Target Learning in Event-Based Prospective Memory

Accepted, Journal of Experimental Psychology: Learning, Memory, and Cognition, 10/13/2020.

© 2020, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: [10.1037/xlm0000900](https://doi.org/10.1037/xlm0000900)

Target Learning in Event-Based Prospective Memory

Luke Strickland¹, Andrew Heathcote² Michael S. Humphreys³, &
Shayne Loft⁴

¹ The Future of Work Institute,
Curtin University, Australia

² The School of Psychology,
The University of Tasmania, Australia

³ The School of Psychology,
The University of Queensland

⁴ The School of Psychological Science,
The University of Western Australia, Australia

Address for Correspondence

Luke Strickland,
Future of Work Institute,
Curtin University,
78 Murray Street,
6000 Perth, Australia
Email: luke.strickland@curtin.edu.au

Author Note

This research was in part supported by an Australian Research Council Discovery Grant (DP160101891) awarded to Heathcote and Loft, and by Australian Research Council Discovery Grant (DP160100575) awarded to Loft. The data and code associated with the manuscript are available at: https://github.com/lukestrickland/Target_Learning. We thank Natalie Griffiths for assistance with data collection.

Abstract

Event-based prospective memory (PM) tasks require individuals to remember to perform a previously planned action when they encounter a specific event. Often, the natural environments in which PM tasks occur are embedded are constantly changing, requiring humans to adapt by learning. We examine one such adaptation by integrating PM target learning with the Prospective Memory Decision Control (PMDC) cognitive model. We apply this augmented model to an experiment that manipulated exposure to PM targets, comparing a single-target PM condition where the target was well learned from the outset, to a multiple-target PM condition with less initial PM target exposure, allowing us to examine the effect of continued target learning opportunities. Single-target PM accuracy was near ceiling whereas multiple-target PM accuracy was initially poorer but improved throughout the course of the experiment. PM response times were longer for the multiple-compared to single-target PM task but this difference also decreased over time. The model indicated that PM trial evidence accumulation rates, and the inhibition of competing responses, were initially higher for single compared to multiple PM targets, but that this difference decreased over time due to the learning of multiple-targets over the target repetitions. These outcomes provide insight into how the processes underlying event-based PM can dynamically evolve over time, and a modelling framework to further investigate the effect of learning on event-based PM decision processes.

Key Words: Prospective Memory, Prospective Memory Decision Control, Target Learning, Evidence Accumulation Models.

We often need to remember to perform a deferred action at some point in the future, for example, to take medicine after dinner, or attend a meeting. Prospective Memory (PM) refers to the cognitive processes required to perform deferred actions. Many daily activities require PM, and PM failures are relatively common (Rummel & Kvavilashvili, 2019). Furthermore, PM is often impaired in the elderly and clinical populations, interfering with critical activities such as medication adherence (Park & Kidder, 1996). Successful PM is also essential to ensure safe and efficient performance in workplace settings, such as in defence, air transportation, and healthcare (Loft Dismukes, & Grundgeiger, 2019), where PM failures can have serious consequences. Given the consequences of PM failures in clinical and safety-critical work contexts, it is essential to understand the cognitive processes that underlie prospective remembering.

Einstein and McDaniel (1990) introduced a paradigm to study PM in the laboratory. This paradigm requires participants to remember to execute their PM action while engaged in an ongoing task (e.g., a lexical decision task, deciding whether strings of letters form a word or a non-word). Often studies examine “event-based” PM, in which participants are required to make a response to PM target items if they occur in the ongoing task (e.g., press ‘9’ if presented the word ‘chair’ during lexical decision). These studies have revealed factors that benefit prospective remembering, including the degree to which attention to stimulus features required to detect a PM target overlap with the processing required for ongoing decisions (target focality; Einstein & McDaniel, 2005), and instructional emphasis on the importance of the PM task (Smith & Bayen, 2004). Recently, Strickland, Loft, Remington and Heathcote (2018) proposed Prospective Memory Decision Control (PMDC), a model that specifies the cognitive process dynamics underlying the race to response selection between PM and ongoing task goals. Unlike prior theories, PMDC can account for the full range of behavioral effects

caused by focality and importance manipulations, including ongoing task accuracy, PM accuracy, and response time (RT) distributions observed for each participant. PMDC can also account for the effects of PM load, importance, and time pressure in more applied domains such as air traffic control (Boag, Strickland, Heathcote, Neal, & Loft, 2019; Boag, Strickland, Loft, & Heathcote, 2019) and maritime surveillance (Strickland et al., 2019).

To effectively respond to the ever-changing environment, the human cognitive system must adapt. However, to date, the learning processes underlying PM have received little attention. Learning is defined as the modification of behavior as a function of task experience (Melton, 1963). Learning is a general feature of cognition (Logan, 1988; Rescorla & Wagner, 1972; Raaijmakers & Shiffrin, 1981) and the brain (e.g., Hebb, 1949), and thus is likely to affect PM. Further, differences in learning opportunities across PM and ongoing tasks are ubiquitous, with ongoing task responses required more frequently than PM responses, and these relative differences in learning may partly underlie PM errors. In line with this, Loft and Remington (2010) found that individuals were less likely to remember to make PM responses to target aircraft features that were more often practiced with the ongoing-task aircraft-acceptance response.

In the current article, we aim to investigate adaptation in PM by studying *target learning*, in which repeated presentations of a PM target lead to better performance with that target. We compare an experimental condition with a single, highly practiced, PM target to performance in a condition with multiple PM targets, where in both conditions the overall amount of time dedicated to establishing PM intentions is equated. In the former condition, target learning is expected to begin at ceiling, whereas in the latter condition learning is expected to improve performance over time. We propose a computational model in which PM and ongoing task decisions race for retrieval, and in which the number of prior learning

opportunities provided linking PM targets to the PM response affects the speed of PM response retrieval relative to the ongoing task response and, therefore, increases PM accuracy. Before proceeding to describe the current study, we introduce Strickland et al. (2018)'s PMDC model, which we then augment to account for learning.

The PMDC Model

Until recently, most PM theories relied on analysis of coarse manifest measures, particularly PM accuracy and PM cost. PM cost refers to the finding that RTs to non-PM trials are often longer in PM blocks of trials, where participants must remember to respond to PM items, as compared with control blocks, where participants do not need to remember to respond to PM targets (e.g., Marsh, Hicks, Cook, Hansen, & Pallos, 2003; Smith, 2003). PMDC allows a more comprehensive quantitative characterization of observed performance, including the entire distribution of ongoing task RTs, ongoing task response choices, as well as the RT distributions and response choices observed on PM trials.

PMDC (Figure 1) assumes a race to response selection between PM and ongoing task decision processes. Each of the ongoing task accumulators and the PM accumulator has its own threshold, which corresponds to the evidence that must be accumulated to make that decision. Upon stimulus presentation, evidence accumulates towards each decision at an accumulation rate, and the first to reach threshold determines the decision made (Brown & Heathcote, 2008). Thus, successful PM hits occur on PM trials where the PM accumulator reaches threshold before the ongoing task accumulators, whereas PM errors occur when the ongoing task accumulators reach threshold before the PM accumulator.

PMDC is a measurement model that can accurately estimate its parameters from data, and these parameters can be mapped to the cognitive mechanisms underlying PM. Two mechanisms prominent in PM theory are capacity sharing and cognitive control. The idea behind capacity sharing, borrowed from cognitive resource theories (Kahneman, 1973;

Navon & Gopher, 1979), is that PM and ongoing task processing compete for limited cognitive resources, and thus devoting resources to one task comes at the expense of the other. Cognitive control refers to the processes that allow humans to break the bounds of automatic, stimulus-driven behavior and act in a goal directed fashion (Miller & Cohen, 2001). PMDC includes two forms of cognitive control (Braver, 2012): *proactive control*, which is active in advance of an event so that the control is applied when the event occurs, and *reactive control*, that occurs when an event is encountered, in order to facilitate responding to that event.

Figure 1. The PMDC model (as depicted in Strickland et al., 2018). Each accumulator begins with a starting amount of evidence drawn from the uniform distribution $U[0, A]$. Over time, evidence accumulates towards threshold b at a rate drawn from a normal distribution with mean v , standard deviation sv . The predicted response is determined by the first accumulator to reach threshold. Total response time is equal to total time to accumulate to threshold, plus a non-decision time parameter included as time for additional processes such as stimulus encoding and motor responding.

Capacity Sharing. Several PM theories assume that PM costs result from a reduction in ongoing task capacity with the addition of a PM task (e.g., Einstein & McDaniel, 2010; Smith, 2010). The evidence accumulation rate parameters of PMDC estimated from non-PM trial responses provide a measure of ongoing task processing capacity. There are two types of ongoing task accumulation rates: *match* accumulation rates, which index the accumulation

towards the correct decision (e.g., the word accumulator's rate for a word stimulus), and *mismatch* accumulation rates, which index the accumulation towards the incorrect decision (e.g., the non-word accumulator's rate for a word stimulus). Either a decrease in match accumulation rates, or an increase in mismatch accumulation rates, could indicate a reduction in processing capacity. Contrary to previous PM theories, evidence accumulation modelling indicates that PM costs are not associated with decreased ongoing task capacity in basic PM paradigms (e.g., Heathcote, Loft & Remington, 2015; Horn & Bayen, 2015; Strickland et al., 2018), but that it is in more demanding paradigms, such as air traffic control (Boag, Strickland, Heathcote et al., 2019; Boag, Strickland, Loft et al., 2019). This indicates that capacity sharing is more likely when overall capacity demands of a task approach an individual total capacity limit, which is less likely to result from simple ongoing tasks such as lexical decision making.

Proactive control. One form of proactive control is to increase ongoing task thresholds in PM blocks of trials, so that when PM items are presented the ongoing task accumulators are less likely to pre-empt the PM accumulator, improving PM accuracy (Loft & Remington, 2013). Many studies find that participants increase ongoing task thresholds under PM conditions, and that these increased thresholds are the underlying cause of PM costs (e.g., Heathcote et al., 2015). However, simulations from PMDC (Strickland et al., 2018), and manipulations of ongoing task thresholds (Strickland, Loft & Heathcote, 2020), indicate that this form of control provides at best secondary support to PM accuracy, as compared with PMDC's other mechanisms. For example, proactive control can also be exerted over the PM threshold; Strickland et al. (2018) found that participants decreased the threshold of the PM accumulator when the importance of the PM task was emphasized, and that this substantially improved PM accuracy. In addition, PM thresholds have been shown to be larger for non-focal PM than focal PM, perhaps due to participants having conscious appreciation of higher

task demands and thus lower confidence in their ability to successfully recognize non-focal PM targets compared to focal PM targets.

Reactive control. In the PMDC model, reactive control is activated by processing of PM stimulus features on PM trials (Figure 2). Processing PM stimulus inputs activate the “PM” detector, causing participants to accrue evidence towards the PM decision (reactive excitation). In addition, PMDC also proposes that the PM detector has inhibitory connections to competing ongoing task decisions (e.g., word and non-word), and thus PM detector activation reduces accumulation rates towards these decisions. Reactive control can be assessed empirically by examining whether PM accumulation rates are faster on PM trials than non-PM trials (indicating reactive excitation), and whether ongoing task accumulation rates slower on PM trials than non-PM trials (indicating reactive inhibition). Strickland et al., (2018) found such evidence of reactive control, and that variations in reactive control were critical to accounting for effects such as that of PM target focality and PM importance.

Figure 2. PMDC’s control architecture (Strickland et al., 2018). Stimulus inputs activate the detectors, which in turn excite their corresponding accumulators, and inhibit other accumulators. The left right panel depicts the potential effects of inputs from the PM dectector – an increase in the PM accumulation rate (via A1), and a decrease in ongoing task accumulation rates (via B1 and B2).

Learning and PMDC

There is a long history of learning theories in the psychological literature. For example, in the search of associative memory theory (SAM; Raaijmakers & Shiffrin, 1980) the probability of memory retrieval is a function of the strength of associations between memory probes and images in long term memory, and each memory retrieval increases the strength of associations between targets and corresponding images. Increased associative strengths improves the probability of subsequent memory retrievals, and suppresses the retrieval of competing memory images. This process bears a resemblance to PMDC's mechanisms shown in Figure 2 in which PM detector activation increases accumulation of PM accumulation and suppresses competing ongoing task accumulators. This core idea of learning by modifying associative strengths is common to many theories (e.g., Hebb, 1949; LaBerge & Samuels, 1974; Rescorla & Wagner, 1972), and has inspired computational models of cognition such as ACT-R (J. R. Anderson & Lebiere, 1998), in which retrieval probability is determined by factors such as the frequency with which productions have been previously retrieved and the attention paid to stimuli meeting the productions conditions (also see Altmann & Trafton, 2002; Altmann & Gray, 2008; Dismukes & Nowinski, 2006; Nowinski & Dismukes, 2005).

Although many theories of learning reference associative strength, alternative cognitive mechanisms could also bring about PM target learning. For example, instance theories identify the individual experience (i.e., the instance) as the primitive unit of knowledge and treat learning as the storage and retrieval of instances from memory (e.g., Brooks, 1978; Kruschke, 1992; Logan, 1988; Nosofsky, 1986). Each experience stores an "instance" in memory and retrieval probability and speed increases with the number of stored instances. Target learning could also occur due to target repetitions invoking a feeling of familiarity, which can result from an increase in ease of stimulus processing (perceptual fluency; e.g.,

Jacoby & Dallas, 1981; Lindsay & Kelley, 1996), and which has been suggested to trigger a “discrepancy” experience (when the actual quality of stimulus processing mismatches an expected quality of processing; Whittlesea & Williams, 2001a, 2001b) that can increase the probability of the retrieval of a PM response (see McDaniel, Guynn, Einstein, & Breneiser, 2004; Lee & McDaniel, 2013). If familiarity increases through stimulus repetitions, the probability of causing a discrepancy experience also increases, which in aggregate could lead to an apparently smooth change in behaviour (see Rickard, 1997, for an analogous process in skill acquisition). These different mechanisms might also interact. For example, familiarity results from an associative learning process in the SAM model (Raaijmakers & Shiffrin, 1980), and in Logan’s (2002) instance theory of attention and memory model familiarity is the sum of the similarities of instances in memory to a presented target (Logan, 2002, see also Nosofsky, 1988).

The current work does not aim to compare the different potential accounts of PM target learning, but rather to demonstrate the presence of learning effects in PM, and to provide a model framework to describe their effects on PM decision processes. To do so, we assume a smooth increasing function that captures a negatively accelerated time course of learning towards an asymptotic level (for a review of such functions see Evans et al., 2018). Such a function is supported by most accounts of learning. This learning could result from better representations of the PM target, either due to the storage of more instances, or better connections between PM targets and PM goal representations, for example due to increased associative strength, and either of these mechanisms could increase target familiarity and/or discrepant target processing. As demonstrated by prior work in the categorization and recognition memory literatures, evidence accumulation models such as PMDC can provide a meeting point between such representational changes and the dynamic decision process, with the rate of evidence accumulation of the decision process a function of both the quality of

knowledge representations, and the strength of connections between presented cues and knowledge representations (Cox & Shiffrin, 2017; Nosofsky & Palmeri, 1997; Ratcliff, Thapar, Gomez, & McKoon, 2004).

Within the PMDC computational architecture, the key learning mechanism affecting behavior is that repeated target presentations result in greater activation of PMDC's PM detector (Figure 2). Increased activation boosts PM accumulation and reduces accumulation to competing ongoing task decisions through inhibition. It is possible that there is a direct mapping between PM targets and the PM detector, but more likely there is mediator between the representation of PM targets and the PM detector. For example, PM targets may be associated with an intermediary process such as recognition, or an arbitrary verbal label such as 'PM target', which is itself associated with the PM detector. The idea of a recognition mediator is closely linked to the outlined theories of memory. For example, a familiarity/discrepancy process brought about by PM targets could initiate recognition which in turn brings about PM responding. Although it is important to recognize the possibility of such relationships between familiarity, discrepancy and recognition processes, the current modelling is not designed to disentangle them.

The presence of a mediator implies that there are two possible types of learning – PM target to mediator, and mediator to PM detector. Learning the mapping from mediator to detector is the same (and relatively easy after some practice) for any PM target regardless of the number of targets, and so is expected to reach asymptote fairly quickly. In addition, such a mapping needs only to be learned once (e.g., at the outset of a PM experiment), rather than varying with PM conditions. Thus, the majority of observable PM target learning is expected to result either from adaptations to the direct mapping PM targets onto the PM detector, or from mapping PM targets onto an intermediary process such as recognition.

In the current study, we test PMDC's account of learning with a multi-session

experiment that manipulates exposure to PM targets, with one condition designed to induce near ceiling initial learning, and another to result in lower initial learning that enables learning to be observed with repeated PM target exposures over time. Although we report analyses of observed performance, our key analyses rely on PMDC as it can titrate the effects of shifts in the rate of evidence accumulation (i.e., on reactive excitation and inhibition) from potential shifts in decision-making strategies (i.e., proactive control), which is not possible based on observed measures such as PM costs and PM accuracy.

The Current Study

Participants performed an ongoing lexical decision task with the intention to make a PM response to certain PM target words. We included two PM conditions, representing different levels of PM target learning, but with the focality of the PM task, total time forming PM intentions, and the frequency of making PM responses equated. In one part of each experimental session participants performed a multiple-target PM task (make an alternative response to any item from a list of eight PM target words), and in another part a single-target PM task (make an alternative response to one specific target word). For each task, participants were provided the same amount of total study time to initially establish PM intentions related to the targets. Consequently, there was greater time for learning of the single PM target than of any individual target in the multiple condition. To further reinforce these differences in target learning, participants were required to successfully distinguish PM targets from non-targets prior to the main experimental blocks, in a recognition memory task format. In the single-target condition, participants were required to correctly identify the PM target sixteen times to bring learning close to asymptote. In the multiple target PM blocks, participants were required to perform two correct identifications of each target, enough to ensure adequate encoding but leaving headroom for further learning during the task.

In order to accurately estimate PM trial parameters, the experiment used a within-

subjects design, with thousands of trials performed over two days. In the single-target PM condition, each target was repeated 64 times throughout the course of each PM block. In the multiple-target PM condition, each target was repeated 8 times in each PM block. PM performance was expected to be at ceiling in the single target condition, so improvements with target repetitions was not expected. In contrast, as learning for the multiple targets does not begin at ceiling, the eight target repetitions in each block were expected to improve PM performance over the course of each block.

Given the greater initial learning in the single-target condition, we expected to observe better initial PM accuracy in the single-target condition than the multiple-target condition. In terms of PMDC, we expected to observe increased PM excitation and increased inhibition of competing ongoing task decisions in the single-target condition. In contrast, the multiple targets were expected to be learned over the eight stimulus presentations in each block, so differences in PM accuracy across single and multiple target PM blocks were predicted to decrease later in each block. We apply the PMDC model to describe this learning effect by estimating a learning function that can potentially increase PM accumulation and PM-induced inhibition as a function of the number of previous times the PM target had previously been presented in the block. We fit an exponential function, which has been shown to provide a good account of individual level learning curves in skill acquisition (Heathcote, Brown, & Mewhort, 2000; Evans et al., 2018), and incorporates the common observation that learning effects proceed towards a ceiling in a smooth, negatively accelerated manner.

PM accuracy may differ between single and multiple-target conditions because of strategic changes in response thresholds. For example, participants might consciously increase their PM thresholds relative to their ongoing task thresholds in the multiple-target condition, either due to an increase in perceived PM task difficulty or concern about elevated

PM false alarm rates. A key benefit of PMDC is that it can parse the effects of such conscious changes in decision strategies from the effects of changes in PM input detector activation.

We expect the multiple-target PM condition to result in longer ongoing task RTs than single-target PM conditions (i.e., increased PM costs), consistent with some previous studies using related paradigms (e.g., Cohen et al., 2008; Humphreys, Li, Burt & Loft, 2020). As reviewed, some PM theories claim that difficult PM tasks are associated with longer ongoing task RTs because the higher capacity requirements of the PM task drain capacity from the ongoing task (Einstein & McDaniel, 2010, Smith, 2010). For example, in the current study, participants could attempt to hold PM targets in working memory so that they can be consciously compared to presented ongoing task items. Consistent with this, working memory has been shown to play a significant role in PM (Ball, Vogel, & Brewer, 2019; Smith, 2003). However, with eight PM targets, this strategy would be highly capacity consuming, clearly exceeding the limits of working memory in the multiple-target condition, and so we think that strategy is unlikely. Participants might attempt such a strategy in the single-target condition, which would not overwhelm working memory capacity.

If participants did use a working memory strategy in the multiple-target condition to improve PM, we would expect to see lower PM accumulation rates than in the single-target condition, consistent with a set size effect (Schubert et al., 2015), but also lower accumulation rates for the ongoing task, in line with substantial capacity demands. An alternative explanation is that PM costs are characterized by increases in ongoing task thresholds, rather than decreases in capacity, as PMDC modelling has indicated is the case for other types of PM cost effects (e.g., Strickland et al, 2018). Fitting PMDC allows us to test the degree to which any PM costs we observe arise from increases in ongoing task thresholds (proactive control), costs to ongoing task accumulation (capacity sharing), or both.

Method

Participants

The study had ethics approval from the UWA Human Research Ethics Office. Participants included members of the UWA community (receiving \$40 AUD dollars as reimbursement for their time), and students participating for course credit. We tested until we reached our target number of 32 participants (11 males, the rest female) with viable data. During this process, two participants did not return for their second session, and one was unable to complete the experiment due to a power failure, and so the data of these participants was replaced. The data of three further participants was replaced. Two because they performed some blocks of trials with near chance ongoing task accuracy (<60%). The other was excluded due to 0% accuracy in discriminating their single-target PM cue from ‘new’ items in a post-block recognition test. The included participant ages ranged from 17 to 43 ($M = 22.06$, $SD = 5.41$).

Design

The experiment focused on two within-subjects conditions, multiple-target PM and single-target PM. To obtain adequate trial numbers, participants completed two sessions on different days. Each day they performed one block of each condition. If they performed the multiple-target condition first on day one, they would perform the single-target condition first on day two, and vice versa. Participants were assigned to one of four possible counterbalances for their response keys (word: ‘d’, nonword: ‘s’, PM: ‘j’, word: ‘s’, nonword: ‘d’, PM: ‘j’, word: ‘j’, nonword: ‘k’, PM: ‘d’, word: ‘k’, nonword: ‘j’, PM: ‘d’). The eight possible combinations of condition orders and key assignments were balanced over the 32 participants.

Materials

Lexical decision and PM stimuli. 1306 words and non-words were randomly

selected from Strickland et al. (2018)'s stimuli (Experiment 2). These words had length ranging between 5 and 10 characters, and were low in written frequency, ranging between 1 and 7 per million written words according to the Sydney Morning Herald Database (Dennis, 1995). Strickland et al. generated non-words with the Wuggy algorithm (Keuleers & Brysbaert, 2010), by replacing two out of three subsyllabic segments of the words, whilst matching both subsyllabic segment lengths and transition frequencies. For each participant, 18 words were randomly selected to be PM targets (2 x 8 for each multiple-target block, and 1 for each single target block). In total, each participant performed 4 blocks of 644 trials. Each block included 64 PM trials, 322 non-word trials and 258 word trials. This required assigning 1288 non-PM non-words (all of the original 1306 generated non-words except for those matching the PM items), and 1032 non-PM words to each participant (there are less non-PM words because some word trials are replaced with PM trials).

The assignment of non-PM trial stimuli to blocks, and the order of non-PM stimulus presentation order, was randomized. In the single-target condition, the PM target would be presented 64 times over the 644 trial block, and in the multiple-target condition each of the 8 PM targets would be presented 8 times over the block. PM targets were not presented in the first 2 trials of each block, or trials 323 and 324 (which followed a mid-block break). Starting after trials 2 and 322, a PM target was presented once every 10 trials. The PM target position was randomly selected from the 10-trial range, except that targets were always separated by at least 3 non-PM trials. The 8 targets in the multiple-target blocks were shuffled, and then assigned to the first 8 PM trials presented in the block, shuffled again, and assigned to the next 8 PM trials presented in the block and so on to fill out the 64 PM trial positions.

Recognition memory tests. Participants performed a recognition memory test prior to each PM block and a recognition memory test after each PM block. The recognition memory test prior to the blocks were to ensure adequate encoding of PM targets, and to provide

enough learning opportunities prior to the experimental blocks for performance to the single PM targets to reach ceiling. In the modelling of the subsequent PM task, we assume that each PM target repetition presents a learning opportunity, which might not be the case if participants failed to initially encode PM targets. Thus, the initial recognition task was looped until performed 100% accuracy to ensure adequate target encoding. Although this looping resulted in more practice recognition blocks for the multiple-target condition, potentially weakening our manipulation, our subsequent analyses reveal that PM learning began at ceiling for single-target blocks and not multiple-target blocks, indicating that our manipulation worked. The recognition test after each PM block was to further confirm that participants adequately recognized their PM targets.

The recognition tests required participants to discriminate between PM words and non-PM words. The non-PM words were drawn from Strickland et al. (2018) and were randomly selected from a pool of the 590 remaining words from that experiment that were not selected to be stimuli in the current study's primary task. The pre-block recognition memory task included 32 trials, 16 corresponding to the PM targets and 16 to non-PM words. In the single-target blocks, this corresponded to 16 presentations of the PM target. In the multiple-target blocks, this corresponded to two presentations of each of the 8 PM targets. To create the 32 trials, 8 non-PM words were selected and randomly shuffled with the 8 PM items once to form the first 16 trials, and this process was repeated to form the next 16 trials. The post-block recognition memory test consisted of 32 trials, 8 corresponding to PM targets (either one of each from the multi-target list or 8 repetitions of the single target), and 24 corresponding to non-PM words.

Procedure

Block procedure. At the start of each day, participants completed 24 practice lexical decisions with no PM component. They then received instructions for their first PM block,

which informed them that in the next block of lexical decision trials they would be required to make an alternative response to a list of target words (multiple-target condition), or to a single target word (single-target condition). On the next instruction screen, participants were presented with their PM words (or word) and asked to spend two minutes memorizing. Participants then performed a 32 trial recognition memory pre-block test. They were presented words and required to press ‘y’ to indicate the word was a PM item, or ‘n’ to indicate that the word was not. If they were not 100% accurate in the pre-block test, they were returned to study their PM word(s) again for another 2 minutes and were re-tested. This process continued until recognition accuracy reached 100%. After the recognition task, participants were instructed that they should press the PM key instead of submitting an ongoing task response (e.g., press ‘j’ instead of ‘d’) when presented with these PM item(s) during the lexical decision task. Participants then performed a three-minute distractor puzzle before commencing the block of lexical decision trials. In the middle of the 644 trial block participants were presented a break screen and asked to rest for 1 minute, to reduce possible fatigue effects. At the end of each block, participants performed a 32 trial post-block recognition memory test to assess their memory for their PM target (s).

Trial procedure. Each trial began with a fixation cross, displayed for 0.5s. This cross was followed by a blank screen for 0.25s, which remained until the stimulus (i.e., the lexical decision item or recognition memory item) was presented. The stimulus remained on screen until a key was pressed. If an incorrect response was submitted, a feedback screen was presented that for 1 second saying ‘INCORRECT!’. This screen was not presented for correct lexical decision responses to PM trials.

Results

The first two trials of each block, and the two trials after the mid-block break, were excluded from all further analyses. These trials were excluded to avoid confounding from

possible start-up costs at the beginning of each block/after each break. Any trial immediately following a PM trial or a PM false alarm was excluded from further analysis. These trials were excluded to avoid any confounding of the non-PM trial analysis caused by post-PM slowing (Meier & Rey-Mermet, 2012; Rummel et al., 2017). For each participant, trials with responses < 0.2 seconds or $>$ their mean RT + 3 times their IQR/1.349 (a robust equivalent to standard deviation) were excluded from further analysis (5.08% of responses). These exclusion criteria aim to exclude responses that are confounded by anticipatory responses (“fast guesses”), or by overall task disengagement (e.g., not attending to stimulus presentation). Before turning to the LBA model analyses, we perform a more traditional analysis of the data. We analyzed the effects of our manipulated factors: condition (multiple target PM/single target PM) and stimulus type (word/nonword). To capture within-block learning effects, we used a four-level ‘trial range’ factor in our analysis, examining trials 2-162, 162-322, 324-483, and 484-644 (note the two trials at the start of the block and after the break are excluded). In each level of this factor each of the multiple targets was presented twice and each of the single targets 16 times. PM target learning effects would be reflected in improvements in PM accuracy (increases) or PM RT (decreases) for later trial ranges as compared with earlier ranges. Finally, we also included a two-level ‘session’ factor (day 1/day 2) to capture any possible long-running effects of practice.

We applied mixed effects models with a random intercept for each participant using the ‘lme4’ package (Bates Mächler, Bolker, & Walker, 2015) implemented in R. To analyze accuracies, we fit a generalized linear model with a binomial probit link function to every observed response. To analyze RTs, we fitted a general linear model to each participant’s mean correct RTs. In text we report the main findings of these models, with tests of each effect tabulated in the supplementary materials. The supplementary materials also contain follow-up contrasts conducted with the R package lsmeans (Lenth, 2017). Our primary

interest was in identifying the strong effects that have theoretical implications and are thus crucial for our model to account for. As we applied many analyses, in some cases with a large amount of data (e.g., for the probit models fitted to every trial), we applied a conservative alpha of $p < .005$ for statistical significance (Benjamin et al., 2018). Only results that achieved this significance level are described in text. Performance on the post-block recognition test, which was used to exclude one participant with low recognition accuracy, is reported in the supplementary materials. Pre-block recognition test performance is not reported, as the pre-block task was run repeatedly until accuracy was 100%. The average number of recognition practice blocks required for perfect performance was 1.86 ($SD = 1.07$) for multiple-target conditions and 1.36 ($SD = 0.70$) for single-target conditions.

Prospective Memory Task

Accuracy

PM responses on non-PM trials were very rare (0.2% of trials) and are not analysed further. As PM stimuli were words, there was no stimulus type variable in our analyses of PM trial performance. We found a main effect of condition, and an interaction between condition and trial range (see Table 1 and supplementary materials). PM accuracy was initially higher in the single-target condition than the multiple-target condition, consistent with stronger initial learning. However, this difference was reduced for later trial ranges because accuracy in the multiple-target condition increased. This is consistent with learning of PM targets over repeated target presentations in the multiple-target condition moving towards the ceiling achieved in the single-target condition. Despite this learning effect, PM accuracy in the single-target condition was still significantly higher than the multiple-target condition even for the last trial range in each block, suggesting that learning had not reached asymptote by the end of the multiple-target blocks. In contrast to the multiple-target condition learning

effect, we found that initial single-target PM accuracy was numerically higher than the single-target PM accuracy for subsequent trial ranges, although this effect was not statistically significant.

Table 1

PM accuracies. Displaying M (SE), with SEs calculated by the Morey (2008) bias-corrected method.

Trial Range	Single	Multiple
(2,162]	0.96 (0.01)	0.81 (0.03)
(162,322]	0.93 (0.02)	0.84 (0.02)
(324,484]	0.93 (0.01)	0.88 (0.02)
(484,644]	0.93 (0.01)	0.89 (0.02)

Response Time

There was a main effect of condition, and the effect of condition interacted with trial range. Generally, PM RTs were longer in multiple-target than single-target conditions, but these differences decreased for later trial ranges (see Table 2 and supplementary materials). Analogous to the PM accuracy results, we found that PM RT for the single-target condition was relatively stable across trial ranges, whereas PM RT decreased in the multiple-target condition. This is consistent with a ceiling effect for the single-target condition, and learning in the multiple-target condition. Again, in contrast to the learning in our multiple-target condition, in the single-target condition we found a numerical advantage for PM RT (faster RT) for the first trial range as compared with other single-target trial ranges. This advantage failed to reach significance, although the difference between the first trial range and the last

trial range was very close ($p = .005$). In addition to the effects of condition and trial range, there was an effect of session, with PM RTs slightly longer on day 1 ($M = 0.71s$, $SE = 0.02s$) than on day 2 ($M = 0.69s$, $SE = 0.02s$).

Table 2

Correct PM Response Times. Displaying M (SE), with SE s calculated by the Morey (2008) bias-corrected method.

Trial Range	Single	Multiple
(2,162]	0.63s (0.01s)	0.78s (0.01s)
(162,322]	0.65s (0.01s)	0.76s (0.01s)
(324,484]	0.64s (0.01s)	0.74s (0.01s)
(484,644]	0.66s (0.01s)	0.73s (0.01s)

Lexical Decision Task

Accuracy

Ongoing task accuracies are plotted in Figure 3. By far the most substantial effect on ongoing task accuracy was that of stimulus type. Accuracy was lower to words, ($M = 0.90$, $SE = 0.01$) than to non-words ($M = 0.96$, $SE = 0.01$). There was an effect of trial range, but post-hoc comparisons of the trial ranges indicated only one difference reached significance (see supplementary materials). We did not find an effect of PM condition on ongoing task accuracy. This has been a common finding in previous PM studies, where costs tend to manifest in RT rather than accuracy (see F. T. Anderson, Strube & McDaniel, 2019).

Figure 3. Observed ongoing task accuracies displayed by condition, stimulus type, experimental session, and trial range. Error bars indicate the Morey (2008) bias-corrected, within-subjects standard errors.

Response Time

Mean correct ongoing task RTs are plotted in Figure 4. Stimulus type, condition, session, and trial range all had significant effects on mean correct ongoing task RTs. RTs were slower on day 1 ($M = 0.72s$, $SE = 0.02s$) than on day 2 ($M = 0.69s$, $SE = 0.02s$), and got slightly faster for later trial ranges (see contrasts in supplementary materials). Stimulus type interacted with condition. RTs were longer in the multiple-target condition for both word trials (single $M = 0.69s$, $SE = 0.01s$; multiple $M = 0.74s$, $SE = 0.01s$) and non-word trials (single $M = 0.69s$, $SE = 0.01s$; multiple $M = 0.71s$, $SE = 0.01s$). The differences in RTs across the PM conditions were larger for word trials than non-word trials, possibly because PM targets were exclusively words (Heathcote et al., 2015).

Figure 4. Observed ongoing task response times displayed by condition, stimulus type, experimental session, and trial range. Error bars indicate the Morey (2008) bias-corrected, within-subjects standard errors.

Model Results

We fit the three accumulator PMDC model (Figure 1). Each possible decision is assigned an accumulator, and evidence in each accumulator begins at a start point drawn from a uniform distribution $U[0, A]$. During stimulus processing evidence is accrued for each decision until one of the accumulators reaches threshold b . The first accumulator to reach threshold determines the response, and the total RT is equal to time for the accumulator to reach threshold plus non-decision time. We report thresholds in terms of B , which is equal to $b - A$. The accumulation rates are drawn from a normal distribution truncated at 0 with mean ν and standard deviation $s\nu$. The design of the experiment includes several factors that model parameters could vary over, including stimulus type (word/non-word/PM), condition (multiple-target PM/single-target PM), experimental session (day 1/day 2), and latent accumulator (one corresponding to the three possible decisions: word, non-word and PM). However, several restrictions were applied to the entire set of models reported.

To constrain parameter estimation, we only allowed one A parameter for each participant across all conditions, and one non-decision time parameter. As thresholds are

assumed to be set prior to stimulus presentation, and not change during stimulus processing, we fixed them so they could not vary over the stimulus factor that varied randomly from trial to trial. Thresholds could vary over all other factors and thus a separate threshold was estimated for each latent accumulator, for each experimental condition, and for experimental session. In order to constrain parameter estimation, thresholds were the only factor allowed to vary over experimental session in line with most previous PMDC models. We included only two standard deviations of accumulation rates in the model, one corresponding to matching accumulation (e.g., accumulation towards PM on a PM trial, accumulation towards word on a word trial), and one to mismatching accumulation (e.g., accumulation towards ‘word’ on a PM trial, accumulation towards non-word on a word trial). The standard deviation of mismatching accumulation was fixed at 1, as a scaling parameter. We estimated a separate mean ongoing task accumulation rate for each stimulus type for each latent accumulator and for each experimental condition. We estimated PM accumulation rates on PM trials separately for each experimental condition. We only estimated one ‘PM false alarm’ accumulation rate (i.e., the rate of PM accumulation on non-PM trials) across the two conditions, because PM false alarms were rarely observed. We included a ‘reactive inhibition’ parameter in our model, which is equal to the difference between ongoing task accumulation rates on non-PM trials and on PM trials, with higher inhibition indicating that the ongoing task accumulation rate was reduced on PM trials. Reactive inhibition was estimated separately for each condition and for each latent accumulator (word/non-word), and for each condition (multiple/single).

In order to model target learning, we allowed PM accumulation and PM-induced inhibition to vary as a function of how many times the PM target had previously been presented. We modelled learning with an exponential function, where the PM accumulation rate for PM target repetition P_N is given by $P_N = asymptote_{P_N} - De^{-N\alpha}$. Similarly,

inhibition is given by $I_N = asymptote_{I_N} - De^{-N\alpha}$. The asymptotes represent the maximum PM accumulation rate and PM inhibition of ongoing task accumulation once learning is complete. These asymptotes did not vary across single and multiple-target conditions. In addition, one learning rate parameter, α , was estimated for each participant. These constraints amount to the assumption that both conditions follow the same learning process, with the difference in conditions being the initial amount of learning. The D parameter represents the difference between the PM accumulation rates and PM inhibition at first target presentation and their asymptotes. The D parameter was the same for PM accumulation and PM induced inhibition, instantiating the assumption that learning equally affects both by activation of the PM detector (Figure 2). The D parameter varied across the single-target and multiple-target conditions to allow for different initial levels of learning of single vs. multiple PM targets. As this is the first paper to implement this PM learning model, we evaluated its parameter recovery properties in detail (see supplementary materials).

Sampling

We estimated model parameters with Bayesian estimation, using the Dynamic Models of Choice suite of R functions (Heathcote et al., 2019). We estimated separate parameters for each individual participant. We could have fit a hierarchical model that assumes a common hyper-parameter distribution across participants, and which uses that distribution to constrain individual-level estimates (“shrinkage”). However, as with our previous applications of PMDC, we found that fitting such a model was prohibitively computationally expensive. Further, we did not have knowledge of the appropriate population level distributions for PM learning parameters. Fitting a hierarchical model with inappropriately specified population distributions could result in inappropriate shrinkage of individual-level parameter estimates. Bayesian estimation requires specifying prior distributions, which detail beliefs about the

parameter values prior to observing the data. Our chosen prior distributions are displayed in Table 1. Generally, they are similar to the priors selected in Strickland et al., (2018).

Table 3

Priors for the parameters of the PMDC model fitted to our data. All prior distributions were either normal (when no lower or upper bounds were specified), or truncated normal (when lower and upper bounds were specified). The v (PM match) and reactive inhibition priors refer to the asymptote values in the learning equation.

Model Parameter	M	SD	Lower	Upper
A	1	1	0	10
B	2	1	0	None
v (Ongoing match)	1	2	None	None
v (Ongoing mismatch)	0	2	None	None
v (PM match)	1	2	None	None
v (PM false alarm)	-1	2	None	None
Reactive inhibition	0	2	None	None
α	0	2	0	None
D	0	2	None	None
s_v	1	1	0	None
$t0$	0.3	1	0.1	1

DMC's sampling algorithm requires running many parallel Markov chains, which share information to efficiently converge to the posterior (Turner, Sederberg, Brown, & Steyvers, 2013). For each participant, we ran 90 parallel chains, which was three times the number of total free model parameters (30). We sampled 3600 iterations for each chain at a time, retaining 180 of them after thinning. We continued to run iterations until the posterior samples appeared stationary, mixed, and converged, which we confirmed with visual

inspection and Gelman's R statistic (Gelman et al., 2013).

Model Fit

Figure 5 and 6 display the posterior predictions of the model. Overall, the model provided a close fit to the observed trends in accuracy and RT. The plots reveal some degree of miss-fit to errors on PM trials. However, as PM trials are quite rare, and PM accuracy quite high, this represents a very small amount of observed data. Figure 7 demonstrates that the model is able to fairly accurately capture trends in PM accuracy and RT over repeated presentations of the PM target, except for some underestimation of single-target PM accuracy for the first eight PM targets. This suggests an adequate learning model.

Figure 5. Posterior predictions for accuracies, averaged over participants. The model predictions correspond to the white circles, the posterior means correspond to the black shaded dots. The error bars display the 95% posterior credible intervals of the predictions.

Figure 6. Posterior predictions for response time (RT), pooled over participants. The model predictions correspond to the white circles, the posterior means correspond to the black shaded dots. The error bars display the 95% posterior credible intervals of the predictions. Three quantiles of RT are depicted. The bottom quantiles on each plot represent the 0.1 quantile, the middle the median RT, and the top the 0.9 quantile of RT.

Figure 7. Posterior predictions for effects of multiple PM target presentations on PM accuracy and correct PM response time (RT). The model predictions correspond to the white circles, the posterior means correspond to the black shaded dots. The error bars display the 95% posterior credible intervals of the predictions. In the RT graphs, three quantiles of RT are depicted. The bottom quantiles on each plot represent the 0.1 quantile, the middle the median RT, and the top the 0.9 quantile of RT. PM Trial Position refers to the position of the PM trial within the experimental block. Within the multiple target PM blocks, each PM trial position contained exactly 1 presentation of each possible PM target. In the single target blocks, only one PM target was presented, and thus each trial position included eight repetitions of the same PM target.

Parameter Estimates

In order to examine parameter estimates across experimental conditions, we created a ‘subject-averaged’ distribution of parameter estimates, which averaged the values of each parameter over all subjects for each posterior sample. The posterior mean of the A parameter was 0.2 ($SD = 0.01$), the mean of the non-decision time parameter was 0.12 ($SD = 0.002$), and the posterior mean of the sv parameter towards decisions matching the correct response was 0.43 ($SD = 0.005$). In the subsequent sections, we review the patterns in other parameters

across experimental conditions, particularly focusing on how they vary across the multiple target and single target PM conditions. To statistically test parameter differences, we use a posterior p value based on the number of times that one parameter was sampled higher in than the other. We report the posterior p in the direction against observed effects, to be consistent with intuition about p values. Thus, if we observed parameter x was mostly larger than parameter y , we would report posterior p as the proportion of samples on which y was larger than x . In the supplementary materials, we explore how the parameter differences reported in text related to the observed performance data with simulations.

PM Learning

As reviewed, learning theory dictates that both PM conditions shared a common asymptote for the overall rate of PM accumulation on PM trials ($M = 2.33$, $SD = 0.06$), the inhibition of word accumulation on PM trials ($M = 1.65$, $SD = 0.05$) and the inhibition of non-word accumulation on PM trials ($M = 2.52$, $SD = 0.14$). In addition, both conditions shared a common learning rate ($M = 0.62$, $SD = 0.09$). Conditions were allowed to vary in terms of D , which represents the difference between the initial accumulation and inhibition rates and the asymptote values. We found that the D parameter was above 0 for the multiple-target condition ($M = 0.21$, $SD = 0.05$), indicating learning, $Z = 4.24$, $p < .001$. However, the D parameter was below 0 for the single-target condition, indicating that PM accumulation and PM-induced inhibition actually decreased with PM target repetitions ($M = -0.29$, $SD = 0.06$), $Z = -4.98$, $p < .001$.

Proactive Control

The obtained threshold estimates are depicted in Figure 8. We expected participants in the multiple-target PM condition might increase their PM thresholds, due to increased

perceived task difficulty or awareness of the weaker learning of targets in the multiple-target condition. Consistent with this, there were very large differences between the PM threshold in the multiple-target condition (Day 1 $M = 1.30$, $SD = 0.02$; Day 2 $M = 1.26$, $SD = 0.02$) and the PM threshold in the single target PM condition (Day 1 $M = 1.18$, $SD = 0.02$; Day 2 $M = 1.14$, $SD = 0.02$), $Z = 11.57$, $p < .001$.

Figure 8. Estimates of thresholds. The shapes indicate the posterior means and the error bars correspond to the mean plus or minus the posterior standard deviation.

In addition to increasing PM thresholds, participants also increased ongoing task thresholds in multiple-target blocks. Thresholds towards making a word response were higher in the multiple-target condition (day 1 $M = 1.03$, $SD = 0.02$; day 2 $M = 0.99$, $SD = 0.02$), than in the single-target condition, (day 1 $M = 0.97$, $SD = 0.02$; day 2 $M = 0.92$, $SD = 0.01$) $Z = 8.47$, $p < .001$). Non-word thresholds were also higher in the multiple target condition (day 1 $M = 0.95$, $SD = 0.02$; day 2 $M = 0.90$, $SD = 0.01$) than the single target condition (day 1 $M = 0.94$, $SD = 0.01$; day 2 $M = 0.85$, $SD = 0.01$), $Z = 4.78$, $p < .001$, although this difference was generally smaller and not substantial on day 1. The finding that non-word threshold increases were smaller than word thresholds is consistent with Heathcote et al. (2015)'s previous modelling of a PM task that only included word PM targets.

Capacity Sharing

Estimates for accumulation rates are depicted in Figure 9. Previous modelling of simple paradigms such as lexical decision have found no capacity effects associated with PM (e.g., Strickland et al., 2018), and thus we did not expect to find reduced ongoing task capacity in the multiple-target condition as compared with the single-target condition. In line with this, correct accumulation rates to word trials were not substantially lower in the multiple-target condition ($M = 1.88$, $SD = 0.03$) compared with the single-target condition, ($M = 1.9$, $SD = 0.03$), $Z = 1.6$, $p = .057$. There were also no substantial differences between the correct non-word accumulation rates across multiple ($M = 1.86$, $SD = 0.02$) and single ($M = 1.87$, $SD = 0.02$) target conditions, $Z = 0.82$, $p = .20$. These results suggest no appreciable loss of ongoing task capacity across multiple and single-target conditions.

Figure 9. Estimates of accumulation rates. The circles indicate the posterior means and the error bars correspond to the mean plus or minus the posterior standard deviation.

We found mixed results regarding the error accumulation rates. For word trials, error accumulation rates were lower in the multiple-target condition ($M = -0.27$, $SD = 0.05$) than the single-target condition ($M = -0.13$, $SD = 0.04$), $Z = 3.95$, $p < .001$, consistent with a gain in capacity in the multiple-target condition. In contrast, for non-word trials error accumulation rates were higher in the multiple condition ($M = -0.73$, $SD = 0.06$) than the single-target condition ($M = -0.86$, $SD = 0.06$), $Z = 2.43$, $p = .009$, consistent with a loss of capacity. Overall, given the error accumulation rates indicated one in favor of more capacity for the multiple-target condition, and one indicative of reduced capacity, they did not provide robust evidence for capacity sharing with multiple target PM.

Discussion

We manipulated the learning of PM targets in a lexical decision PM paradigm. We compared a condition in which participants were required to make a PM response to an overlearned single PM target with a multiple-target condition in which participants were required to respond to any of a list of eight target words. We found that PM accuracy was higher, and PM RT was faster, for the single-target condition. We also measured how PM performance evolved over the course of each experimental block. In the single-target condition, PM performance did not improve later in the block, consistent with target learning beginning at ceiling. In contrast, we found that, later in each block, PM accuracy increased, and PM RT decreased in the multiple-target condition, getting closer to comparability with the single-target PM condition by the end of each block, consistent with target learning proceeding over the course of each block. In addition to these effects on PM performance, we also found slower ongoing task RT in multiple-target conditions compared with single-target conditions.

We found that the PMDC model provided good fits to the observed data. The model indicated that PM excitation and PM-induced inhibition of ongoing task accumulation on PM

trials were initially larger in the single-target PM condition than in the multiple-target condition, consistent with greater activation of the PM detector in the single-target condition (see Figure 2). This is consistent with increased initial learning of the single PM target due to more extensive rehearsal for the single PM target than the multiple PM targets, as well as the provision of more learning opportunities in the pre-block recognition task. In addition, our model measured learning over the course of the PM blocks, and we found evidence that learning over repeated target presentations increased PM accumulation rates and PM-induced inhibition of ongoing task responses in multiple-target conditions. This suggests that learning of PM targets can occur over the time course of an ongoing task, reflecting adaptation to a dynamic task environment. Such learning effects are compatible with multiple psychological explanations. This includes associative learning, in which target repetitions strengthen associations between PM targets and goal representations (e.g., Raaijmakers & Shiffrin, 1980), and instance learning, in which target repetitions are automatically stored as instances in long term memory, with more instances increasing the probability of memory retrieval (Logan, 1988), either of which could have increased target familiarity and the discrepant processing of targets (e.g., McDaniel et al., 2004). The contribution of the current study is not to distinguish between these explanations, but to demonstrate the presence of learning effects in PM, and to provide a model framework to describe their effects on PM decision processes.

We did not find evidence of target learning in the single-target condition, consistent with a ceiling effect. In fact, we found a (non-significant) trend for a decrease in PM accuracy and slowing of PM RT after the first 8 target presentations, both of which run in the opposite direction to a learning effect. Our PMDC model indicated an effect opposed to learning, with decreased PM excitation and PM-induced inhibition for later stimulus presentations. We consider three possible explanations for these findings. One possibility is that fatigue over blocks caused general degradations in performance. However, ongoing task performance

shows little evidence of fatigue over the course of the single-target blocks. Second, although we have focused on target learning, other types of learning could also be important, such as learning associations between responses and experimental context. That is, over the course of the experimental blocks, the association between ongoing task responses and the experimental task context may have grown stronger, increasing ongoing task retrieval speed, and therefore decreasing the probability of PM retrieval (Loft & Remington, 2010, 2013). However, this account predicts a speedup in PM RT over practice, as slow PM decisions should be unlikely to reach threshold before fast ongoing task decisions, but the opposite was observed. Thirdly, some participants may have rehearsed and maintained their single PM target in working memory over the course of the filler tasks, and hence began the experimental blocks with the PM item in an activated and readily-accessible state (McElree, 2006; Oberauer, 2002; Öztekin, Davachi, & McElree, 2010). It is possible that participants then learnt that they could achieve reasonable performance by instead relying on recognition of the single PM target, rather than maintaining the target in working memory, and subsequently abandoned the more demanding, but slightly better performing, working memory strategy.

Given the above suggests that participants may sometimes choose to hold the single PM target in working memory, one could ask whether they also attempt to do so in the multiple-target condition. Given the large capacity demands of holding eight items in memory while performing an ongoing task, such a strategy would be expected to lead to large capacity costs in the multiple-target condition as compared with the single-target condition. However, the current modelling indicated that differences in ongoing task performance across multiple-target and single-target PM conditions were accounted for by shifts in ongoing task thresholds across the conditions, and not shifts in ongoing task accumulation rates (capacity-sharing). This finding suggests that participants did not attempt to hold PM items in working

memory in the multiple-target condition. It is also consistent with our previous findings in lexical decision PM paradigms, where the effects of PM task focality and importance on ongoing task performance were explained by threshold increases (Strickland et al., 2018). Participants might have increased their ongoing task thresholds in multiple-target conditions to allow more time for the slower PM accumulator reach threshold on PM trials, as described by the delay theory of PM cost (Heathcote et al., 2015). However, recent simulations and empirical work suggest this mechanism is relatively ineffective in promoting PM accuracy (Strickland et al., 2020). Thus, thresholds might simply have been raised as a generic response to an increase in perceived task difficulty.

We found that PM thresholds were lower in the single-target condition than in the multiple-target condition. Participants may have raised thresholds in the multiple-target condition due to a perception of higher task difficulty, or in response to concern about false alarms due to poorer discrimination of the multiple PM targets. Interestingly, Strickland et al. (2018) found a similar pattern, where PM thresholds to a single-target focal task were lower than thresholds to a categorical PM task. The underlying mechanism between such PM threshold elevations warrants further investigation.

Our paradigm closely resembled a PM “list length” paradigm, in which PM tasks with longer lists of target words are compared with tasks to shorter lists (e.g., make a PM response to a single target word). Previous findings regarding PM list length have been mixed. Some studies find that longer PM target lists diminish PM accuracy (F. T. Anderson et al., 2019; Cohen, Jaudas, Hirschhorn, Sobin, & Gollwitzer, 2012), others that longer lists of PM targets do not diminish PM accuracy, but instead increase PM costs (e.g., Cohen et al., 2008; Einstein et al., 2005), and others still have reported decreased PM accuracy and increased PM costs (Humphreys et al., 2020). Accordingly, the effects of PM list length have been described as a “puzzle” for PM theories (F. T. Anderson et al. 2019). Previous findings may

have been mixed and theoretical inferences clouded due the types of learning effects that we have identified, with more learning opportunities for smaller target lists in some designs, either directly during testing or indirectly during rehearsal. To our knowledge, the learning of target lists cannot be readily balanced with other factors such as the spacing between target presentations, the level of boredom and fatigue, and retention interval (Bowyer, Humphreys, & Revelle, 1983; Hintzman, 1974). This is why our approach was to push learning to asymptote in our single-target condition and observe learning towards that asymptote in the multiple-target condition.

In our study, we controlled for PM target focality, with both our single and multiple target conditions including a PM task that is focal to lexical decision making (identify specific target words; see the focal/non-focal condition coding in F.T Anderson et al., 2019). Nevertheless, our findings may have implications for studies on PM target focality. When focality is manipulated, the focal task is most often a single-target word that is repeated many times over the course of the block (e.g., Einstein et al., 2005), as in our single-target condition, whereas non-focal conditions typically involve no or fewer PM target repetitions. As focal PM tasks generally do not involve training recognition memory to ceiling the way that we did, it is also possible there is learning over the course of PM blocks as single PM targets are repeatedly presented. In line with this, Hicks Franks, and Spitler (2017) found that performance in a categorical non-focal task (respond to any word that is an animal) became comparable to performance in a focal task when the categorical non-focal task repeated only a single target word, suggesting that learning over stimulus repetitions may underlie focality effects to some degree. At this stage, it is not possible to disentangle whether the participants in Hicks et al. learned to approach their single-target categorical task as if they were performing a focal task (i.e., conscious, explicit learning), or they benefitted from implicit learning of the PM targets in a slower, graded fashion.

A series of previous studies applying PMDC modeling have highlighted the role of cognitive-control processes in determining PM performance (Boag, Strickland, Heathcote et al., 2019; Boag, Strickland, Loft et al., 2019; Strickland et al., 2018, 2019, 2020). The present work shows that learning can decrease PM errors through modulation of the inputs to the control processes and presents a framework for modelling such learning. There are several possible ways this learning framework could be extended. One way would be to add a decay function, to account for reductions over time in the effects of PM target presentations (e.g., Anderson, Fincham, & Douglass, 1999). Investigating such a model would likely require manipulation of the intervals among target presentations. Future research could also use the augmented PMDC model to account for the positive effects of PM context re-instatement (Smith, & Skinner, 2019), or negative effects of proactive interference (Oates & Peynircioğlu, 2016), on prospective remembering. For example, regarding proactive interference, Cook, Marsh, Hicks, and Martin (2006) found that PM targets that were associated with items in a previous experimental task reduced their effectiveness as cues in a subsequent PM task.

Another possible direction relates to discrete shifts in strategy. For example, we speculated that participants might initially favor a conscious strategy for PM performance in our single-target condition in which they attempt to hold items in working memory, but at some point transition to a strategy that is more reliant on automatic, recognition memory processes. As another example, participants may modify their initially-set decision thresholds to speed up responding, and risk the increased possibility of missing PM targets, if they learn over time that targets are infrequent (Loft, Kearney & Remington, 2008). The potential for such strategic shifts is appreciated by the Dynamic Multi-process View (Scullin et al., 2013). This change point idea could be incorporated into PMDC by specifying a model where parameters change at some specific point, with the precise location of that point estimated

from the data. A final future direction relates to the type of PM errors modelled. In some instances, PM errors might occur due to a failure to store a PM target in the first place or perhaps through a failure to retrieve PM intent resulting in no activation occurring in the PM detector, and it would be important to identify how such errors are affected by learning. In the current design, we avoided these possibilities by training participants on PM targets until they could adequately recognize them even in the multiple-target condition. However, future experimental designs could attempt to induce such failures, which might be modeled in PDMC by allowing a probability that the PM accumulator fails to run on a proportion of trials (see Matzke et al., 2017, for a related approach).

References

- Altmann, E. M., & Gray, W. D. (2008). An integrated model of cognitive control in task switching. *Psychological Review*, 115(3), 602–639.
- Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, 26(1), 39–83.
- Anderson, F. T., Strube, M. J., & McDaniel, M. A. (2019). Toward a better understanding of costs in prospective memory: A meta-analytic review. *Psychological Bulletin*, 145(11), 1053–1081.
- Anderson, J. R., Fincham, J. M., & Douglass, S. (1999). Practice and retention: A unifying analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(5), 1120-1136.
- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. (pp. xi, 490). Lawrence Erlbaum Associates Publishers.
- Ball, B.H., Vogel, A., Brewer, G.A. (2019). Individual Differences in Prospective Memory. In J Rummel & M. A. McDaniel (Eds.), *Current Issues in Memory: Prospective Memory* (pp. 116-134). Routledge.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software; Vol 1, Issue 1 (2015)*.
<https://www.jstatsoft.org/v067/i01>
- Benjamin, D. J., Berger, J. O., Johannesson, M., Nosek, B. A., Wagenmakers, E.-J., Berk, R., Bollen, K. A., Brembs, B., Brown, L., & Camerer, C. (2018). Redefine statistical significance. *Nature Human Behaviour*, 2(1), 6–10.
- Boag, R., Strickland, L., Heathcote, A., Neal, A., & Loft, S. (2019). Cognitive control and capacity for prospective memory in complex dynamic environments. *Journal of Experimental Psychology: General*, 148(12), 2181–2206.

- Boag, R., Strickland, L., Loft, S., & Heathcote, A. (2019). Strategic attention and decision control support prospective memory in a complex dual-task environment. *Cognition*, *191*, 103974.
- Bowyer, P. A., Humphreys, M. S., & Revelle, W. (1983). Arousal and recognition memory: The effects of impulsivity, caffeine and time on task. *Personality and Individual Differences*, *4*(1), 41–49.
- Braver, T. S. (2012). The variable nature of cognitive control: A dual mechanisms framework. *Trends in Cognitive Sciences*, *16*(2), 106–113.
- Brooks, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169–211). Hillsdale: Lawrence Erlbaum Associations, Inc.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178.
<https://doi.org/10.1016/j.cogpsych.2007.12.002>
- Cohen, A. L., Jaudas, A., & Gollwitzer, P. M. (2008). Number of cues influences the cost of remembering to remember. *Memory & Cognition*, *36*, 149–156.
<https://doi.org/0.3758/MC.36.1.149>
- Cohen, A. L., Jaudas, A., Hirschhorn, E., Sobin, Y., & Gollwitzer, P. M. (2012). The specificity of prospective memory costs. *Memory*, *20*(8), 848–864.
<https://doi.org/10.1080/09658211.2012.710637>
- Cook, G., Marsh, R., Hicks, J., & Martin, B. (2006). Fan effects in event- based prospective memory. *Memory*, *14*, 890–900.
- Cox, G. E., & Shiffrin, R. M. (2017). A dynamic approach to recognition memory. *Psychological Review*, *124*(6), 795–860.
- Dennis, S. (1995). *The Sydney Morning Herald word database*.

- Einstein, G. O., & McDaniel, M. A. (1990). Normal aging and prospective memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16(4), 717–726.
<https://doi.org/10.1037/0278-7393.16.4.717>
- Einstein, G. O., & McDaniel, M. A. (2005). Prospective memory multiple retrieval processes. *Current Directions in Psychological Science*, 14(6), 286–290.
<https://doi.org/10.1111/j.0963-7214.2005.00382.x>
- Einstein, G. O., & McDaniel, M. A. (2010). Prospective memory and what costs do not reveal about retrieval processes: A commentary on Smith, Hunt, McVay, and McConnell (2007). *Journal of Experimental Psychology-Learning Memory and Cognition*, 36(4), 1082–1088. <https://doi.org/10.1037/a0019184>
- Evans, N. J., Brown, S. D., Mewhort, D. J., & Heathcote, A. (2018). Refining the law of practice. *Psychological Review*, 125(4), 592.
- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., & Rubin, D. (2013). *Bayesian Data Analysis* (3rd ed.). Taylor & Francis.
- hae Lee, J., & McDaniel, M. A. (2013). Discrepancy-plus-search processes in prospective memory retrieval. *Memory & cognition*, 41(3), 443-451.
- Heathcote, A., Brown, S.D. & Mewhort, D.J.K. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin and Review*, 7, 185-207.
- Heathcote, A., Lin, Y.-S., Reynolds, A., Strickland, L., Gretton, M., & Matzke, D. (2019). Dynamic models of choice. *Behavior Research Methods*, 51(2), 961–985.
- Heathcote, A., Loft, S., & Remington, R. W. (2015). Slow down and remember to remember! A delay theory of prospective memory costs. *Psychological Review*, 122(2), 376–410.
<https://doi.org/10.1037/a0038952>
- Hebb, D. O. (1949). *The organization of behavior: a neuropsychological theory*. Wiley.
- Hicks, J. L., Franks, B. A., & Spitler, S. N. (2017). Prior task experience and comparable

- stimulus exposure nullify focal and nonfocal prospective memory retrieval differences. *The Quarterly Journal of Experimental Psychology*, 70(10), 1997–2006.
- Hintzman, D. L. (1974). Theoretical implications of the spacing effect. In R. L. Solso (Ed.), *Theories in cognitive psychology: The Loyola Symposium*. Lawrence Erlbaum.
- Horn, S. S., & Bayen, U. J. (2015). Modeling criterion shifts and target checking in prospective memory monitoring. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(1), 95–117. <https://doi.org/10.1037/a0037676>
- Humphreys, M. S., Li, Y.R., Burt, J. S., Loft, S. (2020). How semantic processing affects recognition memory. *Journal of Memory and Language*, 113, 1-10.
- Kahneman, D. (1973). *Attention and effort*. Prentice-Hall, Inc.
- Keuleers, E., & Brysbaert, M. (2010). Wuggy: A multilingual pseudoword generator. *Behavior Research Methods*, 42(3), 627–633.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of learning. *Psychological Review*, 99, 22–44
- LaBerge, D., & Samuels, S. J. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology*, 6(2), 293–323.
- Lindsay, S. D., & Kelley, C. M. (1996). Creating illusions of familiarity in a cued recall–remember/know paradigm. *Journal of Memory and Language*, 35, 197–211.
- Loft, S., Dismukes, R. K., & Grundgeiger, T. (2019). Prospective memory in safety-critical work contexts. In J Rummel & M. A. McDaniel (Eds.), *Current Issues in Memory: Prospective Memory* (pp. 170–185). Routledge.
- Loft, S., Kearney, R., Remington, R. (2008). Is task interference in event-based prospective memory dependent on cue presentation? *Memory and Cognition*, 36, 139-148.
- Loft, S., & Remington, R. W. (2010). Prospective memory and task interference in a continuous monitoring dynamic display task. *Journal of Experimental Psychology:*

Applied, 16, 145–157.

Loft, S., & Remington, R. W. (2013). Wait a second: Brief delays in responding reduce focality effects in event-based prospective memory. *The Quarterly Journal of Experimental Psychology*, 66(7), 1432–1447.

<https://doi.org/10.1080/17470218.2012.750677>

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological review*, 95(4), 492-527.

Logan, G. D. (2002). An instance theory of attention and memory. *Psychological review*, 109(2), 376-400

Marsh, R. L., Hicks, J. L., Cook, G. I., Hansen, J. S., & Pallos, A. L. (2003). Interference to ongoing activities covaries with the characteristics of an event-based intention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 861–870. <https://doi.org/10.1037/0278-7393.29.5.861>

Matzke, D., Hughes, M., Badcock, J. C., Michie, P., & Heathcote, A. (2017). Failures of cognitive control or attention? The case of stop-signal deficits in schizophrenia. *Attention, Perception, & Psychophysics*, 79(4), 1078–1086.

McDaniel, M. A., Guynn, M. J., Einstein, G. O., & Breneiser, J. (2004). Cue-focused and reflexive-associative processes in prospective memory retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(3), 605-614.

McElree, B. (2001). Working memory and focal attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(3), 817–835.

McElree, B. (2006). Accessing recent events. In B.H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 46, pp. 155–200). San Diego, CA: Academic Press

Meier, B., & Rey-Mermet, A. (2012). Beyond monitoring: After-effects of responding to prospective memory targets. *Consciousness and Cognition*, 21(4), 1644–1653.

- Melton, A. W. (1963). Implications of short-term memory for a general theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 2(1), 1–21.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24(1), 167–202.
- Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychological Review*, 86(3), 214–255. <https://doi.org/10.1037/0033-295x.86.3.214>
- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115, 39–57.
- Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 700–708.
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104(2), 266–300.
- Oates, J. M., & Peynircioğlu, Z. F. (2016). Event-Based Prospective Memory Is Resistant but Not Immune to Proactive Interference. *American Journal of Psychology*, 129(4), 419–427.
- Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 411–421.
- Oberauer, K. (2009). Design for a working memory. *Psychology of Learning and Motivation*, 51, 45–100.
- Öztekin, I., Davachi, L., & McElree, B. (2010). Are representations in working memory distinct from representations in long-term memory? Neural evidence in support of a single store. *Psychological Science*, 21, 1123–1133.
- Park, D. C., & Kidder, D. P. (1996). Prospective memory and medication adherence. In M.

- Brandimonte, G. O. Einstein, & M. A. McDaniel (Eds.), *Prospective memory: Theory and applications* (pp. 369–390). Lawrence Erlbaum Associates.
- Raaijmakers, J. G., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological review*, 88(2), 93-134.
- Ratcliff, R., Thapar, A., Gomez, P., & McKoon, G. (2004). A diffusion model analysis of the effects of aging in the lexical-decision task. *Psychology and Aging*, 19(2), 278–289.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black, & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton Crofts.
- Rickard, T. C. (1997). Bending the power law: A CMPL theory of strategy shifts and the automatization of cognitive skills. *Journal of Experimental Psychology: General*, 126(3), 288-311.
- Rummel, J., & Kvavilashvili, L. (2019). Take the Field! Investigating Prospective Memory in Naturalistic and Real-Life Settings. In J Rummel & M. A. McDaniel (Eds.), *Current Issues in Memory: Prospective Memory* (pp. 157–169). Routledge.
- Rummel, Jan, Wesslein, A.-K., & Meiser, T. (2017). The role of action coordination for prospective memory: Task-interruption demands affect intention realization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(5), 717-735.
- Schubert, A. L., Hagemann, D., Voss, A., Schankin, A., & Bergmann, K. (2015). Decomposing the relationship between mental speed and mental abilities. *Intelligence*, 51, 28-46.
- Scullin, M. K., McDaniel, M. A., & Shelton, J. T. (2013). The Dynamic Multiprocess Framework: Evidence from prospective memory with contextual variability. *Cognitive Psychology*, 67(1), 55–71.

- Smith, R. E. (2003). The cost of remembering to remember in event-based prospective memory: Investigating the capacity demands of delayed intention performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 347–361. <https://doi.org/10.1037/0278-7393.29.3.347>
- Smith, R. E. (2010). What costs do reveal and moving beyond the cost debate: Reply to Einstein and McDaniel (2010). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36, 1089–1095. <https://doi.org/10.1037/a0019183>
- Smith, R. E., & Bayen, U. J. (2004). A Multinomial Model of Event-Based Prospective Memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 756–777.
- Smith, R. E. & Skinner, D. (2019). Prospective memory in context: Methods, findings, and future directions. In J. Rummel & M. McDaniel (Eds.), *Current issues in memory: Prospective memory* (pp 27-43). Routledge.
- Strickland, L., Elliott, D., Wilson, M. D., Loft, S., Neal, A., & Heathcote, A. (2019). Prospective memory in the red zone: Cognitive control and capacity sharing in a complex, multi-stimulus task. *Journal of Experimental Psychology: Applied*.
- Strickland, L., Loft, S., & Heathcote, A. (2020). Investigating the effects of ongoing task bias on prospective memory. *Quarterly Journal of Experimental Psychology*.
- Strickland, L., Loft, S., Remington, R. W., & Heathcote, A. (2018). Racing to remember: A theory of decision control in event-based prospective memory. *Psychological Review*, 125(6), 851–887.
- 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.
- Whittlesea, B. W., & Williams, L. D. (2001a). The discrepancy-attribution hypothesis: I. The

heuristic basis of feelings and familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(1), 3-13.

Whittlesea, B. W., & Williams, L. D. (2001b). The discrepancy-attribution hypothesis: II. Expectation, uncertainty, surprise, and feelings of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(1), 14-33.