be understood to reflect the gradual changes in one’s temporal-spatial and social-emotional environment (i.e., contextual state) as their daily experiences unfold (Sahakyan & Smith, 2014). Even though context drift occurs throughout the passage of time, the two are not synonymous. For instance, context will change rapidly if a participant switches from one task to another, but will drift relatively slowly if a participant performs the same task over an extended period. In fact, it has been shown that experimentally inducing context change in participants leads them to overestimate the temporal duration of an experiment (Sahakyan & Smith, 2014).

The model implements context drift by using a modified form of the Murdock (1997) context drift equation. The main parameter that drives context drift is γ , which can be seen to represent the proportion of context that is preserved from trial to trial. The overall context match ω for a trial with given lag l from a prior trial can be expressed as:

$$\omega_l = \gamma^l \mu_{ss} \quad (6)$$

If $\gamma = 1$, then there is no context drift and the level of context match to the study list context will be consistent across trials. If $0 < \gamma < 1$, then context drift will occur and the context match to the study list context will decrease exponentially over each successive test trial. As such, increases in context drift will result in the magnitudes of both the self-match and item-noise decreasing (see Equation 4).

Across test trials, context drift occurs for both the stored study list items and stored test items. It is assumed that the first test item is bound to a context that matches perfectly to the

study list context (i.e., $\mu_{ss} = 1$) but that for each successive test item, the context match to the study list will decay exponentially. It is also assumed that participants are able to partially distinguish between study and test contexts, meaning that the context cues used at retrieval have an imperfect match to the current test context.

The context match to the most recently stored test item is represented by the $\mu_{ss,test}$ parameter, which ranges between 0 and 1. The motivation behind this free parameter is that test items **are** likely **to** produce lower interference than study items. Initially, we fixed this parameter to 1 (meaning that both study and test items produced equal interference), but found as a result that the model struggled to capture the LLE. This is because the strong influence of test item interference minimized the relative difference in interference between the short and long study lists. It is also assumed that for both the stored study items and stored test items, the rate of context drift γ is consistent. Note that it is also likely that context drifts across lists, and not just across trials. However, the model as implemented by Osth, Jansson et al. (2018) only accounted for context drift across trials.

The Diffusion Decision Model

The means and variances of the log likelihood ratio distributions from the Osth and Dennis (2015) model are then used to determine distributions of the rate of evidence accumulation (i.e., drift rate) for the DDM. In the DDM, evidence begins at a starting point z and accumulates noisily until it reaches one of two response boundaries: an upper boundary, denoted by a , that corresponds to a ‘YES’ response; or a lower boundary, denoted by 0, that corresponds to a ‘NO’ response. The boundary which is reached first in the evidence accumulation process determines the response that is made (refer to Figure 4 for a visual demonstration). The RT of a decision is calculated as the time taken to reach either boundary from the starting point, plus an initial period of non-decision time t_0 (with variability s_{t0}). The DDM’s z parameter can be seen to reflect the bias of a decision, as a higher starting point

increases the likelihood of evidence reaching the ‘YES’ boundary and a lower starting point increases the likelihood of evidence reaching the ‘NO’ boundary. The a parameter is able to capture the speed-accuracy trade-off of a decision, as wider response boundaries lead to slower RTs and reduced errors, while narrower response boundaries lead to faster RTs and increases errors. In the DDM, higher drift rates increase the accuracy and RTs of decisions. For this modelling exercise, drift rates were allowed to randomly vary across trials, as governed by the variances of the log likelihood ratio distributions from the memory model. While the full diffusion model usually includes variability in the starting point z as a uniform distribution with height s_z , we implemented the model without any starting point variability based on previous work demonstrating that such variability is unnecessary for recognition memory tasks in the DDM (Osth, Bora, Dennis, & Heathcote, 2017).

Figure 4. Schematic of the diffusion decision model (Ratcliff, 1978). Evidence begins at starting point z and accumulates noisily at rate v until either the upper boundary a or lower boundary 0 are reached, resulting in a ‘YES’ or ‘NO’ response respectively. Adapted from Osth, Jansson et al. (2018)

Changes in Bias and Speed-Accuracy Thresholds

In addition to the drift rates, the DDM component of the Osth, Jansson et al. (2018) model allows the response boundary parameter a and starting point parameter z (parameterized as a proportion of a) to change across test trials. These changes are modelled

as linear functions with intercept parameters (a, z) and slope parameters (a_{slope} and z_{slope}). Since it is common for evidence accumulation models to allow these parameters to vary between experimental conditions that the participant is aware of (Ratcliff & McKoon, 2008; Nosofsky et al., 2011), each the four intercept and slope parameters above were allowed to vary across the list length conditions in this work.

Model Selection

The standard version of the Osth, Jansson et al. (2018) model, described above, was the first model used in the model selection exercise, and also provided the basis from which each additional model was constructed. Each of the 19 free parameters included in the standard model are described in Table 2. While the standard model accounted for changes across trials within a list – via the context drift γ , z_{slope} and a_{slope} parameters – it did not account for changes across separate lists. As such, the standard model, like existing models of recognition memory, treated each list as if there was no proactive interference that had accumulated over prior lists in the session.

The second model used was the ‘Proactive Interference (PI) model’, which allowed both study and test items from prior lists to contribute proactive interference to future test lists. The extent to which PI from prior lists affected any given list of the experiment was governed by the model’s item mismatch variability parameter σ_{ti}^2 , in addition to a new inter-list context drift parameter $\gamma_{inter-list}$ (see Table 2). Similar to how the γ parameter represents the proportion of context that is preserved from trial to trial, the $\gamma_{inter-list}$ parameter represents the proportion of context that is preserved from list to list. If $\gamma_{inter-list} = 1$, then there is no context change across lists and all PI generated in prior lists will contribute to the global match of the current list. If $0 < \gamma_{inter-list} < 1$, then context change will occur across lists, and the proportion of PI preserved from a particular list will decrease exponentially with each

successive list. If $\gamma_{inter-list} = 0$, then context changes completely between lists and the current context cue would only match the current list. This means that there would be no build-up of PI over the session (c.f., the standard model), even if the item mismatch variability was high. Note that the $\gamma_{inter-list}$ parameter is comparable to the β_{post}^{recall} parameter in the Lohnas et al. (2015) model that governed the rate of context drift between a period of recall and the subsequent study phase.

It is important to highlight that the inclusion of PI from the experimental session in the global match is a form of background-noise as the PI items mismatch on both the item and context dimensions. The exception is for PI from studied lures which constitutes a form of context-noise as the items match on the item dimension but mismatch on the context dimension. The experimental PI contribution, however, is unique as the degree of context match will be greater than that of pre-experimental background-noise and context-noise that participants enter an experiment with; this reflects the notion that the contextual representations of prior list items are likely more similar to the current list context than from items learned prior to the experiment. Updated expressions for the means and variances of the target, lure and studied lure distributions (with the accumulation of PI included) are provided in Appendix B.

Table 2.

Description and bounds for the free parameters included in each model

Free Parameter	Bounds	Description
t_0	$0 : \infty$	Non-decision time
s_{t0}	$0 : \infty$	Non-decision time variability
z_{short}, z_{long}	$0 : 1$	Starting point (bias) for short and long lists
a_{short}, a_{long}	$0 : \infty$	Response boundary height for short and long lists
c_{short}, c_{long}	$-\infty : \infty$	Response criterion for short and long lists
γ	$0 : 1$	Exponential rate of context drift across test trials

$z_{slope,short} , z_{slope,long}$	$-\infty : \infty$	Linear rate of starting point (bias) change across trials
$a_{slope,short} , a_{slope,long}$	$-\infty : \infty$	Linear rate of response boundary change across trials
σ_{ti}^2	$0 : \infty$	Item mismatch variability (item-noise parameter)
σ_{tt}^2	$0 : \infty$	Item match variability
ρ	$0 : \infty$	Context mismatch variability (context-noise parameter)
r	$0 : \infty$	Learning rate
ρ_{WF}	$0 : 1$	Proportion of ρ to vary with word frequency
$\mu_{ss,test}$	$0 : 1$	Baseline of context match to items learned at test
$\gamma_{inter-list}$	$0 : 1$	Exponential rate of context drift across lists (Not included in the standard model)

Notes: t_0 , s_{t0} , z , z_{slope} , a and a_{slope} parameters are DDM parameters; the remaining parameters are implemented in the front-end memory model

As mentioned earlier, a key motivation for the model selection exercise was to determine whether the changes in performance across cycles are the result of interference or decision-based factors. For instance, while the LLE is commonly attributed to differences in item-noise between lists, Donkin and Nosofsky (2012) found that LLEs in a short-term memory task could be completely explained by changes in response boundaries. To investigate this issue, we devised increasingly comprehensive models that, along with PI accumulation, allowed other inter-list changes to be captured. A third model included separate drift criterion c parameters for each study-test cycle in order for experiment-wide changes in criterion position to be captured. The fourth model incorporated separate r parameters to capture learning rate changes across lists; and the fifth model incorporated separate a parameters to capture response boundary (or speed-accuracy trade-off) changes across lists. Further models included different combinations of these three amendments (see Table 3 for a summary).

Each model was fit within a hierarchical Bayesian framework using differential evolution Markov chain Monte Carlo sampling (DE-MCMC; Turner, Sederberg, Brown, & Steyvers 2013). The use of hierarchical Bayesian methods was advantageous as fitting the diffusion component of the model requires a large number of trials per participant to accurately

estimate its parameters, and hierarchical methods are capable of adjusting (potentially uncertain) parameter estimates for individual participants in light of the overall group estimates (Wagenmakers, 2009). The use of DE-MCMC posterior sampling was preferred to regular MCMC algorithms as it is robust to parameter correlations, which tend to be common in mechanistic models (Turner et al., 2013).

Table 3 below outlines each of the models that were fit to the data (refer to Appendix C for the full list of free parameters used in each model). To evaluate each model, we used the widely applicable information criterion (WAIC; Watanabe, 2010). WAIC is a model selection technique that quantifies the ability of a model to fit a dataset, and subtracts a quantity that corresponds to the model's complexity. Smaller WAIC values indicate that the model is more parsimonious; that is, it fits the data well while avoiding unnecessary complexity. Since WAIC values are difficult to interpret in isolation, Table 3 reports the change in WAIC (Δ WAIC) for each model relative to the standard model; the WAIC value of which has been fixed to zero. As such, models with Δ WAIC < 0 can be considered as superior to the standard model. Furthermore, since WAIC is measured on a log likelihood scale, differences in WAIC that are greater than 10 can be considered large.

As predicted, the baseline PI model (second row) and all other models featuring the PI mechanism were favored substantially (Δ WAIC > 1000) over the standard model (first row). However, the most successful and parsimonious model was the model that added inter-list criterion and response boundary changes to the PI model (seventh row, in bold). For the remainder of this article, this model will be referred to as the 'winning PI model'. While changes in learning across lists seemed a plausible candidate to capture the condition order effects in the data, the Δ WAIC values in Table 3 indicate that this amendment generally reduced model parsimony. For instance, the 'runner-up model' (row 9 of Table 3) produced group-level posterior predictive values that were largely indistinguishable from those of the

winning PI model, despite including 12 additional parameters. It should also be noted that inspection of r estimates for each model fit that allowed changes in learning across cycles revealed that r was predicted to rise linearly across the entire course of the hour-long session; a notion which seemed implausible.

Table 3.

List of amendments included in each model and their corresponding Δ WAIC values

Model	Proactive interference build-up across lists	Criterion changes across lists	Learning changes across lists	Boundary changes across lists	No. of free parameters	Δ WAIC
1	No	No	No	No	19	0
2	Yes	No	No	No	20	-1470
3	Yes	Yes	No	No	43	-2164
4	Yes	No	Yes	No	32	-1090
5	Yes	No	No	Yes	43	-3757
6	Yes	Yes	Yes	No	55	-2636
7	Yes	Yes	No	Yes	66	-4697
8	Yes	No	Yes	Yes	55	-3650
9	Yes	Yes	Yes	Yes	78	-4598

Notes: The first row corresponds to the standard model and the second row corresponds to the PI model. The best fitting model (the winning PI model) is listed as Model 7 above (in bold), and the runner-up model is listed as Model 9. The Δ WAIC values are calculated relative to the standard model (fixed at zero).

In the Supplementary Materials section (found on our OSF page: <https://osf.io/7qxgm/>), we performed a parameter recovery for the winning PI model. This was done by fitting the winning PI model to a novel dataset that was created using the median parameter estimates of its original fit. We found that, despite the model's complexity, its parameters were well recovered. The only parameter which was not recovered for the majority of participants was

non-decision time variability s_{t0} . This was not of great concern as the s_{t0} parameter is not relevant to the inferences that we are interested in making from the model fit. In the Supplementary Materials section, we also present correlations between data and model predictions across participants for the winning PI model. In doing so we demonstrate that the model is successful in accounting for individual participants, and not just group-level averages.

Posterior Predictive Values: Overall Performance

To evaluate how well the winning PI model fit the data, we generated posterior predictive values from the model, averaged them across all participants, and collapsed over word frequency. While we also analysed the posterior predictive values for HF and LF words separately, we did not notice any critical interaction effects, thus, for simplicity's sake, we only report the collapsed values here. In each of the data plots below, the winning PI model's posterior predictive values were represented by 95% highest density intervals (HDIs).

Posterior predictive plots for both the standard model and baseline PI model are presented and discussed in the Supplementary Materials section.

Figure 5 presents the mean proportions of 'YES' responses for targets, lures and studied lures across short and long lists (top panel), as well as the 10th, 50th and 90th percentiles of the correct (middle panel) and error (bottom panel) RT distributions for each of those data points. Overall, the winning PI model provided a strong fit to the data as it was able to capture the overall choice probabilities and RT distributions. While the standard model was able to capture some key trends in the data, such as the increase in FAR for long lists (indicative of an LLE), it was unable to capture the choice proportions for studied lures. As highlighted in Appendix B, the standard model had no mechanism for context change across lists, and so it predicted similar choice proportions for both targets and studied lures. The PI model (and others) were able to capture this key trend with the inclusion of the inter-list context change

parameter $\gamma_{inter-list}$. The mean of the winning PI model's posterior distribution for the group-level $\gamma_{inter-list}$ parameter was 0.66, indicating that the context match to study items from the previous list was roughly two thirds the strength of the context match to study items from the present list. This is well reflected in the data as the FAR_{SL} is roughly two thirds the magnitude of the HR.

Figure 5. Choice probabilities (top), correct RT quantiles (middle) and error RT quantiles (bottom) for targets, lures and studied lures in short and long lists. RT quantiles correspond to the 10th, 50th and 90th percentiles of their respective distributions. Winning PI model predictions represented by 95% HDIs. Note that very small HDIs may be visually obscured by the black markers

Posterior Predictive Values: Performance Across Cycles

Figure 6 below presents the choice proportions for each item type across the first 12 cycles of the experiment, both for short lists (left panel) and long lists (right panel). The fit of the winning PI model is mostly strong as it is able to capture the parallel decreases in HR and FAR_{SL} across lists while maintaining a consistent FAR. In fact, all models that included the accumulation of PI were able to capture the trend. Due to the lack of the PI mechanism, the

standard model predicted that performance was consistent throughout all lists. The most challenging aspect of the data in Figure 6 to capture was the high level of accuracy in the first few cycles. Even the winning PI model appears to slightly underestimate the HR and overestimate the FAR over the early cycles of the experiment, particularly for the long lists. Most models, however, fared considerably worse at capturing these data than the winning PI model. This is due to the fact that the PI mechanism only succeeds in capturing a decrease in performance up until a certain point. Since the $\gamma_{inter-list}$ estimate was reasonably low (0.66), only PI from about three lists prior would have any measurable impact on the current list. For instance, while the context match to items from the most recent list would be scaled by 0.66, the context match to items from the 4th most recent list would be scaled by 0.19. In other words, the PI mechanism predicts a decrease in performance for the first few cycles where PI accumulates, but then predicts a plateau in performance as the level of PI accumulation reaches a sort of equilibrium.

The winning PI model is more successful at capturing the data in question as it allows the response boundary parameter a and drift criterion parameter c to vary over cycles. Regarding the model estimates of these parameters, a series of linear regression analyses revealed that participants shift their speed-accuracy trade-off to favor speed as the experiment progresses, but that speed is slightly more favored in long lists than short lists (a_{short} values decreased over cycle with a slope of -0.03 and intercept of 1.72; a_{long} values decreased over cycle with a slope of -0.02 and intercept of 1.57). Furthermore, it was revealed that participants become more conservative with their criterion over the course of the experiment, but that participants are generally much more liberal during long lists (c_{short} values increased over cycle with a slope of 0.06 and intercept of 0.35; c_{long} values increased over cycle with a slope of 0.08 and intercept of 0.04). The emphasis on speed and lower criterion placement in long lists help the model account for the overall LLE that was found throughout the experiment (c.f., Donkin &

Nosofsky, 2012). More importantly, however, the general increase in speed and conservative responding throughout testing allows the model to account for the constant decrease in accuracy across study-test cycles, and consequently the high levels of accuracy in the early study-test cycles (see Figure 6).

Figure 6. Choice probabilities for targets, lures and studied lures in short (top) and long lists (bottom) across the 12 study-test cycles. Winning PI model predictions represented by 95% HDIs

Posterior Predictive Values: Discriminability

Despite the successes of the winning PI model relative to the weaker models, the struggle in capturing the early performance peak is even more evident in Figure 7, which presents d' for short and long lists across the 12 cycles of the experiment. Figure 7, however, also demonstrates the key strength of the winning PI model, in that it was able to capture an LLE on the odd-numbered cycles, and a null LLE on the even-numbered cycles. This trend can be seen more clearly in the right panel of Figure 8, which presents the data and winning PI model predictions for the overall difference in d' between short and long lists, alongside that

for odd-numbered and even-numbered cycles. The left panel of Figure 8 presents the same data but with the predictions of the standard model. In contrast to the winning PI model predictions, the standard model predicts a similarly sized LLE across all cycles of the experiment.

The inclusion of the PI mechanism in the winning PI model allowed the LLE in odd cycles and null LLE in even cycles to be captured. Since a large amount of PI from prior cycles contributed to the future cycles, cycles which equated the overall number of observed items between those currently studying a short list or a long list (i.e., the even-numbered cycles) also led to the rough equation of interference between them, and thus the null LLE. Note, however, that the winning PI model prediction for the LLE in even-numbered cycles is not completely null. This is due to the fact that $\gamma_{inter-list}$ was estimated to be lower than 1, meaning that the total amount of within-experiment interference (i.e., item-noise + PI) was not fully equated in the even-numbered cycles, only closely equated. Despite this, the PI mechanism in the winning PI model provides support for the MINERVA 2 model proposed by Brandt et al. (2019), given that it similarly allowed items from prior lists to contribute interference to future lists. Note that the function of the $\gamma_{inter-list}$ parameter in our model approximates the function of the varying learning rates for prior lists in their model.

Another key aspect of model that allowed the LLE to be captured was the inclusion of the $\mu_{ss,test}$ parameter that constitutes the baseline context match of the context cues to the test list context (prior to context drift). Without this parameter, it was assumed that the context match of the cues to the test items was equivalent to that of the study items (i.e., $\mu_{ss} = 1$). This proved to be an issue, however, as each test item produced the same level of interference as each study item, meaning that the interference difference between short and long lists was reduced and the LLE would be very hard to capture. To clarify, while the difference between 20 and 80 items on short and long lists respectively is large, the assumption of equal context

match to test items effectively increases each list length by 32 items (the length of the test list), making the short and long lists 52 and 112 items long respectively, and thus reducing the relative difference between them. When the $\mu_{ss,test}$ parameter was included, it was estimated to be approximately .40 for each model, meaning that participants were able to partially distinguish the test context from the study list context, and that each test item would produce less than half the interference of each study item. This allowed the LLE to be captured as the bulk of interference stemmed from the study lists, thus emphasizing the relative differences between short and long lists.

Another aspect of the data to note is that the values of d' across cycle eventually reach a plateau (see Figure 7). This contrasts with the constant decrease in the HR found in Figure 6. This trend may be attributable to the model's PI mechanism that, as mentioned before, only allows PI from roughly the three most recent lists to have any noticeable impact on the current global match. To reiterate, since the $\gamma_{inter-list}$ parameter causes the context match to a prior list to decline exponentially as each new list is learnt, the context match to a particular prior list becomes negligibly small after about three more lists have passed. If only the most recent lists produce impactful PI, then an increasing number of prior lists should produce negligible increases in PI and performance should remain relatively constant.

Figure 7. d' values for short and long lists across the 12 study-test cycles. Winning PI model predictions represented by 95% HDIs. Winning PI model predictions for short and long lists are offset to the left and right respectively

Figure 8. The LLE across all cycles, odd cycles and even cycles, as predicted by the standard model (left) and winning PI model (right). Standard and winning PI model predictions represented by 95% HDIs

Estimates of Interference Contributions

While the difference between the standard and winning PI models in their ability to

capture key trends of the data is stark, the difference in conclusions that can be drawn from either are even more notable. Figure 9 below presents the proportions of item-noise, context-noise and background-noise in the first long list of the experiment as predicted by the standard model (left panels) and winning PI model (right panels) when their group-level parameter estimates are input into Equation 5. Predictions for LF words are presented in the top row and predictions for HF words are presented in the bottom row.

In Figure 9, the plots for the standard model (left panels) suggest that there is a dominance of background-noise at retrieval relative to item-noise and context-noise. This pattern is similar to those observed in previous investigations (e.g., Osth & Dennis, 2015; Osth, Jansson et al., 2018). The plots for the winning PI model (right panels), however, suggest that there is a dominance of item-noise at retrieval, and that the contribution of background-noise is relatively minimal. This observation is striking as it completely reverses the conclusions made by Osth and Dennis (2015). It is likely that the high item-noise estimates of the winning PI model are the result of capturing the large LLEs in the odd-numbered cycles (the first cycle in particular). Since the large LLEs were effectively hidden in the data, the standard model **would have** been unable to detect them, and thus provided no explanation for their occurrence. It is also likely that the high item-noise estimates are attributable to the model's prediction of heavy PI accrual and increases in **speedy and biased** responding. Both factors lead to a reduction in performance over time for both short and long lists. A consequence of this is that the relative differences in performance between short and long lists is decreased as well (i.e., the LLE decreases). In order for the winning PI model to jointly account for decreases in performance while maintaining a strong LLE on the odd cycles, large increases in item-noise need to be employed. It is also worth noting that both the standard and winning PI models estimate higher proportions of context-noise for HF words than LF words. This is in accordance with the findings of Osth and Dennis (2015) and

reflects the notion that HF words, by definition, have been witnessed in a greater number of pre-experimental contexts than LF words.

Figure 9. Standard model (left) and winning PI model (right) predictions for the proportions of item-noise, context-noise and background-noise at retrieval in the first long list. Predictions for LF and HF words are presented in the top and bottom rows respectively. Standard and winning PI model predictions represented by 95% HDIs

Test Position Effects in the Final List

As highlighted in the introduction, each session concluded with a lengthy test list of 160 test items, the purpose of which was to provide a measure of the TPE. In our results, we found that the HR decreased throughout the final list, but that the FAR remained constant. This pattern indicates a TPE and is consistent with prior research (Criss et al., 2011). This TPE can be observed in Figure 10 which presents choice proportions (top panel), correct RT

quantiles (middle panel) and error RT quantiles (bottom panel) for targets and lures across blocks of 40 trials in the final list. The figure also demonstrates that the winning PI model was successful in providing a strong fit to each of the data points. Note that the model was fit to the 160 individual test trials and that the blocking is only done here for the purpose of visualization.

The winning PI model's estimates of item mismatch variability σ_{ti}^2 and context drift γ appeared to capture the TPE as the former allowed each additional test item to contribute additional interference to the global match (thus decreasing performance over time), while the latter ensured that with each successive trial the context match to the study list decreased (thus reducing the HR over time). However, the inclusion of slope parameters for a and z across trials was necessary to capture the full extent of the decrease in HR across blocks (i.e., the TPE). This may be due to the fact that estimates of the $\mu_{ss,test}$ parameter (i.e., the context match to the test items) were quite low (.36 for the winning PI model), meaning that each additional test item produced low levels of interference in the global match. The additional parameters, however, were not needed to capture the consistent FAR. Note how the changes across blocks mirror the changes across cycles in Figure 6.

Figure 10. Choice probabilities (top), correct RT quantiles (middle) and error RT quantiles (bottom) for targets and lures in the final list. RT quantiles correspond to the 10th, 50th and 90th percentiles of their respective distributions. Winning PI model predictions represented by 95% HDIs

We also investigated whether the models made different predictions regarding the cause of the TPE in the final list of the experiment. This was determined using a correlational analysis that was derived from Osth, Jansson et al. (2018). Note that Osth, Jansson et al. established that context drift was the primary cause of the TPE. Figure 11 presents the standard (top row) and winning PI model's (bottom row) predictions for the correlations between each participant's TPE magnitude and their respective parameter estimates of σ_{ti}^2 , γ , and $\alpha_{slope, long}$. The magnitude of the TPE was calculated by subtracting the predicted d' value for the 160th trial of the final list from the predicted d' value for the 1st trial (d'_{change}). The predictions were generated by taking 10% of the posterior samples and generating 5,000 simulations for each. As reflected in the Pearson r values in Figure 11 below, the standard model predicted a sizable correlation between d'_{change} and γ , but not between d'_{change} and the other parameters, suggesting that the TPE is primarily explained by context drift. However, the winning PI model demonstrated a similarly sizable correlation between d'_{change} and σ_{ti}^2 , suggesting that increases in item-noise throughout testing also had an influence on the TPE. While the standard model was indeed successful at capturing the TPE – suggesting that context drift was sufficient to drive the effect – the increases in item-noise needed to account for the LLE in the winning PI model meant that item-noise had an impact on the TPE as well. Another explanation is that the build-up of PI from prior lists decreases the extent of the TPE, due to the overall drop in performance, and that an increase in item-noise from the test items was needed to maintain the magnitude of the TPE. Regardless, the winning PI model predictions contradict the conclusions of Osth, Jansson et al., who found that changes

in context drift were the sole cause of the TPE.

Figure 11. Standard model (top) and winning PI model (bottom) predictions for the correlations of the TPE (d'_{change}) with σ_{ti}^2 , γ , and $a_{slope,long}$. Pearson r values for each correlation are included at the top of their respective plots. Red lines represent the line of best fit for each correlation

General Discussion

The first aim of the present investigation was to replicate the finding of Brandt et al. (2019) that the list length effect (LLE) in recognition memory can be masked by a build-up of proactive interference (PI) from prior lists. Given that the design of Brandt et al. employed picture stimuli and just two study-test cycles, we aimed to replicate their finding using word stimuli across 12 study-test cycles. The second, and more important, aim of the present investigation was to fit the results using a computational model that would be able to measure the sources of interference at retrieval (i.e., the contributions of item-nose, context-noise, and

background-noise). In doing so, we would be able to determine whether the conclusion of Osth and Dennis (2015) – that pre-experimental background-noise is the dominant source of interference in recognition – would hold when the model had to account for the build-up of PI across lists. Given that Brandt et al. did not consider how RTs or decision-related factors may have influenced their findings, we analysed our results using variants of the Osth, Jansson et al. (2018) model, that integrates the memory retrieval components of the Osth and Dennis (2015) model with the diffusion decision model (DDM; Ratcliff, 1978) to jointly account for choice probabilities and RT distributions. Allowing the model to be constrained by RTs was important as it enabled performance changes resulting from memory-based factors (e.g., interference and learning) to be distinguished from those resulting from decision-based factors (i.e., changes in response bias and the speed-accuracy trade-off). This distinction is non-trivial as it has been shown in short-term memory that increases in RT with increasing list length are attributable to higher speed-accuracy thresholds and not higher interference (Donkin & Nosofsky, 2012).

Regarding our first aim, the present investigation was successful in replicating the findings of Brandt et al. (2019): a strong LLE was found in the first cycle of the experiment and a null (or reversed) LLE was found in the second cycle. Furthermore, it was shown that the same result could be replicated over the span of several lists: a strong LLE was observed across all 6 of the odd-numbered cycles in the experiment (where within-experiment interference was unequal between those currently studying a short list or a long list) and a null LLE was observed across the 6 even-numbered cycles (where within-experiment interference was equated). Regarding our second aim, the present investigation demonstrated how model estimates of noise proportions (i.e., the proportions of item-noise, context-noise and background-noise at retrieval) are affected when accounting for the build-up of PI throughout an experiment. The first model we fit, referred to as the ‘standard model’, was

similar to existing recognition models in that it had no mechanism to account for proactive interference build-up across lists. The standard model predicted that there was a dominance of background-noise at retrieval, supporting the findings of Osth and Dennis (2015). The best fitting model, referred to as the ‘winning PI model’, was distinct from the standard model in two key ways: 1) it allowed items from prior lists to contribute PI to future lists; and 2) it allowed criterion position and the speed-accuracy trade-off (the DDM’s response boundary parameter a) to vary across lists. In contrast to the standard model, the winning PI model predicted a dominance of item-noise at retrieval, completely reversing the conclusions made by Osth and Dennis (2015).

The fact that the present study found robust LLEs in the odd-numbered lists and not in the even-numbered lists supports the arguments of Brandt et al. (2019) and the assumptions of their MINERVA 2 model: that the LLE occurs naturally, but that the learning of subsequent lists can lead to an equation of interference between groups (i.e., those currently studying a short list or a long list) that disguises the effect. In the odd-numbered cycles of the experiment, participants currently learning a long list would have been exposed to more words throughout the entire experiment than those currently studying a short list, meaning that the amount of interference between the groups would be unequal and an LLE would be observed. In the even-numbered cycles, however, participants would learn the opposite list length to the prior cycle, meaning that all participants would have been exposed to the same total of words throughout the experiment. Consequently, the amount of interference between the groups would be roughly equal and a null LLE would be observed.

Support for the arguments of Brandt et al. (2019) and their MINERVA 2 model also came from our modelling exercise. The best fitting models we applied, particularly the winning PI model, were able to account for the pattern of LLEs across cycles by allowing items from prior lists to contribute PI to later lists. Inspection of these models’ posterior

distributions for the inter-list context drift parameter $\gamma_{inter-list}$ suggested that roughly two-thirds of PI from one cycle would carry over in the next. In contrast, the standard model had no mechanism to account for PI build-up within the experiment, and consequently was only able to predict an LLE of consistent magnitude across all cycles (not the pattern we observed). It should also be noted that while the winning PI model improved model fit by allowing changes in criterion and the speed-accuracy trade-off across cycles, these amendments were not essential in accounting for the trends in our data. In short, it is difficult to comprehend how a model would account for these data without the assumption of PI building-up within the experiment.

Given that the winning PI model estimated high levels of item-noise at retrieval, the present research supports global matching models such as SAM (Gillund & Shiffrin, 1984), MINERVA 2 (Hintzman, 1988), and the retrieving-effectively-from-memory model (REM; Shiffrin & Steyvers, 1997) that make similar predictions. Conversely, the present research is at odds with context-noise models like BCDMEM (Dennis & Humphreys, 2001). The discrepancy between our model predictions and those of Osth and Dennis (2015) – which suggested a dominance of background-noise – most likely lies in the fact that our winning PI model was able to detect large LLEs that would be effectively hidden in the data from models that don't account for the accrual of PI across lists. Increases in item-noise were necessary for the winning PI model to be able to capture these large LLEs (particularly that of the first cycle). Since the standard model had no PI accrual mechanism, it had no means of accounting for the alternating LLEs and null LLEs at each cycle, and instead, was only capable of accounting for the experiment-wide LLE. As shown in Figure 8, the overall LLE for the experiment was considerably smaller than the LLE for the odd-numbered cycles. It follows that smaller levels of item-noise would have been required to produce the effect. The discrepancy between our findings and those of Osth and Dennis (2015) may also lie in the

fact that the winning PI model allowed criterion position and the speed-accuracy trade-off to change across lists. Since the model predicted that responding got speedier and more biased over lists, both things which lead to performance decreases, higher levels of item-noise would have been needed to maintain the magnitude of the large LLEs in the odd-numbered lists.

The present results dovetail with recent work on global similarity effects in recognition memory (Osth, Shabahang, Mewhort, & Heathcote, in revision). Item-noise models naturally predict that performance declines as the global similarity between probes and the contents of memory is increased. This prediction is commonly tested using category length designs, where it is found that increases in the length of studied semantic categories degrades performance (e.g., Shiffrin, Huber, & Marinelli, 1995). Context-noise models have no natural means for predicting such effects and instead have appealed to additional mechanisms such as generation of implicit associative responses (IARs: Dennis & Humphreys, 2001) or strategies where participants cue with the labels of the studied categories when the categorical structure is noticed (Dennis & Chapman, 2010). However, Osth et al. (in revision) found that graded global similarity effects, which were measured using semantic representations from the BEAGLE model (Jones & Mewhort, 2007), were still found with lists of unrelated words that lack obvious categorical structure. Such results additionally undermine the viability of a pure context-noise account of recognition memory.

A reviewer highlighted the similarity of the PI mechanism in our model to the temporal cueing mechanism of the scale-independent memory, perception, and learning model (SIMPLE; Brown, Neath, & Chater, 2007). According to SIMPLE, items in memory are organized on a temporal continuum that ranges from past to present. This continuum is logarithmically compressed, such that temporally distant memories become more compacted and less discriminable than temporally recent memories. When recall of items from a particular list is attempted, a temporal cue associated with that list is implicitly generated.

Such cues, however, can be imprecise, meaning that lists with high temporal proximity to the target list may be confused with each other, allowing for extra-list item intrusions to occur. Items from much older lists may also be retrieved, but the likelihood of doing so decreases as the temporal match between them and the cue decreases (c.f., the PI mechanism of our model). Despite this similarity, it is worth noting that it would be difficult to apply the SIMPLE model to recognition datasets such as ours. This is because the last temporal occurrence of any given lure item (barring studied lures) is unknown. Additionally, it would be challenging for SIMPLE to account for the accumulation of PI in our experiment. In SIMPLE, the discriminability of any prior list is a function of the amount of time that has elapsed between them. However, the group parameter estimates of our model indicated that context changes at a greater rate between lists ($\gamma_{inter-list} = .660$) than within lists ($\gamma = .996$). A purely time-based model, such as SIMPLE, would not be able to account for this phenomenon without additional mechanisms.

Despite the results of the present experiment being convincing, there is at least one limitation with our methodology that could have negatively impacted the conclusions drawn. Namely, our list length design did not control for contextual reinstatement by employing a delay period after both short and long lists. As demonstrated by Dennis et al. (2008), the use of a delay period can lead to an overall null LLE in recognition. While the present study suggests that the LLE is naturally very strong, and that null LLEs are the result of averaging artefacts, it could be the case that the strong LLEs we observed were partly caused by the lack of a delay period, and not simply a dominance of item-noise. In defence of our results, Brandt et al. (2019) utilized a delay period in their design (roughly 90 seconds for the short lists and 30 seconds for the long lists) while still reporting a strong LLE in the initial cycles, suggesting that the LLE is not a merely a consequence of not controlling for contextual reinstatement. Additionally, it could be argued that the reduction in the LLE by using a delay

period is not the result of controlling contextual reinstatement, but rather a disguise effect. Since lengthy delay periods negatively affect accuracy at test (Cary & Reder, 2003), it is possible that differences in accuracy between short and long lists are simply hidden. In other words, the LLE may not be an artefact of a particular design; instead it may be the null LLE that is a design artefact.

Implications for Explanations of the Test Position Effect

The present experiment also involved a lengthy test list that participants completed at the end of each session, the purpose of which was to provide a measure of the test position effect (TPE). One of the aims of the modelling exercise was to determine the cause of the TPE that we measured. Item-noise models posit that the TPE is caused by a build-up of interference from item presentations during the test list (Criss et al., 2011; Gillund & Shiffrin, 1984; Shiffrin & Steyvers, 1997). In contrast, context-noise models posit that the TPE is caused by changes in the context representation that occur as a consequence of memory retrieval (i.e., context drift; Osth & Dennis, 2015). Osth, Jansson et al. (2018) sought to resolve this issue by simultaneously measuring the influence of both item-noise and context drift, as well as changes in the speed-accuracy trade-off, on the TPE. Their model was applied to data that provided numerous constraints; in particular, their first experiment manipulated list length and found an approximately null list-length effect and a substantial TPE. Parameter estimates from their model suggested that the TPE was primarily due to the influence of context drift and not the accumulation of item-noise. However, a limitation of their analyses was that they did not incorporate the build-up of PI into their model.

In the present investigation, we performed a similar analysis to Osth, Jansson and colleagues but included models that allowed for a build-up of PI throughout the session. While the fit of the standard model suggested that the TPE was attributable to context drift alone, the fit of the winning PI model revealed that the TPE was attributable to *both* item-

noise and context drift. Why is this the case? Note that while higher levels of values of item mismatch variability σ_{ti}^2 in our model increase item-noise and thus increase the magnitude of both the LLE and TPE, they also increase the PI from prior lists. If there is large PI, the relative impact of learning additional test items would produce relatively weak levels of interference. Mathematically, this is because increasing the number of items increases variance in a linear fashion, and since d' is on a standard deviation scale, the number of items bears a square root relationship to memory performance, such that each item produces progressively *less* interference.

While increasing item mismatch variability increases PI, context drift mitigates against the build-up of PI. Changes in context prevent prior learned items from matching the current context and reduces their interference, such that more recently learned items will exert greater interference. Thus, it's possible that both mechanisms are required to explain why the TPE can be large despite the accumulation of an entire session's worth of PI.

Conclusion

In summary, the present study established evidence for an LLE in recognition memory by replicating Brandt et al. (2019) with word stimuli, several lists and a two-session design. More crucially, by allowing a computational model to account for the build-up of proactive interference across lists, parameter estimates were derived that were indicative of more item-noise being present at retrieval than previously believed. Overall, it is concluded that while the LLE is primarily caused by interference from other list items, the TPE is caused by both item-noise and context drift. Future research should consider the impact of proactive interference and condition order effects on data trends, as well as develop computational models that can account for these differences.

References

- Bowles, N. L., & Glanzer, M. (1983). An analysis of interference in recognition memory. *Memory & Cognition*, 11(3), 307-315.
- Brandt, M., Zaiser, A. K., & Schnuerch, M. (2019). Homogeneity of item material boosts the list length effect in recognition memory: A global matching perspective. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45(5), 834-850.
- Brown, G. D., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, 114(3), 539-576.
- Brown, J. (1958). Some tests of the decay theory of immediate memory. *Quarterly Journal of Experimental Psychology*, 10(1), 12-21.
- Cary, M., & Reder, L. M. (2003). A dual-process account of the list-length and strength-based mirror effects in recognition. *Journal of Memory and Language*, 49(2), 231-248.
- Clark, S. E., & Gronlund, S. D. (1996). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin & Review*, 3(1), 37-60.
- Criss, A. H., & Shiffrin, R. M. (2005). List discrimination in associative recognition and implications for representation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(6), 1199-1212.
- Criss, A. H., Malmberg, K. J., & Shiffrin, R. M. (2011). Output interference in recognition memory. *Journal of Memory and Language*, 64(4), 316-326.
- Davis, C. J. (2005). N-Watch: A program for deriving neighborhood size and other psycholinguistic statistics. *Behavior Research Methods*, 37(1), 65-70.
- Dennis, S., & Chapman, A. (2010). The inverse list length effect: A challenge for pure exemplar models of recognition memory. *Journal of Memory and Language*, 63, 416-424.

- Dennis, S., & Humphreys, M. S. (2001). A Context Noise Model of Episodic Word Recognition. *Psychological Review*, 108(2), 452-478.
- Dennis, S., Lee, M. D., & Kinnell, A. (2008). Bayesian analysis of recognition memory: The case of the list-length effect. *Journal of Memory and Language*, 59(3), 361-376.
- Donkin, C., & Nosofsky, R. M. (2012). The structure of short-term memory scanning: An investigation using response time distribution models. *Psychonomic Bulletin & Review*, 19(3), 363-394.
- Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin & Review*, 9(1), 59-79.
- Gillund, G., & Shiffrin, R. M. (1984). A Retrieval Model for Both Recognition and Recall. *Psychological Review*, 91(1), 1-67.
- Glanzer, M., & Adams, J. K. (1985). The mirror effect in recognition memory. *Memory & Cognition*, 13(1), 8-20.
- Gronlund, S. D., & Elam, L. E. (1994). List-length effect: Recognition accuracy and variance of underlying distributions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(6), 1355-1369.
- Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, 95, 528-551.
- Hintzman, D. L., Caulton, D. A., & Levitin, D. J. (1998). Retrieval dynamics in recognition and list discrimination: Further evidence of separate processes of familiarity and recall. *Memory & Cognition*, 26(3), 449-462.
- Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: a theory for episodic, semantic, and procedural tasks. *Psychological Review*, 96(2), 208-233.

- Humphreys, M. S., Pike, R., Bain, J. D., & Tehan, G. (1989). Global matching: A comparison of the SAM, Minerva II, Matrix, and TODAM models. *Journal of Mathematical Psychology*, 33(1), 36-67.
- Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114, 1-37.
- Keppel, G., & Underwood, B. J. (1962). Proactive inhibition in short-term retention of single items. *Journal of Verbal Learning and Verbal Behavior*, 1(3), 153-161.
- Kinnell, A., & Dennis, S. (2011). The list length effect in recognition memory: An analysis of potential confounds. *Memory & Cognition*, 39(2), 348-363.
- Kinnell, A., & Dennis, S. (2012). The role of stimulus type in list length effects in recognition memory. *Memory & Cognition*, 40(3), 311-325.
- Lohnas, L. J., Polyn, S. M., & Kahana, M. J. (2015). Expanding the scope of memory search: Modeling intralist and interlist effects in free recall. *Psychological Review*, 122(2), 337-363.
- Mewhort, D. J. K., Shabahang, K. D., & Franklin, D. R. J. (2018). Release from PI: An analysis and a model. *Psychonomic Bulletin & Review*, 1-19.
- Murdock, B. B., & Kahana, M. J. (1993a). Analysis of the list-strength effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(3), 689.
- Murdock, B. B., & Kahana, M. J. (1993b). List-strength and list-length effects: Reply to Shiffrin, Ratcliff, Murnane, and Nobel (1993). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(6), 1450-1453.
- Murdock, B. B. (1997). Context and mediators in a theory of distributed associative memory (TODAM2). *Psychological Review*, 104(4), 839-862.

- Murnane, K., & Shiffrin, R. M. (1991). Interference and the representation of events in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(5), 855-874.
- Nosofsky, R. M., Little, D. R., Donkin, C., & Fific, M. (2011). Short-term memory scanning viewed as exemplar-based categorization. *Psychological Review*, 118(2), 280-315.
- Ohrt, D. D., & Gronlund, S. D. (1999). List-length effect and continuous memory: Confounds and solutions. In C. Izawa (Ed.), *On human memory: Evolution, progress, and reflections on the 30th anniversary of the Atkinson-Shiffrin model* (pp. 105–125). Mahwah, NJ: Lawrence Erlbaum Associates.
- Osth, A. F., Bora, B., Dennis, S., & Heathcote, A. (2017). Diffusion vs. linear ballistic accumulation: Different models, different conclusions about the slope of the zROC in recognition memory. *Journal of Memory and Language*, 96, 36-61.
- Osth, A. F., & Dennis, S. (2015). Sources of Interference in Item and Associative Recognition Memory. *Psychological Review*, 122(2), 260-351.
- Osth, A. F., Dennis, S., & Heathcote, A. (2017). Likelihood ratio sequential sampling models of recognition memory. *Cognitive Psychology*, 92, 101-126.
- Osth, A. F., Fox, J., McKague, M., Heathcote, A., & Dennis, S. (2018). The list strength effect in source memory: Data and a global matching model. *Journal of Memory and Language*, 103, 91-113.
- Osth, A. F., Jansson, A., Dennis, S., & Heathcote, A. (2018). Modeling the dynamics of recognition memory testing with an integrated model of retrieval and decision making. *Cognitive Psychology*, 104, 106-142.
- Osth, A. F., Shabahang, K. D., Mewhort, D. J. K., & Heathcote, A. (in revision). Global semantic similarity effects in recognition memory: Insights from BEAGLE representations and the diffusion decision model.

- Peterson, L., & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology*, 58(3), 193-198.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59-108.
- Ratcliff, R., Clark, S. E., & Shiffrin, R. M. (1990). List-strength effect: I. Data and discussion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16(2), 163-178.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873-922.
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56(5), 356-374.
- Sahakyan, L., & Smith, J. R. (2014). “A long time ago, in a context far, far away”: Retrospective time estimates and internal context change. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(1), 86-93.
- Schulman, A. I. (1974). The declining course of recognition memory. *Memory & Cognition*, 2(1), 14-18.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM-retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4(2), 145-166.
- Shiffrin, R. M., Huber, D. E., & Marinelli, K. (1995). Effects of category length and strength on familiarity in recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 267-287.
- Strong Jr, E. K. (1912). The effect of length of series upon recognition memory. *Psychological Review*, 19(6), 447-462.
- Turner, B. M., Sederberg, P. B., Brown, S. D., & Steyvers, M. (2013). A method for efficiently sampling from distributions with correlated dimensions. *Psychological Methods*, 18(3), 368-384.

- Underwood, B. J. (1978). Recognition memory as a function of length of study list. *Bulletin of the Psychonomic Society*, 12(2), 89-91.
- Wagenmakers, E. J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, 14(5), 779-804.
- Wagenmakers, E. J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology*, 21(5), 641-671.
- Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11(Dec), 3571-3594.

Appendix A

Analytic Derivations of the Osth and Dennis (2015) Model

As highlighted by Humphreys, Pike, Bain and Tehan (1989), if a target item is presented at test, the memory matrix M of Equation 2 can be decomposed to reflect all the types of matches to stored memories:

$$\begin{aligned}
 s &= (I'_t \otimes C'_s) \cdot [r(I_t \otimes C_s) && \text{Self-Match} \\
 &+ \sum_{i \in L, i \neq t} r(I_i \otimes C_s) && \text{Item-Noise} \\
 &+ \sum_{u \in P, u \neq s} (I_t \otimes C_u) && \text{Context-Noise} \\
 &+ \sum_{u \in P, u \neq s, z \notin L} (I_z \otimes C_u)] && \text{Background-Noise}
 \end{aligned} \tag{A1}$$

where i represents items on the study list other than the target item t , u represents a context from a set of all prior contexts (P) to the study list s , and z represents items from prior contexts that were not featured on the study list.

The equation can then be expressed in terms of 1) matches and mismatches between the cued item vectors and the stored item vectors, and 2) matches and mismatches between the cued context vectors and stored context vectors:

$$\begin{aligned}
 s &= r(I'_t \cdot I_t)(C'_s \cdot C_s) && \text{Self-Match} \\
 &+ \sum_{i \in L, i \neq t} r(I'_t \cdot I_i)(C'_s \cdot C_s) && \text{Item-Noise} \\
 &+ \sum_{u \in P, u \neq s} (I'_t \cdot I_t)(C'_s \cdot C_u) && \text{Context-Noise} \\
 &+ \sum_{u \in P, u \neq s, z \notin L} (I'_t \cdot I_i)(C'_s \cdot C_u) && \text{Background-Noise}
 \end{aligned} \tag{A2}$$

The dot products that comprise the matches and mismatches between the item vectors and between the context vectors are then parameterized as normal distributions:

$$\begin{aligned}
I'_t \cdot I_t &\sim \text{Normal}(\mu_{tt}, \sigma_{tt}^2) && \text{Item Match} \\
I'_t \cdot I_i &\sim \text{Normal}(0, \sigma_{ti}^2) && \text{Item Mismatch} \\
C'_s \cdot C_s &\sim \text{Normal}(\mu_{ss}, 0) && \text{Context Match} \\
C'_s \cdot C_u &\sim \text{Normal}(0, \sigma_{su}^2) && \text{Context Mismatch}
\end{aligned} \tag{A3}$$

The means and variances of each of the distributions are the parameters of the model.

Note that the means of the mismatch distributions and the variance of the context match distribution are set to zero in order to avoid covariances.

Equations for the means and variances of the target and lure distributions can be derived by substituting the distribution expressions of Equation A3 into Equation A2. Since the noise terms of Equation A2 (i.e., self-match, item-noise, context-noise and background-noise) each involve the multiplication of two normal distributions, they each result in a modified Bessel function of the third kind with the following means and variances:

$$\begin{aligned}
E(X_1 X_2) &= \mu_1 \mu_2 \\
V(X_1 X_2) &= \mu_1^2 \sigma_2^2 + \mu_2^2 \sigma_1^2 + \sigma_1^2 \sigma_2^2
\end{aligned} \tag{A4}$$

The means and variances for the target and lure distributions are thus:

$$\begin{aligned}
\mu_{old} &= r \mu_{tt} \mu_{ss} \\
\mu_{new} &= 0
\end{aligned} \tag{A5}$$

$$\begin{aligned}
\sigma_{old}^2 &= r^2 (\mu_{ss}^2 \sigma_{tt}^2) + && \text{Self-Match} \\
&+ r^2 (l - 1) (\mu_{ss}^2 \sigma_{ti}^2) + && \text{Item-Noise} \\
&+ m (\mu_{tt}^2 \sigma_{su}^2 + \sigma_{tt}^2 \sigma_{su}^2) + && \text{Context-Noise} \\
&+ n (\sigma_{ti}^2 \sigma_{su}^2) && \text{Background-Noise}
\end{aligned} \tag{A6}$$

$$\begin{aligned}
\sigma_{new}^2 &= r^2 l (\mu_{ss}^2 \sigma_{ti}^2) + && \text{Item-Noise} \\
&+ m (\mu_{tt}^2 \sigma_{su}^2 + \sigma_{tt}^2 \sigma_{su}^2) + && \text{Context-Noise} \\
&+ n (\sigma_{ti}^2 \sigma_{su}^2) && \text{Background-Noise}
\end{aligned} \tag{A7}$$

where l represents the length of the study list, m represents the number of prior occurrences of the target item, and n represents the number of learned experiences in memory that do not correspond to a study list item or the study list context.

Equations A6 and A7 can be simplified by 1) replacing the μ_{ss}^2 factor with the ω^2 factor that corresponds to the mean level of context match after context drift has occurred (refer to the section below ‘The Context Drift Equation’ for details on the calculation of ω); 2) combining the m and σ_{su}^2 factor into a ρ factor that corresponds to the total context mismatch variability; 3) and replacing the entire background-noise term with a single β parameter. The simplified equations are as follows:

$$\begin{aligned} \sigma_{old}^2 = & r^2(\omega^2 \sigma_{tt}^2) + & \text{Self-Match} \\ & r^2(l-1)(\omega^2 \sigma_{ti}^2) + & \text{Item-Noise} \\ & (\mu_{tt}^2 \rho + \sigma_{tt}^2 \rho) + & \text{Context-Noise} \\ & \beta & \text{Background-Noise} \end{aligned} \quad (A8)$$

$$\begin{aligned} \sigma_{new}^2 = & r^2 l (\omega^2 \sigma_{ti}^2) + & \text{Item-Noise} \\ & (\mu_{tt}^2 \rho + \sigma_{tt}^2 \rho) + & \text{Context-Noise} \\ & \beta & \text{Background-Noise} \end{aligned} \quad (A9)$$

In addition to the items from the study list, new items (i.e., lures) are also learnt throughout the course of testing. Test items can be modelled by including additional item-noise terms (row 2 of Equation A8, row 1 of Equation A9) for each new test item as it is learnt, with $l = 1$. For each new term, the context match ω will initially equal 1, but will decay as further items are added due to the effect of context drift.

Finally, HF and LF words are modelled using separate ρ factors. This is accomplished by sampling an additional ρ_{WF} parameter than corresponds to the proportion of context mismatch variability that varies with respect to word frequency. The calculations for the separate HF and LF factors for ρ are as follows:

$$\begin{aligned}\rho_{HF} &= p + (pp_{WF}) \\ \rho_{LF} &= p - (pp_{WF})\end{aligned}\tag{A10}$$

The Likelihood Ratio Transformation

In the Osth and Dennis (2015) model, recognition decisions are formed on the basis of likelihood ratios. The purpose of this inclusion is to allow the model to produce the mirror effect (Dennis & Humphreys, 2001; Shiffrin & Steyvers, 1997). Following the method of Osth, Dennis and Heathcote (2017), log-likelihood ratios λ for targets and lures are derived from normal distributions that are calculated in terms of discrimination d and the relative variability of the target distribution S . The expressions for d and S are as follows:

$$\begin{aligned}d &= \mu_{old}/\sigma_{new} \\ S &= \sigma_{old}/\sigma_{new}\end{aligned}\tag{A11}$$

Using d and S , the means and standard deviations of the λ distributions can be calculated:

$$\begin{aligned}\mu_{\lambda_{new}} &= -\left(\frac{d^2}{2}\right)\left(\frac{S^2 + 3}{4S^2}\right) + \log(S) \\ \mu_{\lambda_{old}} &= d^2 \frac{S^2 + 1}{2S^2} + \mu_{\lambda_{new}} \\ \sigma_{\lambda_{new}} &= d \frac{S^2 + 1}{2S^2} \\ \sigma_{\lambda_{old}} &= S\sigma_{\lambda_{new}}\end{aligned}\tag{A12}$$

The Context Drift Equation

Context drift is governed by the γ parameter, a scalar than ranges from zero to one. The overall context match ω for a trial a with given lag l from a prior trial can be expressed as:

$$\omega_l = \gamma^l \mu_{ss}\tag{A13}$$

If $\gamma = 1$, then context will not drift from trial to trial. If $0 < \gamma < 1$, then the overall context match ω will decrease exponentially as lag l increases. To capture context match to the study list items on the first test trial, μ_{ss} is fixed to 1. To capture context match of a test item to the

immediately preceding test item, μ_{ss} is replaced with a free parameter $\mu_{ss,test}$ that ranges from zero to one.

Appendix B

Proactive Interference Build-Up Across Lists

Each model, other than the standard model, allows study and test items from prior lists to contribute PI to future test lists. The proportion of PI that carries over from prior lists is governed by the $\gamma_{inter-list}$ parameter, a scalar that ranges from zero to one. For any given study-test cycle other than the first, Equations A8 and A9 can be rewritten to incorporate the amount of proactive interference (in this case, within-experiment background-noise) that is preserved from a set of all prior cycles P :

$$\begin{aligned}
 \sigma_{old}^2 = & r^2(\omega^2\sigma_{tt}^2) + && \text{Self-Match} \\
 & r^2(l-1)(\omega^2\sigma_{ti}^2) + && \text{Item-Noise} \\
 & (\mu_{tt}^2\rho + \sigma_{tt}^2\rho) + && \text{Context-Noise} \\
 & \beta + && \text{Background-Noise (Pre-Experiment)} \\
 & \sum_{c \in P} \gamma_{inter-list}^L (r^2 N_c (\omega^2 \sigma_{ti}^2)) && \text{Background-Noise (Within-Experiment)}
 \end{aligned} \tag{B1}$$

$$\begin{aligned}
 \sigma_{new}^2 = & r^2 l (\omega^2 \sigma_{ti}^2) + && \text{Item-Noise} \\
 & (\mu_{tt}^2 \rho + \sigma_{tt}^2 \rho) + && \text{Context-Noise} \\
 & \beta + && \text{Background-Noise (Pre-Experiment)} \\
 & \sum_{c \in P} \gamma_{inter-list}^L (r^2 N_c (\omega^2 \sigma_{ti}^2)) && \text{Background-Noise (Within-Experiment)}
 \end{aligned} \tag{B2}$$

where c represents each prior cycle, L represents the lag between the current cycle and the prior cycle c , and N_c represents the number of non-probe items in the study list from prior cycle c . If $\gamma_{inter-list} = 1$, then all PI generated in prior lists will be preserved. If $0 < \gamma_{inter-list} < 1$, then the proportion of PI from a prior list that is preserved will decrease exponentially as lag L increases. Note that the background-noise term β is unchanged from Equations A8 and

A9, but that we refer to it as pre-experimental background-noise so as to distinguish it from the build-up of proactive interference (within-experiment background-noise).

Modelling Studied Lures

As the standard model has no mechanism to account for context change across lists, the mean and variance of the self-match for studied lures is assumed to mirror those of the self-match for targets. As the PI model incorporates the $\gamma_{inter-list}$ parameter, a measure of context change across lists, expressions for the mean and variance of the studied lure distribution can be updated.

$$\mu_{SL} = \gamma_{inter-list}(r\mu_{tt}\mu_{ss}) \quad (B3)$$

$$\begin{aligned} \sigma_{SL}^2 = & \gamma_{inter-list}(r^2(\omega^2\sigma_{tt}^2)) + & \text{Self-Match} \\ & r^2l(\omega^2\sigma_{ti}^2) + & \text{Item-Noise} \\ & (\mu_{tt}^2\rho + \sigma_{tt}^2\rho) + & \text{Context-Noise} \\ & \beta + & \text{Background-Noise (Pre-Experiment)} \\ & \sum_{c \in P} \gamma_{inter-list}^L(r^2N_c(\omega^2\sigma_{ti}^2)) & \text{Background-Noise (Within-Experiment)} \end{aligned} \quad (B4)$$

The incorporation of the $\gamma_{inter-list}$ parameter thus allows the context matches to targets and studied lures to be differentiated. Note that the self-match to studied lures is a form of proactive interference; specifically, it is a form of within-experiment context-noise as there will be a match on the item dimension but a mismatch on the context dimension (albeit a weak mismatch).

Appendix C

Free Parameters Used in Each Model

Table C1 presents the full list of free parameters used in each model.

Table C1.

List of free parameters used in each model

Model	Free Parameters
1	$t_0, s_{t0}, z_{short}, z_{long}, a_{short}, a_{long}, c_{short}, c_{long}, \gamma, z_{slope,short}, z_{slope,long}, a_{slope,short}, a_{slope,long}, \sigma_{ti}^2, \sigma_{tt}^2, \rho, r, \rho_{WF}, \mu_{ss,test}$
2	$t_0, s_{t0}, z_{short}, z_{long}, a_{short}, a_{long}, c_{short}, c_{long}, \gamma, z_{slope,short}, z_{slope,long}, a_{slope,short}, a_{slope,long}, \sigma_{ti}^2, \sigma_{tt}^2, \rho, r, \rho_{WF}, \mu_{ss,test}, \gamma_{inter-list}$
3	$t_0, s_{t0}, z_{short}, z_{long}, a_{short}, a_{long}, \gamma, z_{slope,short}, z_{slope,long}, a_{slope,short}, a_{slope,long}, \sigma_{ti}^2, \sigma_{tt}^2, \rho, r, \rho_{WF}, \mu_{ss,test}, \gamma_{inter-list}, c_{short,1} \dots c_{short,12}, c_{long,1} \dots c_{long,12}, c_{final}$
4	$t_0, s_{t0}, z_{short}, z_{long}, a_{short}, a_{long}, c_{short}, c_{long}, \gamma, z_{slope,short}, z_{slope,long}, a_{slope,short}, a_{slope,long}, \sigma_{ti}^2, \sigma_{tt}^2, \rho, \rho_{WF}, \mu_{ss,test}, \gamma_{inter-list}, r_1 \dots r_{12}, r_{final}$
5	$t_0, s_{t0}, z_{short}, z_{long}, c_{short}, c_{long}, \gamma, z_{slope,short}, z_{slope,long}, a_{slope,short}, a_{slope,long}, \sigma_{ti}^2, \sigma_{tt}^2, \rho, r, \rho_{WF}, \mu_{ss,test}, \gamma_{inter-list}, a_{short,1} \dots a_{short,12}, a_{long,1} \dots a_{long,12}, a_{final}$
6	$t_0, s_{t0}, z_{short}, z_{long}, a_{short}, a_{long}, \gamma, z_{slope,short}, z_{slope,long}, a_{slope,short}, a_{slope,long}, \sigma_{ti}^2, \sigma_{tt}^2, \rho, \rho_{WF}, \mu_{ss,test}, \gamma_{inter-list}, c_{short,1} \dots c_{short,12}, c_{long,1} \dots c_{long,12}, c_{final}, r_1 \dots r_{12}, r_{final}$
7	$t_0, s_{t0}, z_{short}, z_{long}, \gamma, z_{slope,short}, z_{slope,long}, a_{slope,short}, a_{slope,long}, \sigma_{ti}^2, \sigma_{tt}^2, \rho, r, \rho_{WF}, \mu_{ss,test}, \gamma_{inter-list}, c_{short,1} \dots c_{short,12}, c_{long,1} \dots c_{long,12}, c_{final}, a_{short,1} \dots a_{short,12}, a_{long,1} \dots a_{long,12}, a_{final}$
8	$t_0, s_{t0}, z_{short}, z_{long}, c_{short}, c_{long}, \gamma, z_{slope,short}, z_{slope,long}, a_{slope,short}, a_{slope,long}, \sigma_{ti}^2, \sigma_{tt}^2, \rho, \rho_{WF}, \mu_{ss,test}, \gamma_{inter-list}, r_1 \dots r_{12}, r_{final}, a_{short,1} \dots a_{short,12}, a_{long,1} \dots a_{long,12}, a_{final}$
9	$t_0, s_{t0}, z_{short}, z_{long}, \gamma, z_{slope,short}, z_{slope,long}, a_{slope,short}, a_{slope,long}, \sigma_{ti}^2, \sigma_{tt}^2, \rho, \rho_{WF}, \mu_{ss,test}, \gamma_{inter-list}, c_{short,1} \dots c_{short,12}, c_{long,1} \dots c_{long,12}, c_{final}, r_1 \dots r_{12}, r_{final}, a_{short,1} \dots a_{short,12}, a_{long,1} \dots a_{long,12}, a_{final}$

Notes: Model 1 is the standard model, Model 2 is the PI model, Model 7 is the winning PI model, and Model 9 is the runner-up model.