Using response time distributions and race models to characterize primacy and recency effects in free recall initiation

Adam F. Osth

University of Melbourne

Simon Farrell

University of Western Australia

Address correspondence to:

Adam Osth (E-mail: adamosth@gmail.com)

Author Note

We would like to give special thanks to Mike Kahana and colleagues for generously sharing these datasets and Karl Healey for assistance with extracting response times from the PEERS datasets. This work was supported by ARC DE170100106 awarded to Adam Osth, and an ARC Future Fellowship (FT130100149) awarded to Farrell. Model code can be found at <https://osf.io/bkjqn/>

Abstract

Memory models have typically characterized retrieval in free recall as multi-alternative decision making. However, the majority of these models have only been applied to mean response times (RTs), and have not accounted for the complete RT distributions. We show that RT distributions carry diagnostic information about how items enter into competition for recall, and how that competition impacts on the dynamics of recall. We jointly fit RT distributions and serial position functions of free recall initiation with both a racing diffusion model and the linear ballistic accumulator (LBA: Brown & Heathcote, 2008) model in a hierarchical Bayesian framework while factorially varying different assumptions of how primacy and recency are generated. Recency was either a power law or an exponential function. Primacy was treated either as a strength boost to the early list items so that both primacy and recency items jointly compete to be retrieved, a rehearsal process whereby the first item is sometimes rehearsed to the end of the list to make it functionally recent, or due to reinstatement of the start of the list. While serial position curves do not distinguish between these accounts, they make distinct predictions about how RT distributions vary across serial positions. Results from a number of datasets strongly favor the reinstatement account of primacy with an exponential recency function. These results suggest that models of free recall can be more constrained by considering complete RT distributions.

Keywords: free recall; evidence accumulation models; RT distributions; serial position effects

Using response time distributions and race models to characterize primacy and recency effects in free recall initiation

How do participants initiate free recall from a list of studied items? Empirical investigations have found that the first response in a participant's free recall sequence shows the same qualitative trends as their performance across the entire sequence, namely an advantage for the early list items (the primacy effect) and a stronger advantage for the end-of-list items (recency effect: Howard & Kahana, 1999; Laming, 1999). These first responses have been shown to be consequential for the recall of the entire list: When participants initiate at the beginning of the list they often show strong primacy effects and little recency, whereas initiation at the end of the list often results in the opposite pattern, namely strong recency and little primacy (Ward, Tan, & Grenfell-Essam, 2010). Process models of free recall have characterized recall initiation as an instance of multi-alternative decision making, specifically a competition among a set of to-be-retrieved items using Luce's choice rule (Raaijmakers & Shiffrin, 1981). More recently, models such as the temporal context model with accumulators (TCM-A: Sederberg, Howard, & Kahana, 2008) and its successor, the context maintenance and retrieval model (CMR: Polyn, Norman, & Kahana, 2009), have leveraged the successes of the highly successful evidence accumulation models of decision making in the domain of free recall. In evidence accumulation models, each decision alternative consists of an accumulator which races toward a decision threshold; the alternative that reaches the threshold first is chosen and the time taken to reach the threshold, plus some additional time for non-decision processes, is the response time (RT). The strength of the evidence accumulation framework lies in their ability to jointly account for RT distributions and choice probabilities across decision alternatives and conditions. In this article, we address the extent to which a joint account of RT distributions and recall probabilities in first recall constrains theoretical explanations of primacy and recency in free recall initiation.

Initially, recency effects were theorized to be the consequence of the end-of-list items

being accessible in a short-term store (STS). In dual store models of recall, participants initiate recall by "dumping" the contents of STS before initiating retrieval from a long term store (LTS, e.g., Atkinson & Shiffrin, 1968; Raaijmakers & Shiffrin, 1981). When recall is delayed by a demanding distractor task, participants can only initiate recall from the LTS, predicting no advantage for recency items which is what is generally found in delayed free recall (DFR) paradigms (Glanzer & Cunitz, 1966; Howard & Kahana, 1999). The discovery of long-term recency demonstrated the dual store account was insufficient to explain the full range of recency effects. This phenomenon was demonstrated in the continual distractor free recall (CDFR) paradigm, where demanding distractor activity was placed not just at the end of the list but between each of the presentations of the study list items. Under such conditions, participants show recency advantages (Bjork & Whitten, 1974; Howard & Kahana, 1999). An increasingly popular account of recency effects that addresses this phenomenon involves associations of items to a gradually changing temporal context representation that changes during each item presentation and during filled delay activity (Estes, 1955; Mensink & Raaijmakers, 1988; Howard & Kahana, 2002). At test, the current context is used as a cue for each of the list items, and this cue is more likely to match the recency items. Current dual store models have similarly appealed to a temporal context representation to explain long-term recency (Davelaar, Goshen-Gottstein, Ashkenazi, & Haarmann, 2005; Lehman & Malmberg, 2013).

Considerably less consensus has emerged on the locus of primacy effects. Candidate mechanisms include decreasing attention paid to the list items during presentation of the study list (Serruya, Sederberg, & Kahana, 2014), greater temporal distinctiveness for the first item (G. D. A. Brown, Neath, & Chater, 2007), enhanced encoding of the first item due its novelty (Davelaar, 2007; Elhalal, Davelaar, & Usher, 2014; Farrell & Lewandowsky, 2002), extra resources devoted to the first item in a capacity-limited buffer (Atkinson & Shiffrin, 1968; Lehman & Malmberg, 2013; Raaijmakers & Shiffrin, 1981), or emerging from the dynamics of a competitive lateral inhibition network (Davelaar et al., 2005; Grossberg

& Pearson, 2008). In this article, we argue that despite the psychological differences between these explanations, what unites them is a perspective that primacy and recency items jointly compete to be retrieved, with strength being the only factor that distinguishes the different serial positions. Therefore, throughout this article we will refer to such models as *strength models*.

Strength models stand in contrast to other accounts that assume that the primacy advantage arises from a mixture of different strategies or processes. One version of this idea is the argument that the primacy advantage occurs because on some trials primacy items are rehearsed through the duration of the list until the recall period, making them functionally recent (*rehearsal models*; Tan & Ward, 2000). Alternatively, other models assume that participants initiate recall with different cues for the start and end of the list (Farrell, 2012; Lehman & Malmberg, 2013; Metcalfe & Murdock, 1981). For instance, in Metcalfe and Murdock’s model, items are associated to each other and to a cue that indicates the beginning of the list. This start cue can be reinstated, which will provide a strong match to the beginning of list items, otherwise recall can be cued with the end-of-list context, producing a recency advantage. Start of list reinstatement is also a standard assumption in positional models of serial recall, a task where it is more crucial to initiate recall at the beginning of the list (J. Anderson & Matessa, 1997; Henson, 1998; G. D. A. Brown, Preece, & Hulme, 2000; Farrell, 2006). To date, little work has discriminated between these mechanisms of the primacy effect.

Perhaps the most thorough and recent account of primacy and recency in probability of first recall (PFR) curves came from Healey and Kahana (2014). Healey and Kahana collected a very large number of participants in a multi-session free recall study and found considerable variability across participants in the shapes of their PFR curves, with most showing strong recency but some participants showing strong primacy. They accounted for these differences with CMR, which explains recency with a changing temporal context representation. Variation in the sharpness of the recency gradient can be explained via

variation in the rate of contextual change, with rapid contextual change producing strongly peaked recency functions. Primacy in CMR is explained via an ad hoc strength gradient that enhances the strengths of the primacy items, and the primacy items directly compete with the recency items during recall (i.e.; a strength model). The model gave a strong account of the individual differences in their data, which led Healey and Kahana to conclude that the model was providing a successful account of free recall initiation.

However, there are two possible insufficiencies with the Healey and Kahana (2014) account of free recall initiation. The first is that temporal context models generally produce exponential recency functions with an asymptote of zero (Howard, 2004). To date, a large literature quantifying the form of recency functions over a range of memory tasks has concluded that a power law function provides a vastly superior account over such exponential functions (J. R. Anderson & Schooler, 1991; Averell & Heathcote, 2011; Donkin & Nosofsky, 2012a; Rubin & Wenzel, 1996; Wickelgren, 1974; Wixted & Ebbesen, 1991, 1997). While Healey and Kahana (2014)'s model fits showed a strong account of recency effects across individual participants, it remains possible that a power law function provides a better account. In addition, despite CMR being capable of making RT predictions by usage of its evidence accumulation process to produce a decision, their modeling exercise did not consider the constraints imposed by RT distributions. We will demonstrate in this article that three aforementioned classes of primacy (strength, rehearsal, and reinstatement) make distinct predictions about how response time (RT) distributions vary over serial position despite making nearly identical predictions about the shapes of PFR curves.

To date, computational models have yet to jointly model RT distributions and serial position in free recall initiation. A number of investigations have modeled RT distributions using ex-Gaussian distributions and random search models (Rohrer & Wixted, 1994; Wixted & Rohrer, 1993, 1994; Unsworth & Engle, 2007), but this work did not jointly consider serial position and thus cannot address the questions we raise here about how

primacy and recency items compete during the course of retrieval. The CMR model (and its predecessor TCM-A) jointly accounted for latency and serial position, but to date has only accounted for mean RTs and not the entire distribution of latencies. Historically, RT distributions have been found to exert considerably more constraint on cognitive models than mean RTs. A seminal demonstration of this came from Ratcliff and Murdock (1976), who found that a number of models of the Sternberg short-term memory paradigm that were successful in predicting mean RTs were unable to account for the shapes of RT distributions. A number of other examples exist in the literature where RT distributions were used to constrain or even falsify cognitive models (e.g., Donkin & Nosofsky, 2012b; Donkin, Nosofsky, Gold, & Shiffrin, 2013; Fific, Little, & Nosofsky, 2010; Gronlund & Ratcliff, 1991; Ratcliff & Smith, 2004; Ratcliff, Van Zandt, & McKoon, 1999; Starns, Ratcliff, & McKoon, 2012).

Here, we show that the distributions of initial recall times can be used to determine how primacy and recency items compete at recall initiation by directly comparing the strength, rehearsal, and reinstatement mechanisms of primacy within evidence accumulation models. In answering these questions, we also address a number of interesting related questions about the recall process. One concerns the shape of the recency function: while theories often assume that strength falls off exponentially from the end of the list, a power law function might better characterise the gradient, and we directly compare both by implementing these as functional forms of drift rates. A second question is about when people actually begin the recall process. In most experiments, the onset of the cue to recall is predictable, and participants could feasibly initiate recall prior to cue onset. Evidence accumulation models allow us to answer this question by measurement of the duration of nondecision time that is added to the latency of the decision process to produce the RT. Nondecision time is traditionally positive because in many other tasks participants require information present in the cue to begin their decision process. However, in free recall participants already have access to pool of potential responses in advance of the cue, which

allows for anticipatory accumulation that would be evidenced by negative estimates of nondecision time. In this work, we factorially manipulated each of the components relevant to these questions (primacy: strength vs. rehearsal vs. reinstatement, recency: power vs. exponential, accumulation onset: anticipatory vs. cued) in a large scale model selection procedure across multiple datasets within two dominant evidence accumulation models.

Evidence Accumulation Models

We use two different evidence accumulation models in this article: the racing diffusion model (Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014; Tillman & Logan, 2017; Usher, Olami, & McClelland, 2002) and the second is the linear ballistic accumulator model (LBA: S. D. Brown & Heathcote, 2008). Both of these frameworks have been highly successful in explaining response times and accuracy across a range of different tasks. These models were chosen because both models' first passage time distributions can be expressed with closed-form analytic solutions. The leaky competitive accumulator (LCA) model (Usher & McClelland, 2001)—which was used as the back-end of the TCM-A and CMR models—is intractable, requiring lengthy simulation times and additionally shows poor parameter recovery (Miletic, Turner, Forstmann, & van Maanen, 2017). Both the racing diffusion and LBA are instances of independent race models, and under general conditions such models conveniently produce identical predictions about choice probabilities to Luce's choice rule (Bundesen, 1993), which is commonly used in free recall models to transform memory strengths into response probabilities (e.g., Davelaar et al., 2005; Morton & Polyn, 2016; Raaijmakers & Shiffrin, 1981). The diffusion and LBA models have not yet been applied to free recall, but have been highly successful in accounting for a variety of phenomena in short and long-term recognition memory (Criss, 2010; Donkin & Nosofsky, 2012b, 2012a; Osth, Dennis, & Heathcote, 2017; Osth, Bora, Dennis, & Heathcote, 2017; Osth, Jansson, Dennis, & Heathcote, 2018; Ratcliff, 1978; Ratcliff & Smith, 2004; Ratcliff & Starns, 2009; Starns et al., 2012). We use both models

here not for the purposes of placing them in competition with each other, but instead to assess whether the results of our model comparison produce similar conclusions across both the racing diffusion and LBA models.

A diagram illustrating the two models can be found in Figure 1 for the case of a six item list containing the words "METAL", "BUBBLE", "WOOD", "AIR", "CRASH", and "QUICK." Each item on the list is represented by its own accumulator with evidence threshold b . Evidence begins to accumulate from a random point that is sampled from a uniform distribution with height A . Evidence accumulates for each word with a drift rate v , with higher drift rates producing faster accumulation to the threshold. The response time is the time taken to reach the boundary, plus additional nondecision time t_0 that corresponds to encoding the recall cue and producing the response. The parameter B represents the distance from the height of the starting point distribution to the upper threshold b , and provides more efficient parameter estimation as it avoids estimation of cases where the threshold is lower than the height of the starting point distribution.

In the racing diffusion model (top row), evidence accumulates in a noisy fashion by adding zero-centered Gaussian noise on each timestep (with standard deviation s). This noise makes it such that items with lower drift rates (such as "METAL" or "BUBBLE") can sometimes receive large boosts and reach the threshold before the accumulator with the highest drift rate. In the LBA model (bottom row), evidence accumulation is linear and deterministic, but the drift rate on each trial is a sample from a truncated normal distribution with standard deviation s with a lower bound of zero. The between-trial variability in the LBA serves a similar purpose as the noise in the diffusion model, in that on some trials an item with a low mean drift rate can sample a high drift rate from the drift rate distribution by chance and win the competition over an accumulator with a higher mean drift rate. The s parameter in each model is fixed (usually to 1) for scaling, but only needs to be fixed for one accumulator for the models to be identifiable (Donkin, Brown, & Heathcote, 2009).

Figure 1. Diagram of the racing diffusion model (top) and linear ballistic accumulator (LBA) model (bottom row). In both models, each list item receives its own accumulator and evidence races to a retrieval threshold b from a starting point that is drawn from a uniform distribution A . In the racing diffusion model, evidence accumulates with drift rate v and is noisy, making it such that accumulators with low drift rates can still win the competition due to noise pushing it toward the top boundary. In the LBA model, evidence accumulates deterministically but each drift rate is sampled from a distribution, such that accumulators with low mean drift rates can by chance sample a high drift rate and win the competition.

Ordinarily, it is quite common for response thresholds to vary across accumulators to capture biased responding (S. D. Brown & Heathcote, 2008; Usher & McClelland, 2001). As response thresholds for an accumulator are decreased, it becomes more likely for that accumulator to win the competition as it has less distance to travel. While it might be sensible to allow thresholds across particular words to capture the fact that some words that are more accessible (such as high frequency words) require less evidence to retrieve, words and serial positions in free recall experiments are usually de-confounded by the randomization in the experimental design. Thus, we fixed the response threshold across all accumulators, which is additionally the standard assumption in the TCM-A and CMR

models. Variation in thresholds between words is captured by the starting point variability parameter A (in a linear model, starting points and thresholds make identical predictions). Due to the noise in the accumulators, in some preliminary fits the models often produced recall rates above zero even if the drift rate for the accumulator was zero, which imposed difficulties in fitting participants' data when recall rates were at zero for some serial positions. To solve this problem, in Appendix A we describe how the noise for each accumulator was determined to be proportional to the accumulator's drift rate.

The work in the present article is focused on participants' first recalls in their free recall sequences (Hogan, 1975; Howard & Kahana, 1999; Laming, 1999). A major issue that precludes using full sequences of recalls is that free recall is dominated by sequential dependencies: a response in free recall is often followed by report of a nearby item on the list (Kahana, 1996). A complete account of free recall sequences requires a mechanism for how the drift rates for each accumulator change throughout the recall process. In addition, the model would require mechanisms for intrusions from prior lists, erroneous repetitions, and the decision to terminate recall (e.g.; Lohnas, Polyn, & Kahana, 2015; Miller, Weidemann, & Kahana, 2012). Some of these mechanisms have been detailed in current process models of free recall; we elaborate more on potential mechanisms to extend to complete sequences in the General Discussion.

Nevertheless, the first recall may be the single most diagnostic response in free recall sequences. Howard and Kahana (1999) demonstrated that the same patterns of primacy and recency effects found in the serial position curves for full sequences can also be found in participants' probability of first recall (PFR) curves. Specifically, they compared immediate free recall (IFR), delayed free recall (DFR), and continual distractor free recall (CDFR). The patterns of recency and primacy, namely weak primacy, strong recency in IFR and CDFR and attenuated recency DFR, were evident in both levels of overall recall and first recalls in the data of Howard and Kahana (1999). This suggests that the application of evidence accumulation models to participants' first recalls can yield

diagnostic information about the mechanisms of primacy and recency in free recall. In the next section, we will discuss how the various theoretical treatments of primacy and recency can be implemented in the racing diffusion and LBA models.

Recency Functions

In order to implement primacy and recency in our evidence accumulation models, we adopt various functional forms and mixtures of such forms to implement different psychological mechanisms of recency and primacy. In terms of recency, perhaps the most empirically supported functional form is the power function:

$$v_i = \alpha_i t^{-\beta} \quad (1)$$

where v_i is the drift rate for the accumulator corresponding to the item studied in serial position i , α_i is a scaling parameter for accumulator i , t is the study-test lag (measured as $L - i$, with L being the number of items on the list), and β is the decay rate. α is allowed to vary by serial position to capture the effects of primacy (described later).

Models that employ a gradually changing context to capture recency, such as TCM and CMR, implicitly use exponential functions instead of power functions, where the rate of contextual change determines the sharpness of the exponential decay (Howard, 2004). In TCM, each studied list item enters a temporal context vector which depresses the activity of all other items in the context by a scalar, producing event-driven exponential decay. Exponential decay is of the form:

$$v_i = \alpha_i e^{-\beta t} \quad (2)$$

where the parameters have the same interpretation as Equation 1. Unlike some prior tests

of empirical forgetting functions (e.g., Averell & Heathcote, 2011; Rubin & Wenzel, 1996), we did not include an asymptote parameter for the exponential function as there is no mechanism in current temporal context models for such an asymptote¹. The two functions can be distinguished on how they decay with respect to study-test lag: the exponential’s proportional rate of decay is constant with respect to lag, while the power law function’s proportional rate of decay decreases as lag increases. This property accords with Jost (1897)’s second law, which states that older memories are more resistant to forgetting than newer ones (Wixted, 2004).

The widespread support for the power law function over the exponential might make the comparison seem like a foregone conclusion. However, models such as TCM and CMR produce convincing fits to PFR curves despite their reliance on exponential functions of recency (Healey & Kahana, 2014; Polyn et al., 2009; Sederberg et al., 2008). In addition, Wixted and Ebbesen (1997) found that forgetting functions in free recall were well described by exponential curves.

Nonetheless, a limitation of many comparisons of empirical forgetting functions is that they draw conclusions about the form of the recency gradient based on recall probabilities alone. A problematic aspect of this type of model comparison is that the functional form of the curve depends on the measurement variable (such as probability of recall); transformation of such variables can potentially produce support for a different function. For this reason, Donkin and Nosofsky (2012a) advocated comparing recency functions on a latent psychological variable within a process model, such as drift rates in an evidence accumulation model. This is exactly the approach we take here, which further extends on previous work by jointly accounting for the effect of lag on recall probabilities and RT distributions in free recall initiation.

In the case of free recall initiation, a further reason that simple curve fitting of recall

¹While this is not a standard assumption in current models of free recall, temporal context models may be able to introduce an analog to an asymptote parameter by assuming that a proportion of the context elements do not change.

probability can be misleading is that noise in the decision process can obfuscate the relationship between memory strength and the probability of recall. To illustrate this point, the left panel of Figure 2 shows a power law (blue) and exponential (red) function on the latent dimension of drift rates (i.e., memory strength). Although both functions are strongly peaked at the final list item, the power law function shows a more gradual decay with respect to study-test lag and the earliest items are of moderate strength compared to the earliest items in the exponential function. We used these functions to generate probability of first recall (PFR) curves with three different levels of the response threshold parameter B using the racing diffusion model, which can be seen in the middle panel. At the lowest level of the threshold ($B = .2$, depicted by the circles), retrievals are dominated by noise in the decision process, namely the variability in the starting point (A) and the within-trial noise (in the case of the racing diffusion model). This is because the threshold is sufficiently low that each accumulator does not have a large distance to travel, meaning that a spuriously high starting point can be sufficient for an accumulator to win the retrieval competition even if it has a very low drift rate. Because retrieval is dominated by noise, decision making resembles random guessing and consequently the PFR functions resemble straight lines with slight upticks for the recency items. In this case both PFR curves show the very gradual recession in recall probability characteristic of a power law function, even when the generating function was an exponential.

As the response threshold is raised ($B = 2.0$, depicted by triangles), the higher response thresholds make it such that noise in the decision process is less likely to drive an accumulator with a low drift rate to the response threshold. Consider a potential case where the accumulator for serial position 1 (which has the lowest drift rate) starts with a very high starting point relative to the other accumulators. Since this accumulator still has a substantial distance to travel to reach the response threshold, the accumulators with higher drift rates have time to overtake this accumulator; their faster rate of travel implies that accumulator #1 has a low probability of reaching the threshold first. As the response

Figure 2. The left panel shows an exponential and power law function of drift rates. The middle panel shows probability of first recall predictions (PFR) from each function from three different levels of the response threshold parameter B (circles = .2, triangles = 2.0, and stars = 6.0) in the racing diffusion model. The right panel shows the RT distribution predictions from each recency function and response threshold, summarized using the .1, .5 and .9 quantiles.

threshold is raised even higher ($B = 6.0$, depicted by stars), noise in the starting point are unlikely to affect the decision process at all, producing retrievals only for the items with the highest drift rates. In this case, both the exponential and power law recency functions in the left panel produced strongly peaked PFR curves in the middle panel that are often characteristic of an exponential recency function.

If PFR curves can vary considerably depending on the nature of the response thresholds, so that it is difficult to infer backwards from the shape of the PFR to the shape of the underlying latent variable of memory strength, how can recency functions be identified? The right panel depicts the RT distribution predictions from each combination of model and parameterization collapsed across all serial positions. The RT distributions are summarized using the .1, .5, and .9 quantiles of the RT distribution, which are the 10th, 50th, and 90th percentiles of the distribution and are arranged from bottom to top in the figure. The .1 quantile is referred to as the *leading edge* of the RT distribution and

represents the fastest responses, while the .9 quantile is the upper tail of the RT distribution and represents the slowest responses. One can see that as the response threshold is increased, it produces a large increase in the leading edge of the RT distribution and increases the overall variability. Changes in the drift rate parameter, in contrast (depicted later in the article), produce large changes in the upper tail of the RT distribution but comparatively little change in the leading edge (Ratcliff & McKoon, 2008). In this work, we jointly measured the response threshold, decision noise, and recency function within evidence accumulation models. By estimating the response threshold and decision noise, we are better able to estimate the underlying recency function, overcoming the limitations of prior curve fitting approaches.

Mechanisms for Generating Primacy

As mentioned previously, a recall advantage for primacy items can arise from different mechanisms. The simplest implementation assumes that the beginning of list items have extra strength via a primacy gradient (Donkin & Nosofsky, 2012b; Farrell & Lewandowsky, 2002; Healey & Kahana, 2014; Sederberg et al., 2008). This can be incorporated into Equation 1 or 2 by setting an exponential primacy gradient on the α parameters of the recency functions:

$$\alpha_i = a(r_s \exp(-r_d(i - 1)) + 1) \quad (3)$$

where i is the serial position of the accumulator, r_s and r_d are the primacy scale and decay parameters, and a is the base recency scale parameter. If r_s is zero, α_i reverts to a . In addition, if r_d is sufficiently large, it produces a relatively steep primacy gradient such that α_i will be approximately equal to a for the later list items. We refer to this class of models as the strength model of primacy because it assumes primacy items are merely stronger than the items that directly follow them; strength is additionally the only variable that

differentiates between the primacy and recency items. Equation 3 is identical to the primacy gradient function used by the TCM-A and CMR models.

Tan and Ward (2000) proposed an alternative account of primacy whereby rehearsal can sometimes cause the primacy items to be rehearsed all the way to end of the study list, making them functionally recent. When participants are instructed to rehearse out loud during presentation of the study list, it appears that a significant proportion of the primacy effect is due to the recency and frequency of their rehearsals (Brodie, 1975; Brodie & Murdock, 1977; Rundus, 1971; Tan & Ward, 2000). We can implement a simple version of such a rehearsal model by assuming a probability mixture of a) the serial position functions described in Equations 1 and 2 and b) permuted serial position functions where the primacy item is considered to be the most recent on the serial position function, and all the other items are pushed backward one position. This introduces an additional mixing parameter p that determines the proportion of times that recall is initiated with the permuted serial position function. Primacy items that are functionally recent are often stronger than recency items that are functionally recent due to the additional frequency of their rehearsals (Tan & Ward, 2000); for this reason we add an additional drift rate parameter v_r to the first item in the permuted recency function.

When the mixture parameter $p = 0$, the first item is never functionally recent and the model reverts to the predictions of the strength model. Equation 3 still allows for a primacy advantage even if $p = 0$; in the framework of a rehearsal model, this is explained as the primacy items being rehearsed more often than later list items even if they are not rehearsed to the end of the study list (Tan & Ward, 2000). One should note that other early list items such as the second or third items are also often rehearsed to the end of the list. We have chosen to implement a relatively simple model where such a process only occurs for the first item to avoid proliferation of a large number of mixing parameters to represent each possible rehearsal sequence.

The final account we consider assumes that primacy is due to a reinstatement of the

start-of-list context to initiate retrieval (Farrell, 2012; Laming, 1999; Lewandowsky & Murdock, 1989; Metcalfe & Murdock, 1981). This can be implemented in the models considered here by assuming a probability mixture of a recency function and a primacy function. Monotonicity is imposed on the recency function by setting $\alpha_i = a$. Drift rates for the primacy function are an exponentially decreasing function of serial position:

$$v_i = r_s \exp(r_a(i - 1)) \quad (4)$$

The reinstatement model also has a mixing parameter p , but this mixing parameter serves a different function than the one in the rehearsal model. In the reinstatement model, p determines the proportion of times that recall is initiated with the primacy gradient in Equation 4; otherwise, recall is initiated with a recency gradient, which can be either exponential or power law in shape.

Table 1 provides a complete listing and description of the parameters for each of the discussed models. Predicted drift rates, PFR curves, and RT distributions from the strength, rehearsal, and reinstatement models can be seen in Figure 3 using the racing diffusion model to generate predictions (predictions are qualitatively similar for the LBA model). In the rehearsal and reinstatement models, two drift rate functions are depicted: a recency-heavy function is depicted in black which contributes on a proportion $1 - p$ trials, while the primacy-oriented function is shown as a dotted line which contributes on a proportion p trials. For the rehearsal model, the primacy function is a permutation of the recency function such that the first item has the strength of the final item in the recency function, with the exception that there is an additional strength v_r added to the first item that reflects the additional strength from the fact that the primacy item was rehearsed more frequently than the last item. All other items in the permuted function are pushed “backward” such that the last item in this function has the strength of the second to last

item in the recency function, and so on. For the reinstatement model, the primacy function is an exponential function that is independent of the recency function. Parameters of the models were chosen in such a way as to produce PFR curves typically seen in immediate free recall experiments, namely a strong recency effect and a small primacy effect (e.g., Howard & Kahana, 1999). One can see that each model produces virtually identical PFR curves, reinforcing that these models cannot be distinguished on the basis of recall probabilities alone.

Table 1

Parameter names and descriptions. Notes: Param = parameter.

Param	Description
t_0	Nondecision time constant which is added to the decision time to produce the response time.
B	Distance between the height of the starting point distribution and the response boundary.
A	Height of the starting point distribution.
s_{min}	Minimum diffusion noise in the racing diffusion model/minimum cross-trial drift rate variability in the LBA (functional form and motivation are described in Appendix A).
α	Scale parameter for the recency function (power or exponential).
β	Decay parameter for the recency function (power or exponential).
r_s	Scale parameter for the primacy function.
r_d	Decay parameter for the primacy function.
p	Mixture parameter that determines the proportion of primacy trials. For the rehearsal model, this is the probability that rehearsal makes the first item functionally recent. For the reinstatement model, this is the probability of initiating retrieval with the primacy function instead of the recency function.
v_r	Drift rate increase for the first item when it is functionally recent (which occurs with probability p) in the rehearsal model.

The bottom row shows RT distributions for each serial position, summarized using the .1, .5, and .9 quantiles (arranged from bottom to top). Each model produces a different

Figure 3. Predicted drift rates (top row), probability of first recall (PFR) curves (middle row), and .1, .5, and .9 quantiles of the RT distribution for each serial position (bottom row) for the strength (left column), rehearsal (middle column), and reinstatement (right column) models. Model predictions were generated using the racing diffusion model. For the rehearsal and reinstatement models, $p = .15$.

patterns of RTs across serial positions. The strength model is distinguished from the other two models because it predicts very similar response times for each serial position. This is counter-intuitive when one considers the fact that drift rates vary considerably by serial position, with the last item having a much higher drift rate and thus a substantial recall advantage. Given that increases in drift rate produce increases in speed, shouldn't the item with the highest drift rate exhibit much faster response times, and shouldn't the weaker items exhibit much slower times? In race models, this is not the case: for the accumulators

with lower drift rates to win the competition, *they have to beat the accumulator with the highest drift rate*, producing similarly fast decisions.

Figure 4 illustrates three different drift rate functions where the drift rate for the strongest item is varied across three levels along with simulations from the racing diffusion (left) and LBA (right) models. Raising the drift rate for the strongest item has the ironic effect of *decreasing* the response time for each of the other accumulators (bottom row), an effect most noticeable in the upper tail of the RT distribution (the .9 quantile). There is comparatively little change in the leading edge (.1 quantile), in contrast to changes in the response threshold parameter B seen in Figure 2. When the strongest item receives an increase in strength, it wins the race more frequently. Due to its increase in speed, it restricts recall of the other accumulators to cases where they were able to outperform the strongest accumulator. These may be cases where such accumulators received spuriously high starting points from the starting point distribution, or cases where the noise happened to produce very high accrual of evidence (in the racing diffusion model), or where the weaker accumulators spuriously sampled a high drift rate (in the LBA model). Thus, the RTs for each serial position in the strength model are strongly constrained by the drift rate for the strongest item in the competition. In Supplementary Materials B, we present simulations from the leaky competitive accumulator (LCA) model and demonstrate that across a range of values of lateral inhibition and leakage that were employed by Sederberg et al. (2008) and Polyn et al. (2009), RTs are predicted to be very similar for each serial position despite large differences in the drift rates for each accumulator. Thus, the LCA may yield similar qualitative predictions as the racing diffusion and LBA models in this context, altho

The mixture models, namely the rehearsal model and the reinstatement model, have predictions that diverge from the strength model because they involve two different drift rate functions with different maximum drift rates. In the rehearsal model, there is a proportion p of cases where the primacy item has the highest drift rate, and this drift rate

Figure 4. Illustration of how changes in drift rate for one accumulator can affect the RTs of the other accumulators. The top row shows three different drift rate functions that differ with their strength for the final and strongest item. The left and right columns show predictions from the racing diffusion and LBA models, respectively. The middle row shows the PFR curves while the bottom row shows predicted RT distributions for each serial position summarized by the .1, .5, and .9 quantiles.

is higher than the highest drift rate in the recency-based function that occurs with probability $1 - p$. If one focuses on the RTs for each of these functions individually, there is a similar pattern as the strength model, namely similar RTs for each serial position. However, when a weighted average of the RTs from each function is taken, a different pattern emerges - the RT for the first item is predicted to be the fastest.

The reinstatement model, in contrast, has a strongly peaked recency function with a high maximum drift rate as well as a less well peaked primacy function with a lower maximum drift rate. Similar to the rehearsal model, each of these drift rate functions individually yield very similar RT distributions for each serial position with the difference being that the primacy function yields slower RTs due to its lower maximum drift rate.

When a weighted average is taken across both functions, it yields a pattern where RTs are very similar for the later list items but get much slower for the earliest items on the list, especially in the tail of the distribution. The reinstatement model thus produces a dissociation between accuracy and RT, with the first item exhibiting the slowest RT despite the fact that it has a recall advantage over several of the later items on the list.

One should note that the predictions of each of the models depend on the values of the free parameters. In particular, increases in starting point variability produce faster responses for accumulators with lower drift rates relative to ones with higher drift rates (Ratcliff & Rouder, 1998). Increases in the response threshold have the opposite effect, producing faster responses for items with higher drift rates (Tillman & Logan, 2017). Thus, observation of a qualitative pattern in the RT data that matches the predictions of one of the models depicted in Figure 3 is insufficient to draw strong conclusions. In the next section, we describe how each of the models were quantitatively fit to the data.

Cued vs. Anticipatory Evidence Accumulation

In addition to the factorial manipulation of the recency and primacy factors, we also tested two different assumptions about the non-decision time parameter t_0 that is added to the RTs. In free recall experiments with predictable list lengths, it is possible for participants to begin evidence accumulation before the experimenter provides the cue to begin recalling the list items, especially given that participants already have access to the potential responses prior to the onset of the cue. In what we call *cued accumulation* models, evidence does not begin accumulating until time t_0 the recall cue occurs (t_0 can only take on positive values), which is the way t_0 is traditionally measured in evidence accumulation models. In *anticipatory accumulation* models, we allow t_0 to take on negative values, which allows for the possibility that participants can begin accumulating evidence *before* they are exposed to the recall cue.

The Model Fit

To investigate each of the aforementioned possibilities we conducted a $2 \times 3 \times 2$ factorial model comparison, varying the shape of the recency function (exponential vs. power), the mechanism producing primacy (strength vs. rehearsal vs. reinstatement), and assumptions about nondecision time (cued vs. anticipatory). Each cell of the design was fit using both the racing diffusion and LBA models; conclusions that are consistent across both models are more likely to be robust.

Because we are modeling only the first response on each trial, and each trial can last up to a minute or more (depending on the time allowed for recall), each participant provides a relatively limited set of useful data. Accordingly, each of the models was fit using hierarchical Bayesian modelling (see Lee, 2011; Rouder & Lu, 2005, for introductions), which simultaneously estimate both group and subject level parameters. When there is little data at the participant level to constrain the model parameters, estimates of the participant parameters get pulled toward the group estimate, a phenomenon referred to as “shrinkage”. Vandekerckhove, Tuerlinckx, and Lee (2011) demonstrated the specific advantage of using hierarchical modeling for diffusion models. To impose minimal constraint on the values of the parameters, we used approximately non-informative prior distributions on each of the parameters (see Appendix A for more details).

Estimation of the posterior distribution in hierarchical Bayesian modelling typically requires the use of Markov chain Monte Carlo (MCMC) algorithms. One challenge associated with evidence accumulation models is that the parameter estimates are often correlated with each other (Ratcliff & Tuerlinckx, 2002; Turner, Sederberg, Brown, & Steyvers, 2013), which is problematic for conventional MCMC algorithms. For this reason, we used differential evolution Markov chain Monte Carlo (DE-MCMC: Turner et al., 2013), a method of posterior sampling that is robust to parameter correlations. Details of the DE-MCMC procedure can be found in Appendix A.

The datasets used in our model comparison are summarized in Table 2. All of the

datasets come from the University of Pennsylvania Computational Memory Lab website ². We restricted consideration to datasets and conditions that presented unique items and where participants are required to recall only the items from the immediately preceding list. Each of the experiments used vocalized recall, and the response time was scored at the beginning of the onset of articulation PEERS refers to the Pennsylvania Electrophysiology of Encoding and Retrieval Study (PEERS), which is a large scale, multi-session free recall study used in several publications. PEERS Experiment 1 was previously reported by Healey and Kahana (2014, 2016); Lohnas et al. (2015), Experiment 2 was reported by Lohnas and Kahana (2014), and Experiment 4 was reported by ? (?). Age groups such as college-aged and older participants were fit separately because despite exhibiting similar PFR curves (Kahana, Howard, Zaromb, & Wingfield, 2002), we found that older participants often exhibited slower RTs, which is consistent with findings in other tasks such as recognition memory (Ratcliff, Thapar, & McKoon, 2004). Additional details on each of the datasets can be found in Supplementary Materials A while response and participant exclusions are described in Appendix A.

Table 2
Summary of the datasets fit by the model.

Number	Dataset	Type	N	Obs.	LL
1	PEERS E1 Younger	Immediate	172	93.93	16
2	PEERS E1 Older	Immediate	38	88.95	16
3	PEERS E2	Delayed	156	26.46	16
4	PEERS E2	Continual distractor	156	26.51	16
5	PEERS E4	Delayed	98	527.49	24
6	Howard et al. (2007)	Immediate	292	22.73	10
7	Murdock & Okada (1970)	Immediate	72	19.75	20
8	Polyn et al. (2009)	Immediate	45	27.77	24

²<http://memory.psych.upenn.edu>

9	Kahana et al. (2002): E1 Young	Immediate	29	27.82	10
10	Kahana et al. (2002): E1 Older	Immediate	26	26.96	10
11	Kahana et al. (2002): E2 Young	Delayed	28	18.8	10
12	Kahana et al. (2002): E2 Older	Delayed	25	17.04	10
13	Golomb et al. (2008) Young	Immediate	36	33.72	10
14	Golomb et al. (2008) Older	Immediate	36	33.91	10

Notes: E = experiment, N = number of participants, Obs. = mean number of observations per participant, LL = study list length. E = experiment.

The data from Dataset 5 (Kahana, Aggarwal, & Phan, 2018) and 7 (Murdock & Okada, 1970) are unusual in that the PFR curve in each dataset is nonmonotonic; recall peaks three positions before the final item, a result that is inconsistent with models that assume monotonic recency functions such as ours. For these specific experiments we added two additional drift rate parameters for the final (L) and penultimate ($L - 1$, where L is list length) list items that are sampled on a $(-\infty, \infty)$ interval to allow for the deflections from a monotonic recency function. While some researchers have argued that this nonmonotonicity is evidence for a short-term store (Davelaar et al., 2005) or access to a final group (Farrell, 2012), the drift rate parameters we include are agnostic to these psychological possibilities and were purely used to capture the data.

It should be noted that each of the models vary in their degree of complexity and allowance to capture the data. The mixture models of primacy, for instance, have more freedom to capture RT variability across serial position. Thus, we compared the models using the widely applicable information criterion (WAIC: Watanabe, 2010), an approximation to leave-one-out cross validation that takes into account the trade-off between goodness-of-fit and model complexity. More complex models receive harsher penalties, such that any improvement in fit has to outweigh the complexity penalty for a

model to be preferred. In Supplementary Materials D, we performed a model recovery with each of the models using the fitted parameters to Dataset 1 and demonstrate that WAIC is able to recover the generating model for the models examined here.

All model code for the project can be found online at <https://osf.io/bkjqn/>.

Nondecision Time Estimates

Estimates of the group mean of the nondecision time parameter (M_{t0}) for the cued and anticipatory models of each dataset can be seen in Figure 5 using the winning models from the primacy and recency comparisons (described below). A consistency across each model is that in most of the datasets where recall is delayed (Datasets 3 and 4, which use delayed and continual distractor free recall respectively, and datasets 11 and 12, which use delayed free recall) produce positive estimates of M_{t0} even for the pre-emptive accumulation models. For the LBA, most of the other datasets show negative estimates of M_{t0} for the immediate free recall datasets (with the exception of Datasets 1 and 2), with some estimates of nondecision time nearing -1 second. For the racing diffusion model, nondecision times are consistently higher and only some of the datasets show negative estimates of nondecision time for the preemptive accumulation models. The higher estimates of nondecision time in the racing diffusion model are consistent with prior investigations that have compared the diffusion and LBA models (Donkin et al., 2009).

Supplementary Materials C presents a model selection comparison between the cued and anticipatory accumulation models. Most of the datasets where recall was delayed, namely Datasets 3, 11, and 12 (delayed free recall) along with Dataset 4 (continual distractor free recall) showed a preference for cued accumulation models, which are simpler than the anticipatory accumulation models due to their greater restriction on the t_0 parameter. However, for the immediate free recall datasets, several datasets showed preference for anticipatory accumulation models. To simplify the number of models under consideration, for the remaining duration of the article predictions and model selection

Figure 5. Comparison of the cued and anticipatory evidence accumulation models' estimates of the group mean of the nondecision time parameter M_{t_0} for the racing diffusion (left panel) and LBA (right panel) models.

results will be restricted to the winning models from this comparison.

Negative estimates of M_{t_0} from the anticipatory accumulation models were found in several of the immediate free recall datasets. While Dataset 1 shows positive estimates, one should note that M_{t_0} is a group mean parameter - for this model, some individual participants exhibited negative estimates of t_0 . Lower estimates of M_{t_0} were also found in the LBA than the racing diffusion model, which is consistent with prior results (Donkin et al., 2009; Tillman & Logan, 2017). These results suggest that during immediate free recall with predictable list lengths, participants are initiating retrieval prior to the cue to initiate recall; this is sensible when one considers that free recall differs from many other tasks in that participants are presented the acceptable responses in advance of the decision cue. In the most extreme cases, accumulation is estimated to be between 500-1,000 ms in advance. This can explain why some immediate free recall datasets show minimum RTs that are extremely short (RTs can be as low as .003 seconds for valid responses) despite the high variability in RTs. In delayed and continual distractor free recall, however, the long and demanding distractor tasks that take place prior to the recall cue are likely sufficient to prevent participants from anticipating the recall cue and initiating early retrieval. One exception to this is Dataset 5, which showed strong preference for anticipatory accumulation despite 24 seconds of a demanding distractor task in advance of the recall

period. However, participants in this dataset were also highly practiced, having done over 500 trials over 23 experimental sessions, so it’s possible that the participants learned to anticipate the recall cue after the delay period.

Model Selection Results: Primacy and Recency Comparisons

To facilitate comparison between models, we present WAIC difference scores (Δ WAIC), where the winning model for each evidence accumulation model class (racing diffusion or LBA) has a score of zero and all other models have positive scores, indicating the difference in WAIC between each model and the best model. These results can be found in Table 3. By convention, differences of 10 or larger between models are considered “large”. Each model employs the assumptions the winning assumptions about nondecision time (anticipatory vs. cued).

Table 3
WAIC difference scores for each model.

Dataset	P	Primacy	Racing Diffusion		LBA	
			Exponential	Power	Exponential	Power
DS1	8	Strength	1030	2532	710	2240
	10	Rehearsal	1329	2831	925	2352
	9	Reinstatement	0	1601	0	1563
DS2	8	Strength	59	344	75	294
	10	Rehearsal	109	412	103	383
	9	Reinstatement	0	310	0	247
DS3	8	Strength	137	312	232	447
	10	Rehearsal	132	267	48	237
	9	Reinstatement	0	213	0	222
DS4	8	Strength	208	245	70	75
	10	Rehearsal	247	314	56	93
	9	Reinstatement	0	110	0	32
DS5	10	Strength	5065	5615	5381	5921
	12	Rehearsal	3011	2177	2058	436
	11	Reinstatement	0	306	297	0

Dataset	P	Primacy	Racing Diffusion		LBA	
			Recency		Recency	
			Exponential	Power	Exponential	Power
DS6	8	Strength	504	689	535	896
	10	Rehearsal	619	770	598	898
	9	Reinstatement	0	182	0	130
DS7	10	Strength	13	56	24	200
	12	Rehearsal	42	119	226	385
	11	Reinstatement	0	44	0	40
DS8	8	Strength	126	151	155	184
	10	Rehearsal	131	175	140	170
	9	Reinstatement	0	41	0	37
DS9	8	Strength	18	76	12	78
	10	Rehearsal	31	87	17	93
	9	Reinstatement	0	53	0	80
DS10	8	Strength	0	38	0	46
	10	Rehearsal	16	52	1	52
	9	Reinstatement	2	42	1	23
DS11	8	Strength	17	18	24	26
	10	Rehearsal	14	16	32	27
	9	Reinstatement	1	0	0	3
DS12	8	Strength	0	12	0	21
	10	Rehearsal	11	21	16	24
	9	Reinstatement	10	18	7	18
DS13	8	Strength	0	89	12	99
	10	Rehearsal	110	164	160	214
	9	Reinstatement	9	92	0	92
DS14	8	Strength	34	119	20	111
	10	Rehearsal	97	179	98	181
	9	Reinstatement	0	105	0	89

Across both the datasets *and* evidence accumulation model classes, two striking consistencies emerge: support for exponential models of recency, and support for

reinstatement models of primacy. There are some exceptions; for the recency comparison, Dataset 5 in the LBA and Dataset 11 for the racing diffusion model showed support for power models of recency. For Dataset 11, the differences between the exponential and power models was small in both the racing diffusion and LBA models. The ambiguity between evidence accumulation model classes for these datasets is not terribly surprising when one considers that both of these datasets used delayed free recall and show considerably attenuated recency, as we will discuss in the next section.

There were also exceptions in the primacy comparison. For Datasets 10 and 12 there was support for a strength model of primacy in both model classes and in Dataset 13 there was support for a strength model of primacy in the racing diffusion model. However, the largest of these differences was 9 points, whereas in the other datasets the advantage of the reinstatement model ranged from large (> 10 points) to massive (> 100 points) and become as large as 500 points for Datasets 1, 5 and 6, which are three of the largest datasets. In the coming sections, we will show the predictions from the models along with the data to help understand why models were preferred in each of the datasets.

While we restricted consideration to the winning models from the nondecision time comparison (cued vs. anticipatory accumulation) to simplify exposition, the model selection result from the complete set of models can be found in Supplementary Materials C. For several of the datasets, both the cued and anticipatory classes of models produced the same preferences, namely support for an exponential model of recency and a reinstatement model of primacy. The cases where cued and anticipatory accumulation models resulted in different preferences for primacy and recency mechanisms largely occurred in the smaller datasets and in cases where support for a particular model was small in magnitude (< 10 points).

Probability of First Recall Curves

As mentioned previously, the probability of first recall (PFR) curves are unlikely to discriminate between the various assumptions of primacy. However, the advantage of the exponential recency function is evident in such curves, which can be seen in Figure 6 and 7 for the racing diffusion and LBA models, respectively. Depicted in these figures are a comparison between the power law and exponential models of recency combined with the winning primacy mechanism. Because the width of the posterior predictive distribution was often small, for convenience we simply show the average of the posterior predictive distribution. In the immediate free recall conditions, the data show a strong recency effect that quickly drops to a level slightly above zero about 4 or 5 positions from the end of the list, along with a small one position primacy effect, similar to the pattern observed in Howard and Kahana (1999). This pattern of recency is more natural for the exponential function which decays to zero. The power function, in contrast, exhibits slower decay as time progresses and thus does not approach zero as quickly. It is also interesting to note that the models are successful in accounting for near-zero recall probabilities for mid-list items. This is partly due to the scaling in Equation 5 (in the Appendix), which ensures that items with lower drift rates also have lower noise to prevent them from reaching the response threshold.

One can see from the figures that the exponential model is better able to account for the recency items in each of the immediate free recall datasets. The delayed free recall datasets (Datasets 3, 5, 11 and 12), show attenuated recency and relatively strong primacy. The ambiguity in Datasets 5, 11 and 12 over whether the exponential or power function provides a better account may stem from the fact that the recency effect was not strong enough to distinguish between these two models. While Dataset 3 showed strong advantage for an exponential model of recency, there were over five times as many participants in Dataset 3 as Datasets 11 and 12, while Dataset 5 does not show a monotonic recency function.

Figure 6. Probability of first recall curves for each dataset along with the predictions of the best racing diffusion models with exponential and power-shaped recency curves.

A comparison of how the different PFR predictions of the strength, reinstatement, and rehearsal models of primacy can be found in Supplementary Materials E. Each model

Figure 7. Probability of first recall curves for each dataset along with the predictions of the best LBA models with exponential and power-shaped recency curves.

produced extremely similar PFR curves, which concords with the earlier analysis that the models are indistinguishable on the basis of PFR curves alone.

RT Distributions

To evaluate the goodness of fit of a model to the RT distributions in the data, it is ordinarily common to summarize the RT distribution using quantiles, such as the 10th, 50th, and 90th percentile, for each participant in each condition and average them together. This approach is faced with two difficulties. The first is that some participants never initiate recall with certain serial positions; most participants initiate recall with an item at the end of the list (Healey & Kahana, 2014). An additional difficulty is the relatively small number of observations per participant, which prevent the calculation of stable quantile estimates. To counter-act both of these hurdles, we parameterically smoothed each participant’s RT distributions with a hierarchical ex-Gaussian model; the details can be found in Appendix A. The ex-Gaussian distribution has been shown to provide an excellent account of RT distributions in not only two choice tasks (Ratcliff & Murdock, 1976) but in free recall as well (Wixted & Rohrer, 1993, 1994; Unsworth & Engle, 2007). Given the relative lack of recalls in mid-list positions, we collapsed positions 2 through $L - 5$, where L is the length of the list, and estimated a single ex-Gaussian distribution for the midlist items. Each ex-gaussian distribution was then used to generate the 10th, 50th, and 90th percentiles of the RT distribution for each participant - these quantiles were subsequently compared to the predictions of the models in the text to evaluate their goodness of fit. In this context, the ex-Gaussian distribution can be considered a “data model” because it is only used to smooth the data and compare it to the predictions of the models in the figures. To be clear, all of the model fitting was done on the raw RTs; the ex-Gaussian distributions were purely used for visually assessing the fit of models to the data.

As Figure 3 illustrated, the different conceptions of primacy can be discriminated on the basis of how the RT distributions change across serial positions. Figures 8 and 9 show the group-averaged RT distribution quantile predictions of the racing diffusion and LBA models, respectively, along with the quantiles derived from the ex-Gaussian data model. Prior to calculating the group averages in each position or bin (in the case of the midlist

items) we restricted consideration to participants with at least two recalls in that position or bin.

What each dataset has in common is that the leading edge of the RT distribution is relatively flat across each serial position. Instead, differences in RTs between serial positions are primarily restricted to the median and the upper tail (the 50th and 90th percentiles). One may note that the racing diffusion and LBA models had differing degrees of fit to the RT distributions; the LBA misses on the leading edge in several cases, but produces tighter predictions than the racing diffusion model. Regardless, the qualitative behavior of the strength and reinstatement models of primacy were similar in both classes of models.

Some of the datasets demonstrate the qualitative predictions of the reinstatement model of primacy depicted in Figure 3, namely the first item having the longest RT despite having a recall advantage. Datasets 1, 3, 4, 7, and 8 clearly demonstrate this pattern and demonstrated large to enormous model selection advantages for the reinstatement model. While Dataset 5 shows the slowest RTs for the midlist items, both the midlist items and the first item show considerably longer latencies than the recency items. It is interesting to note here that even after having been fit to the data, the strength model shows the same qualitative pattern as its predictions in Figure 3; the model's RT distributions are fairly similar across each serial position. This suggests that variation in the model parameters such as the threshold parameter B and starting point variability A , which allow for differences in RT across accumulators, were sufficiently constrained to prevent the model from fitting the RT differences across serial position.

However, the reinstatement model is able to capture the much longer RTs for observed at position 1 and the mid-list items. It is also interesting to note that this tendency was also observed in Datasets 3 and 5, which utilized delayed free recall. These datasets exhibited attenuated recency and stronger primacy than several of the other datasets. Despite the fact that the first item was among the most frequently recalled in that dataset, RTs were slower than the recency items.

Figure 8. Group-averaged predicted RT distributions from the strength, rehearsal, and reinstatement variants of the racing diffusion model along with distributions from the data model. RT distributions are summarized by the 10th, 50th, and 90th percentile. Error bars indicate the 95% highest density interval (HDI).

Some datasets did not show the slowest RTs for the first item and still produced an advantage for the reinstatement model. In particular, the reinstatement model received an enormous model selection advantage for Dataset 5, where RTs are slowest for the midlist items. The reinstatement model is better able to capture the slow RTs for such items, whereas the other models predict much faster RTs for such items.

Some of the delayed free recall datasets (DS 11 and 12) did not indicate a substantial advantage for the reinstatement model. Instead, these data show relatively fast RTs for the first and last item of the list. Both the strength and reinstatement models are able to capture this pattern, although the strength model provides a better account in some of these cases. There is also considerably more uncertainty in the predictions for these datasets due to the relatively low number of participants combined with the low numbers of observations, so it is possible that these datasets are not highly diagnostic.

A similar comment applies to the immediate free recall datasets that showed advantages for a strength model, such as Datasets 9, 10, and 13. Each of these datasets did not employ a large number of participants or items. In addition, the observed RT was relatively constant across the serial positions, which is more consistent with a strength model than a reinstatement model.

None of the datasets matched the qualitative predictions from the rehearsal model shown in Figure 3, which predicts the fastest RTs for the first item. While Datasets 11 and 12 show relatively fast RTs for the first item, the last item on the list shows similarly fast RTs. Nonetheless, for those datasets, the rehearsal model appears to provide the strongest account of the RTs for the first serial position. The fact that none of the datasets demonstrate this critical prediction of the model may indicate why it performed relatively poorly in the model selection. Inspection of Figures 8 and 9 demonstrate that the rehearsal model performed quite similarly to the strength model in many of the datasets but predicted even faster RTs for the first item.

Conclusions about group averages can be contaminated by averaging artifacts. This

Figure 9. Group-averaged predicted RT distributions from the strength, rehearsal, and reinstatement variants of the LBA model along with distributions from the data model. RT distributions are summarized by the 10th, 50th, and 90th percentile. Error bars indicate the 95% highest density interval (HDI).

is especially prevalent for analyses of first recalls in free recall, given that some participants never initiate recall with certain regions of the list (Healey & Kahana, 2014), meaning that RTs for the primacy and recency portions may reflect different subgroups of participants rather than indicating a structural relationship between RT distributions at each serial position. For this reason, we depict the fits to the individual participant quantiles in Supplementary Materials F. For the majority of the datasets the reinstatement model yields the best account of the RTs across individuals.

General Discussion

The bulk of the free recall literature has focused on response proportions. While some work has approached response latencies, this work has been limited to a consideration of mean RTs in conjunction with response proportions (Polyn et al., 2009; Sederberg et al., 2008) or RT distributions alone (Unsworth & Engle, 2007; Wixted & Rohrer, 1993). In the present work, we applied two evidence accumulation models, namely the racing diffusion model and the linear ballistic accumulator (LBA) model, to provide a joint account of first recall serial position curves and complete RT distributions in thirteen datasets. In these models, we factorially manipulated various assumptions about recency and primacy and tested them in a large scale model selection procedure. In particular, we evaluated whether the recency gradient is best described by an exponential or power function and whether primacy items receive a strength boost, are sometimes treated like end-of-list items due to being rehearsed to the end of the list, or whether primacy items race without competition from the recency items due to start-of-list reinstatement. While each of the primacy mechanisms produce similar predictions about PFR curves, they make distinct predictions about RT distributions, demonstrating the need to consider complete RT distributions in models of free recall. In addition, we additionally varied whether accumulation begins prior to or after the recall cue, which would be impossible without consideration of response times.

Exponential vs power gradients of memory strength

The vast majority of the datasets produced evidence for an exponential recency function. This stands in contrast to a large and established literature favoring power functions of forgetting (J. R. Anderson & Schooler, 1991; Averell & Heathcote, 2011; Donkin & Nosofsky, 2012b; Rubin & Wenzel, 1996; Wickelgren, 1974; Wixted & Ebbesen, 1991). When considering this contrast, one important point is that previous studies have often specified relationships between the study-test lag and a particular measurement variable such as proportion correct or mean RT, and thus may not generalize to different measurement variables or transformations of these variables. Donkin and Nosofsky (2012b) argued that a solution to this problem is to instead establish relationships between study-test lag and latent memory strength. In their model, they compared exponential and power functions of memory strength which were used in an LBA model to produce proportions of correct responses and RT distributions for each study-test lag. The results of their modeling found strong evidence for a power function of memory strength using 12 item lists in a single probe item recognition task.

Our investigation similarly conducted the comparison on the basis of a latent memory strength variable through models that considered both response proportions and complete RT distributions using list lengths that ranged from 10 to 24 items, which are within the range of what was tested by Donkin and Nosofsky. Thus, it was somewhat surprising to find evidence for an exponential function of recency. However, our finding is not unprecedented, as Wixted and Ebbesen (1997) found support for an exponential retention function in their analysis of complete sequences of free recall, while finding evidence for power functions in other tasks such as recognition memory. In addition, there are theoretical reasons to expect an exponential function: current free recall models such as TCM (Howard & Kahana, 2002) implicitly use exponential functions, as the process of contextual drift produces exponential decay where the rate of contextual drift is proportional to the decay rate in an exponential function (Howard, 2004).

This prompts the question: why would free recall exhibit a recency function different to that seen in other tasks such as recognition memory? The answer is not immediately obvious, and goes considerably beyond the scope of the present work, but one possible explanation comes from a consideration of how power functions might be generated. Despite the prevalence of evidence for power functions in memory tasks, generation of a power function within a process model is non-trivial, as it is not immediately obvious how a memory's rate of decay should slow as time progresses. Instead, it is more common to assume that power functions arise from a sum or weighted average of exponential functions with different decay rates (e.g., R. B. Anderson & Treeney, 1997; Sikström, 2002). Using this line of reasoning, Howard, Shankar, Aue, and Criss (2015) showed how a power function of memory strength could be generated by using a large number of temporal context vectors which all show exponential decay but at different rates; these are subsequently subjected to an inverse Laplace transform to generate a scale invariant representation of time. Howard et al. (2015) argued that PFR curves provide strong support for their model, but did not perform a model comparison with a simple exponential model.

The Howard et al. (2015) model and other models that rely on multiple exponential scales may be able to provide stronger accounts of PFR curves by assuming that participants are able to “target” specific timescales that are relevant to the study list. One consequence of focussing on a particular time scale would be that the memory strength function would be more exponential in nature. Most free recall experiments are on the scale of seconds or less than a minute, so that tracking longer time scales may not only be unnecessary, but detrimental as participants have to avoid recalling items from prior lists. It is arguably advantageous to restrict cuing to scales that are specific to the most recent study list and avoid longer timescales which subsume prior study lists. Evidence for targeting of relevant timescales in memory retrieval comes from Moreton and Ward (2010), in which participants were asked to recall autobiographical memories from the scale of

either weeks, months, or years. Participants were not only successful in retrieving events from each of these scales, but also exhibited very similar numbers of items recalled, similar patterns of output interference, and similar patterns of contiguity. The Howard et al. (2015) model represents one way in which power or exponential functions might be observed in different tasks or conditions.

Evidence for Reinstatement as a Mechanism of Primacy

In our comparison of mechanisms that generate the primacy effect, the evidence strongly favored a reinstatement approach to primacy over a strength or rehearsal account. A strength account of primacy assumes that primacy items receive extra strength than the items that follow them, but that primacy and recency items jointly compete to be retrieved. The rehearsal account assumes that primacy items are also stronger, but that on some occasions the first item behaves in a similar manner as the recency items due to its rehearsal through to the end of the list. The reinstatement model assumes that participants are able to cue with the beginning of the list instead of the end of the list, implying that on some occasions the primacy items race without any competition from the recency items.

While each of these models appears to make similar predictions about PFR curves, the models make different predictions about how RT distributions vary across serial position. This is due to a counter-intuitive prediction from evidence accumulation models that the RT distributions for each serial position is strongly determined by the fastest competitor in the race, meaning that RTs can be very similar for each serial position despite very different recall rates. We have demonstrated that this prediction applies to both the racing diffusion model and the LBA, and in Supplementary Materials B we demonstrate in simulations that this also applies to the LCA model across a range of values of lateral inhibition and leakage that encompass those used in previous applications of TCM-A and CMR. Because of this property, the strength model is considerably more constrained than the other classes of models, such as the rehearsal model which assumes a

mixture of a fast race of the recency items and an even faster race where the primacy item is the strongest, and the reinstatement model which assumes a mixture of a fast recency race and a slower primacy race.

The model selection results strongly favored the reinstatement model, in large part because the RT distributions exhibited a pattern that was in accordance with the qualitative predictions of that model: namely, slower latencies for the first item than the subsequent items, a finding first reported by Laming (1999) in a reanalysis of the Murdock and Okada (1970) dataset. This is interesting when one considers that the first item exhibits a recall advantage over its successors and better recall is very often associated with faster latencies (Kahana, 1996; Madigan, Neuse, & Roeber, 2000; Nobel & Shiffrin, 2001; Sederberg et al., 2008; Wixted & Rohrer, 1993). This finding is easily captured by the reinstatement model even when there is a very strong recency effect; provided the strongest item in the primacy gradient is weaker than the strongest item in the recency gradient, the primacy race will generally produce slower RTs than the recency race. Most of the recalls of the recency items will stem from the fast recency race while most of the recalls of the primacy item will stem from the slower primacy race, producing relatively slow recalls of the primacy item.

The support for the reinstatement model is intriguing when one considers the prevalence of strength models in the literature. A number of dual store models assume that the primacy item receives extra strength due to additional time in a capacity limited buffer (Atkinson & Shiffrin, 1968; Lehman & Malmberg, 2013; Raaijmakers & Shiffrin, 1980). Temporal distinctiveness models assume that the first item receives additional strength due to an “edge” effect in the temporal representation (G. D. A. Brown et al., 2007). Models that employ prediction-error encoding, whereby each item’s encoding strength is proportional to its similarity to the contents of memory, produce the strongest encoding for the first item due to its novelty (Davelaar, 2007; Elhalal et al., 2014; Farrell & Lewandowsky, 2002).

While uncommitted to the underlying mechanism, TCM-A and CMR are strength models in that they apply an exponential memory strength boost for the early list items. A recent investigation attempting to explain individual differences in PFR curves with CMR concluded that variability across individuals reflects parameters of a single strength model rather than variability in memory search strategies or cues (Healey & Kahana, 2014). However, their investigation did not consider the constraints imposed by RT distributions. The support we found for the reinstatement model—which is only obvious when considering the full RT distribution—is not consistent with their conclusions, and suggests that participants employ different retrieval cues to access items at the start versus the end of the list. This is consistent with prior evidence which has found that participants can initiate free recall at different points in the list when instructed to (Dalezman, 1976), and that participants can initiate free recall or serial recall when prompted to after having encoded the list (Bhatarah, Ward, & Tan, 2008). Such findings are problematic for a strength model, which has to assume that a dominance of primacy is due to the fact that the primacy items are stronger than the recency items. In contrast, a reinstatement model can alter the effective memory strength of the early list items at the beginning of recall by using a start-of-list cue. The current findings suggest that if TCM-A and CMR were to be extended to complete RT distributions, reinstatement of the beginning of the list context may be required to address the differences in RT across serial positions. Such a process is not incompatible with the frameworks of these models, indeed some variants of CMR have employed such a process (Kragel, Morton, & Polyn, 2015; Morton & Polyn, 2016).

One example of a computational model that selectively focusses on primacy or recency items at recall initiation is the grouping model of Farrell (2012), which provides a comprehensive account of both free recall and serial recall by implementing reinstatement. In the grouping model, study list items are sequentially clustered into groups that vary in size. A strong recency effect in free recall can be generated because participants already have access to the most recent group, and so will preferentially recall that group when

given the option. Primacy effects in free recall and initiation of serial recall are both assumed to be the product of reinstating the context of the first group and using that to initiate recall, which will generate high memory strength for items at the beginning of the list and lower memory strength for items at the end of the list. Farrell (2012) showed that his model could account for some aspects of latency data in serial and free recall (e.g., the exponential shape of the cumulative recall function), and an interesting question for future research is whether the model could account for the full pattern of RTs observed here.

The class of rehearsal models generally performed quite poorly in addressing the RT distributions in each of the datasets. How is this finding reconciled with the fact that recall can often be predicted on the basis of participants' rehearsals when rehearsal data are recorded (Laming, 2006, 2008)? There are two considerations here. The first is that the overt rehearsal method often employs slow rates of presentation, with as much as five seconds per item, which both increases rehearsals and the extent to which the recency of rehearsals influences the resulting serial position curve (Brodie, 1975; Brodie & Murdock, 1977; Tan & Ward, 2000). Similarly, Tan and Ward (2008) compared rehearsal schedules data in immediate serial recall across presentation rates and found it was only under long presentation rates that participants engaged in a rehearsal schedule where items prior to the currently presented item were rehearsed; as presentation rate became faster, participants increasingly adopted a strategy where only the most recently presented item was rehearsed. Presentation rates for our datasets are described in Supplementary Materials A. While some of the PEERS datasets employed slower rates of presentation (Datasets 1–4; 3,000 ms per item with a jittered 800–1,200 ms blank interstimulus interval), these were also often accompanied by rating tasks which should reduce the participant's ability to rehearse the items. Several of the other datasets employ more rapid rates of presentation (1,000–1,200 ms per item), making it unlikely for participants to engage in cumulative rehearsal strategies.

Some additional findings from the literature have also demonstrated insufficiencies

with a primacy-as-recency account of primacy effects. Perhaps the most challenging aspect is the fact that increases in study-test delay attenuate the recency effect but have little effect on the primacy effect (Tan & Ward, 2000). In addition, attempts to restrict rehearsal to the just presented item reduce but do not eliminate primacy (Fischler, Rundus, & Atkinson, 1970; Glanzer & Meinzer, 1967). Finally, the ability to predict recall on the basis of rehearsals does not specify the mechanism by which rehearsal influences recall. For instance, it is possible that rehearsal of the primacy items enhances recall not by making them functionally recent, but instead further strengthens their associations to the beginning-of-list context that is reinstated in reinstatement models.

At first glance, it may appear that the differences in latency across serial positions might be well explained by dual-store models. If retrieval from the short term store (STS) is faster than retrieval from the long term store (LTS), this would produce relatively fast RTs for the recency items (which remain in STS) and slower RTs for the earlier items (which are in the LTS). However, there are two potential issues with this account. First, the latency differences across serial position were still found in Datasets 3 (delayed free recall) and 4 (continual distractor free recall). In both datasets recall was delayed by a demanding distractor task which should clear the contents of the STS. Second, this account requires that recall sometimes initiates from the LTS instead of the STS, whereas traditional accounts state that retrieval from the LTS proceeds only after the STS has been cleared (Davelaar et al., 2005; Raaijmakers & Shiffrin, 1981). In these accounts, first retrievals should only stem from the STS, providing no explanation as to why early items are retrieved slower.

Extension to Complete Free Recall Sequences

The present work addressed first responses in participants' free recall sequences. Here, we elaborate on how the models presented here could be extended to complete sequences of recalls. As mentioned previously, a challenge for such a modeling effort will be

addressing the strong sequential dependencies in participants' free recall sequences, whereby participants are likely to follow recall of an item with an item that was studied in close proximity to the previously recalled item (Kahana, 1996): an effect termed the *contiguity effect*. The contiguity effect is measured by estimating conditional response probabilities for each lag; these are the conditional probabilities of retrieving an item with a positional lag from the just-recalled item. CRP curves are noted for their asymmetry; participants are more likely to proceed in the forward direction from the just recalled item than the backward direction. The contiguity effect is also manifest in RTs: retrievals with short lags are fast while longer lags are often considerably slower (Kahana, 1996; Sederberg, Miller, Howard, & Kahana, 2010).

Addressing the constraints imposed by the contiguity effect in conjunction with complete RT distributions may be possible by using a front-end model such as SAM (Raaijmakers & Shiffrin, 1981), TCM (Howard & Kahana, 2002), or the grouping model (Farrell, 2012). In each of these models, the sequential effects are addressed by updating the cues after each recall. For instance, in TCM the retrieved item and its retrieved context representation are combined to form a retrieval cue; the retrieved item produces a forward-order asymmetry in the CRP while the retrieved context produces a symmetric component (Howard & Kahana, 2002). While the model conventionally assumes these cues are additively combined after the retrieval of each item, participants might instead only rely on one cue or the other, potentially because participants are unable to retrieve context on some outputs. Similar to our discussion of the primacy models in this text, these different conceptions of how retrieval cues are employed following each recall may make quite distinct predictions about the shapes of RT distributions at each lag. Specifically, the conventional model which assumes the two cues are combined will predict approximately similar RT for each lag, whereas a probability mixture of retrieved items and retrieved contexts is less constrained and has greater ability to produce variability in the latencies across each retrieval lag.

A challenge for extending the present model to full sequences is output interference. Murdock and Okada (1970) observed that RTs increase exponentially across output position. In the CMR model, this pattern was captured by assuming that when an already recalled item hits its threshold, it resets accumulation for all items. In addition, the lateral inhibition present in the LCA model (which is CMR’s back-end) makes it such that recalled items continue to inhibit other items in the decision stage (Polyn et al., 2009). Farrell’s (2012) grouping model produces this slowing by implementing output interference between successively recalled groups. While these are plausible mechanisms, they are not tractable analytically. An alternative to these mechanisms would be to assume that retrieval causes feed-forward inhibition (e.g., Shadlen & Newsome, 2001), which would decrement the drift rates of each of the other accumulators. A mechanism that could instantiate this inhibition is as follows: if each retrieved item receives an increase in strength (as is done in the SAM model; Raaijmakers & Shiffrin, 1981) and the drift rates are normalized to sum to a constant value, the increase in strength for retrieved items would decrease the strength for the other list items. This mechanism would preserve analytic tractability and is compatible with the current models.

A final consideration would be the decision to terminate recall, which can have an explicit latency when participants are given the option to terminate when they feel they can no longer recall additional items (Dougherty & Harbison, 2007; Miller et al., 2012). This could be instantiated by having a “quit” accumulator that races with the other to-be-retrieved items and causes a termination when it hits its threshold (e.g., Moran, Zehetleitner, Liesefeld, Müller, & Usher, 2016). We did not have to consider such an approach here because it was extremely rare for participants to be unable to recall one item in the datasets we considered.

Conclusions

In this work, we presented the first application of evidence accumulation models to complete RT distributions and first recall serial position curves. In addition to providing support for current models of free recall that use exponential recency functions, the evidence accumulation models were able to adjudicate between three different mechanisms of the primacy effect on the basis of their different predictions about how RT distributions vary across serial position. This work demonstrates the value of considering RT distributions in assessing models of free recall.

References

- Anderson, J., & Matessa, M. (1997). A production system theory of serial memory. *Psychological Review*, 104.
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2(6), 396–408.
- Anderson, R. B., & Tweeney, R. D. (1997). Artifactual power curves in forgetting. *Memory & Cognition*, 25, 724–730.
- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. *Psychology of Learning and Motivation*, 2, 89–195.
- Averell, L., & Heathcote, A. (2011). The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology*, 55(1), 25–35.
- Bhatarah, P., Ward, G., & Tan, L. (2008). Examining the relationship between free recall and immediate serial recall: The serial nature of recall and the effect of test expectancy. *Memory & Cognition*, 36(1), 20–34.
- Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology*, 6, 173–189.
- Brodie, D. A. (1975). Free recall measures of short-term store: Are rehearsal and order of recall data necessary? *Memory & Cognition*, 3, 653–662.
- Brodie, D. A., & Murdock, B. B. (1977). Effect of presentation time on nominal and functional serial-position curves of free recall. *Journal of Verbal Learning and Verbal Behavior*, 16, 185–200.
- Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, 114, 539–576.
- Brown, G. D. A., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, 107, 127–181.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57, 153–178.

- Bundesen, C. (1993). The relationship between independent race models and Luce's choice axiom. *Journal of Mathematical Psychology*, 37, 446–471.
- Criss, A. H. (2010). Differentiation and response bias in episodic memory: Evidence from reaction time distributions. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 36(2), 484–499.
- Dalezman, J. J. (1976). Effects of output order on immediate, delayed, and final recall performance. *Journal of Experimental Psychology: Human Learning and Memory*, 2(5), 597–608.
- Davelaar, E. J. (2007). Sequential retrieval and inhibition of parallel (re)activated representations: A neurocomputational comparison of competitive queuing and resampling models. *Adaptive Behavior*, 15(1), 51–71.
- Davelaar, E. J., Goshen-Gottstein, Y., Ashkenazi, A., & Haarmann. (2005). The Demise of Short-Term Memory Revisited: Empirical and Computational Investigations of Recency Effect. *Psychological Review*, 112, 3–42.
- Donkin, C., Brown, S. D., & Heathcote, A. (2009). The overconstraint of response time models: Rethinking the scaling problem. *Psychonomic Bulletin & Review*, 16(6), 1129–1135.
- Donkin, C., & Nosofsky, R. M. (2012a). A power-law model of psychological memory strength in short- and long-term recognition. *Psychological Science*, 23(6), 625–634.
- Donkin, C., & Nosofsky, R. M. (2012b). The structure of short-term memory scanning: an investigation using response time distribution models. *Psychonomic Bulletin & Review*, 19, 363–394.
- Donkin, C., Nosofsky, R. M., Gold, J., & Shiffrin, R. M. (2013). Discrete-Slots Models of Visual Working-Memory Response Times. *Psychological Review*, 120, 873–902.
- Dougherty, M. R. P., & Harbison, J. (2007). Motivated to retrieve: How often are you willing to go back to the well when the well is dry? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 1108–1117.

- Elhalal, A., Davelaar, E. J., & Usher, M. (2014). The role of the frontal cortex in memory: an investigation of the Von Restorff effect. *Frontiers in Human Neuroscience*, 8, 1–20.
- Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, 62, 145–154.
- Farrell, S. (2006). Mixed-list phonological similarity effects in delayed serial recall. *Journal of Memory and Language*, 55, 587–600.
- Farrell, S. (2012). Temporal clustering and sequencing in short-term memory and episodic memory. *Psychological Review*, 119(2), 223–271.
- Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin and Review*, 9, 59–79.
- Fific, M., Little, D. R., & Nosofsky, R. M. (2010). Logical-rule models of classification response times: A synthesis of mental-architecture, random-walk, and decision-bound approaches. *Psychological Review*, 117(2), 309–348.
- Fischler, I., Rundus, D., & Atkinson, R. C. (1970). The effects of intralist activity on free recall. *Psychonomic Science*, 19, 249–250.
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91(1), 1–67.
- Glanzer, M., & Cunitz, A. R. (1966). Two storage mechanisms in free recall. *Journal of Verbal Learning and Verbal Behavior*, 5(4), 351–360.
- Glanzer, M., & Meinzer, A. (1967). The effects of intralist activity on free recall. *Journal of Verbal Learning and Verbal Behavior*, 6, 928–935.
- Golomb, J. D., Peelle, J. E., Addis, K. M., Kahana, M. J., & Wingfield, A. (2008). Effects of adult aging on utilization of temporal and semantic associations during free and serial recall. *Memory & Cognition*, 36(5), 947–956.
- Gronlund, S. D., & Ratcliff, R. (1991). Analysis of the Hockley and Murdock decision model. *Journal of Mathematical Psychology*, 35(3), 319–344.

- Grossberg, S., & Pearson, L. R. (2008). Laminar cortical dynamics of cognitive and motor working memory, sequence learning and performance: Toward a unified theory of how the cerebral cortex works. *Psychological Review*, 115(3), 677–732.
- Healey, M. K., & Kahana, M. J. (2014). Is memory search governed by universal principles or idiosyncratic strategies? *Journal of Experimental Psychology: General*, 143, 575–596.
- Healey, M. K., & Kahana, M. J. (2016). A four-component model of age-related memory change. *Psychological Review*, 123, 23–69.
- Henson, R. N. A. (1998). Short-term memory for serial order: The start-end model. *Cognitive Psychology*, 36, 73–137.
- Hogan, R. M. (1975). Interitem encoding and directed search in free recall. *Memory & Cognition*, 3, 197–209.
- Howard, M. W. (2004). Scaling behavior in the temporal context model. *Journal of Mathematical Psychology*, 48(4), 230–238.
- Howard, M. W., & Kahana, M. J. (1999). Contextual Variability and Serial Position Effects in Free Recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(4), 923–941.
- Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46, 268–299.
- Howard, M. W., Shankar, K. H., Aue, W. R., & Criss, A. H. (2015). A distributed representation of internal time. *Psychological Review*, 122(1), 24–53.
- Howard, M. W., Venkatadass, V., Norman, K. A., & Kahana, M. J. (2007). Associative processes in immediate recency. *Memory & Cognition*, 35, 1700–1711.
- Jost, A. (1897). Die Assoziationsfestigkeit in ihrer Abhängigkeit von der Verteilung der Wiederholungen [The strength of associations in their dependence on the distribution of repetitions]. *Z. Psychol. Physiol. Sinnesorgane*, 16, 436–472.
- Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory and*

- Cognition*, 24, 103–109.
- Kahana, M. J., Aggarwal, E. V., & Phan, T. D. (2018). The variability puzzle in human memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*.
- Kahana, M. J., Howard, M. W., Zaromb, F., & Wingfield, A. (2002). Age dissociates recency and lag recency effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(3), 530–540.
- Kragel, J. E., Morton, N. W., & Polyn, S. M. (2015). Neural activity in the medial temporal lobe reveals the fidelity of mental time travel. *Journal of Neuroscience*, 35, 2914–2926.
- Laming, D. (1999). Testing the idea of distinct storage mechanisms in memory. *International Journal of Psychology*, 34, 419–426.
- Laming, D. (2006). Predicting free recalls. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 32, 1146–1163.
- Laming, D. (2008). An improved algorithm for predicting free recalls. *Cognitive Psychology*, 57, 179–219.
- Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, 55, 1–7.
- Lehman, M., & Malmberg, K. J. (2013). A buffer model of memory encoding and temporal correlations in retrieval. *Psychological Review*, 120(1), 155–189.
- Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review*, 96, 25–57.
- Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review*, 121, 66–95.
- Lohnas, L. J., & Kahana, M. J. (2014). Compound cuing in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(1), 12–24.
- 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.
- Madigan, S., Neuse, J., & Roeber, U. (2000). Retrieval latency and "at-risk" memories. *Memory & Cognition*, 28(4), 523–528.
- Mensink, G. J., & Raaijmakers, J. G. W. (1988). A model for interference and forgetting. *Psychological Review*, 95(4), 434–455.
- Metcalfe, J., & Murdock, B. B. (1981). An encoding and retrieval model of single-trial free recall. *Journal of Verbal Learning and Verbal Behavior*, 20, 161–189.
- Miletic, S., Turner, B. M., Forstmann, B. U., & van Maanen, L. (2017). Parameter recovery for the Leaky Competing Accumulator model. *Journal of Mathematical Psychology*, 76, 25–50.
- Miller, J. F., Weidemann, C. T., & Kahana, M. J. (2012). Recall termination in free recall. *Memory & Cognition*, 40, 540–550.
- Moran, R., Zehetleitner, M., Liesefeld, H. R., Müller, H. J., & Usher, M. (2016). Serial vs. parallel models of attention in visual search: accounting for benchmark RT-distributions. *Psychonomic Bulletin & Review*, 23, 1300–1315.
- Moreton, B. J., & Ward, G. (2010). Time scale similarity and long-term memory for autobiographical events. *Psychonomic Bulletin & Review*, 17, 510–515.
- Morton, N. W., & Polyn, S. M. (2016). A predictive framework for evaluating models of semantic organization in free recall. *Journal of Memory and Language*, 86, 119–140.
- Murdock, B. B., & Okada, R. (1970). Interresponse times in single-trial free recall. *Journal of Experimental Psychology*, 86(2), 263–267.
- Nobel, P. A., & Shiffrin, R. M. (2001). Retrieval processes in recognition and cued recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(2), 384–413.
- Osth, A. F., Bora, B., Dennis, S., & Heathcote, A. (2017). Diffusion versus 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), 260–311.
- Osth, A. F., Dennis, S., & Heathcote, A. (2017). Likelihood ratio sequential sampling models of recognition memory. *Cognitive Psychology*, 92, 101–126.
- 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.
- Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, 116(1), 129–156.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 14). Academic Press.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of Associative Memory. *Psychological Review*, 88(2), 93–134.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59–108.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20, 873–922.
- Ratcliff, R., & Murdock, B. B. (1976). Retrieval Processes in Recognition Memory. *Psychological Review*, 83(3), 190–214.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9(5), 347–356.
- Ratcliff, R., Sheu, C. F., & Gronlund, S. D. (1992). Testing global memory models using ROC curves. *Psychological Review*, 99(3), 518–535.
- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential-sampling models for two choice reaction time. *Psychological Review*, 111, 333–367.

- Ratcliff, R., & Starns, J. J. (2009). Modeling confidence and response time in recognition memory. *Psychological Review*, 116(1), 59–83.
- Ratcliff, R., Thapar, A., & McKoon, G. (2004). A diffusion model analysis of the effects of aging on recognition memory. *Journal of Memory and Language*, 50, 408–424.
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9, 438–481.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, 106(2), 261–300.
- Rohrer, D., & Wixted, J. T. (1994). An analysis of latency and interresponse time in free recall. *Memory & Cognition*, 22, 511–524.
- Rouder, J. N., & Lu, J. (2005). An introduction to Bayesian hierarchical models with an application in the theory of signal detection. *Psychonomic Bulletin & Review*, 12(4), 573–604.
- Rubin, D. C., & Wenzel, A. E. (1996). One hundred years of forgetting: A quantitative description of retention. *Psychological Review*, 103(4), 734–760.
- Rundus, D. (1971). An analysis of rehearsal processes in free recall. *Journal of Experimental Psychology*, 89, 63–77.
- Sederberg, P. B., Howard, M. W., & Kahana, M. J. (2008). A context-based theory of recency and contiguity in free recall. *Psychological Review*, 115(4), 893–912.
- Sederberg, P. B., Miller, J. F., Howard, M. W., & Kahana, M. J. (2010). The temporal contiguity effect predicts episodic memory performance. *Memory & Cognition*, 38(6), 689–699.
- Serruya, M. D., Sederberg, P. B., & Kahana, M. J. (2014). Power shifts track serial position and modulate encoding in human episodic memory. *Cerebral Cortex*, 24, 403–413.
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex. *Journal of Neurophysiology*, 86, 1916–1936.

- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM - retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4(2), 145–166.
- Sikström, S. (2002). Forgetting curves: implications for connectionist models. *Cognitive Psychology*, 45(1), 95–152.
- Starns, J. J., Ratcliff, R., & McKoon, G. (2012). Evaluating the unequal-variance and dual-process explanations of the zROC slopes with response time data and the diffusion model. *Cognitive Psychology*, 64, 1–34.
- Tan, L., & Ward, G. (2000). A recency-based account of the primacy effect in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(6), 1589–1625.
- Tan, L., & Ward, G. (2008). Rehearsal in immediate serial recall. *Psychonomic Bulletin & Review*, 15, 535–542.
- Tillman, G., & Logan, G. D. (2017). The racing diffusion model of speeded decision making. *Manuscript under revision*.
- 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.
- Unsworth, N., & Engle, R. (2007). The nature of individual differences in working memory capacity: Active maintenance in primary memory and controlled search from secondary memory. *Psychological Review*, 114(1), 104–132.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108, 550–592.
- Usher, M., Olami, Z., & McClelland, J. L. (2002). Hick’s Law in a stochastic race model with speed-accuracy tradeoff. *Journal of Mathematical Psychology*, 46, 704–715.
- Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (2011). Hierarchical diffusion models for two-choice response times. *Psychological Methods*, 16(1), 44–62.
- Ward, G., Tan, L., & Grenfell-Essam, R. (2010). Examining the Relationship Between Free

- Recall and Immediate Serial Recall: The Effects of List Length and Output Order. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(5), 1207–1241.
- Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *The Journal of Machine Learning Research*, 11, 3571–3594.
- Wickelgren, W. A. (1974). Single-trace fragility theory of memory dynamics. *Memory & Cognition*, 2, 775–780.
- Wixted, J. T. (2004). The psychology and neuroscience of forgetting. *Annual Review of Psychology*, 55, 235–269.
- Wixted, J. T., & Ebbesen, E. B. (1991). On the form of forgetting. *Psychological Science*, 2(6), 409–415.
- Wixted, J. T., & Ebbesen, E. B. (1997). Genuine power curves in forgetting: A quantitative analysis of individual subject forgetting functions. *Memory & Cognition*, 25, 731–739.
- Wixted, J. T., & McDowell, J. J. (1989). Contributions to the functional analysis of single-trial free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(4), 685–697.
- Wixted, J. T., & Rohrer, D. (1993). Proactive interference and the dynamics of free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(5), 1024–1039.
- Wixted, J. T., & Rohrer, D. (1994). Analyzing the dynamics of free recall: An integrative review of the empirical literature. *Psychonomic Bulletin & Review*, 1(1), 89–106.

Appendix

Additional Details of the Fitting Procedure

Varying the s Parameter Across Accumulators

A challenge we found in applying such models with large number of accumulators to data is that the models often had difficulty producing recall probabilities at floor even if the drift rate for a given accumulator was zero, preventing good fits to data in conditions where participants never recalled a particular item. In the racing diffusion model, the model predicts recall probabilities above zero when drift rate is zero because noise alone can drive the accumulator to the threshold. In the LBA, the same phenomenon can occur via cross-trial variability in drift rate; it is possible that a high drift rate is sampled from the drift rate distribution even if the mean drift rate is zero. To prevent accumulators with zero drift from hitting the threshold, we scale the noise parameter s parameter to be proportional to the drift rate, using the following linear function:

$$s_i = s_{min} + [(1 - s_{min}) \frac{v_i}{max(v)}] \quad (5)$$

where s_{min} is a parameter that varies between 0 and 1 and $max(v)$ is the highest drift rate across all accumulators and conditions. This ensures that s is fixed to 1 for the accumulator with the highest drift rate, an accumulator with a drift rate of zero has $s = s_{min}$, and all other accumulators have values of s that are between those two values. This fixes scaling for each model as the s parameter only has to be fixed for one accumulator for the models to be identifiable (Donkin et al., 2009), and also allowed each model to scale well to large numbers of accumulators. This assumption is also consistent with several models of episodic memory in which the variability in memory strength is often proportional to the mean (Gillund & Shiffrin, 1984; Ratcliff, Sheu, & Gronlund, 1992; Shiffrin & Steyvers, 1997; Osth & Dennis, 2015). If $s_{min} = 1$, all accumulators have $s = 1$.

Data Exclusions

We excluded all trials where participants failed to recall an item, committed an intrusion error, or where RTs for the first response were over 20 seconds. We additionally excluded two participants each from Datasets 6 and 7 for having mean RTs below 200 ms, which in some cases prevented convergence of model parameters. These criteria resulted in the exclusion of 6.0% and 6.4% of the data from Datasets 6 and 7, whereas for all other datasets it resulted in the exclusion than less of 1% of the data.

Prior Distributions on Model Parameters

Participant parameters are sampled from group level mean and standard deviation parameters M and ς . Several of the bounded parameters were sampled from truncated normal distributions. Unbounded parameters were sampled from normal distributions. The memory strength parameters of the primacy and recency functions were sampled from lognormal distributions:

$$B \sim TN(M_B, \varsigma_B, 0, \infty) \quad (6)$$

$$A \sim TN(M_A, \varsigma_A, 0, \infty) \quad (7)$$

$$\log(r_s) \sim Normal(M_{rs}, \varsigma_{rs}) \quad (8)$$

$$\log(r_d) \sim Normal(M_{rd}, \varsigma_{rd}) \quad (9)$$

$$\log(a) \sim Normal(M_a, \varsigma_a) \quad (10)$$

$$\log(\beta) \sim Normal(M_b, \varsigma_b) \quad (11)$$

$$v_r \sim TN(M_{vr}, \varsigma_{vr}, 0, \infty) \quad (12)$$

Priors on the group level distributions for all parameters in each model were vague. This was partly motivated by the fact that these models have not been applied to the free recall task previously. Priors were defined as:

$$M_A \sim TN(2, 2, 0, \infty) \quad (13)$$

$$M_B \sim TN(3, 3, 0, \infty) \quad (14)$$

$$M_{rs}, M_{rd}, M_a, M_b \sim Normal(0, 10) \quad (15)$$

$$\varsigma_B \sim Gamma(1, 1) \quad (16)$$

$$\varsigma_A, \varsigma_{rs}, \varsigma_{rd}, \varsigma_a, \varsigma_b, \varsigma_{vr} \sim Gamma(1, 3) \quad (17)$$

As mentioned in the main text, we made two allowances for the non-decision time. One class of models we referred to as *cued accumulation* models assumed that evidence accumulation could not begin until after participants were cued to recall the list. These models assume a positive value of non-decision time t_0 which is added to the RT for each decision. Because t_0 is always positive in these models, we specified that it was sampled from a truncated normal distribution:

$$t_0 \sim TN(M_{t_0}, \varsigma_{t_0}, 0, \infty) \quad (18)$$

$$M_{t_0} \sim TN(.5, .5, 0, \infty). \quad (19)$$

Another class of models we referred to as *preemptive accumulation* models assumed that evidence accumulation can begin prior to the recall cue. In these models, we allowed nondecision time t_0 to go negative. Because t_0 is unbounded, the parameter is sampled from a normal distribution:

$$t_0 \sim Normal(M_{t_0}, \varsigma_{t_0}) \quad (20)$$

$$M_{t_0} \sim Normal(.5, .5) \quad (21)$$

Both of these cases use the same prior on ς :

$$\varsigma_{t0} \sim \text{Gamma}(1, 3) \quad (22)$$

Parameters that were bounded between 0 and 1 (s_{min} and p) were sampled from a beta distribution. Improved sampling for the beta distribution can be obtained by parameterizing it in terms of its mean μ and sample size v :

$$\alpha = \mu v \quad (23)$$

$$\beta = (1 - \mu)v \quad (24)$$

A completely flat prior for a beta distribution can be obtained by setting $\mu = .5$ and $v = 2$. This was done for the group means of both parameters:

$$s_{min} \sim \text{Beta}(M_{smin}, v_{smin}) \quad (25)$$

$$p \sim \text{Beta}(M_p, v_p) \quad (26)$$

$$M_{smin}, M_p \sim \text{Beta}(.5, 2) \quad (27)$$

$$v_{smin}, v_p \sim \text{Gamma}(1, 3) \quad (28)$$

For the data from Murdock and Okada (1970), additional drift rate parameters v_d that were sampled from normal distributions were added to the last two positions of the recency function to account for the nonmonotonic PFR curve:

$$v_d \sim \text{Normal}(M_{vd}, \varsigma_{vd}) \quad (29)$$

$$M_{vd} \sim \text{Normal}(.5, 2) \quad (30)$$

$$\varsigma_{vd} \sim \text{Gamma}(1, 3) \quad (31)$$

Details on MCMC Estimation Procedure

For each dataset, the number of chains was set equal to five times the number of participant parameters for each dataset except for Dataset 5, where the very large size of the data combined with the large number of accumulators prevented using such a large number of chains. Instead, for Dataset 5 the number of chains was set equal to three times the number of participant parameters. Chains were heavily thinned to reduce autocorrelation such that one in every 20 MCMC iterations was recorded; this process continued until 1,500 samples were collected. This occurred after 25,000 burn-in iterations were discarded. For Dataset 3 (Howard, Venkatadass, Norman, and Kahana (2007)), the very large number of participants required a lengthy burn-in period of 80,000 iterations to achieve convergence of the hyperparameters.

A model was considered converged if its Gelman-Rubin (G-R) statistic was below 1.20 for all parameters. For some cases, such as Dataset 1 for the LBA models, this standard was not possible even with much larger number of burn-in iterations and collected samples as there were some individual participants with poorly mixing chains. Thus, subsequent to each model fit, we removed chains with median absolute deviations that were 3.5 times greater than the overall median absolute deviation, or 3.5 times less than the overall median absolute deviation, on any of the participant parameters. The removal of such chains brought the G-R statistic lower than 1.2 for all models, with values quite close to 1.0 in the majority of the cases.

Posterior Predictive Generation

For each posterior sample a total of 15,000 simulations from the model were generated to achieve stable estimates of the predicted RT distributions. This was done for one in every 50 MCMC iterations from the model.

Details of the Hierarchical Ex-Gaussian Procedure

The ex-gaussian model consists of three parameters, the two parameters of the normal distribution (μ and σ) and the τ parameter of the exponential distribution. Because the σ and τ parameters have a lower bound of zero, these parameters were sampled on a log scale. Each parameter was sampled from the following group distributions:

$$\mu \sim Normal(M_\mu, \varsigma_\mu) \quad (32)$$

$$\log(\sigma) \sim Normal(M_\sigma, \varsigma_\sigma) \quad (33)$$

$$\log(\tau) \sim Normal(M_\tau, \varsigma_\tau) \quad (34)$$

Non-informative priors were placed on the group level distributions:

$$M_\mu \sim Normal(.5, 2.5) \quad (35)$$

$$M_\sigma, M_{tau} \sim Normal(0, 10) \quad (36)$$

$$\varsigma_\mu, \varsigma_\sigma, \varsigma_\tau \sim Gamma(1, 1) \quad (37)$$

A total of six ex-gaussian distributions were applied to each dataset, with parameters corresponding to the first serial position, serial positions 2 through $L - 5$ (where L is the length of the list), along with serial positions $L - 4$, $L - 3$, $L - 2$, and L .

Each model used three times the number of chains as the number of participant

parameters. After 20,000 burn-in iterations were discarded, one sample in every 20 iterations was recorded until 1,500 samples were collected for each chain. For each model, the G-R statistic was below 1.05 for each parameter for each participant, indicating a strong degree of model convergence.

For each participant, 15,000 simulations were generated for each ex-gaussian distribution. For each simulated distribution, the 10th, 50th, and 90th percentile of the distribution was assessed. The quantiles were subsequently averaged across all of the posterior samples for each participant.