



## Original Articles

## Two distinct mechanisms of selection in working memory: Additive last-item and retro-cue benefits



Marcel Niklaus, Henrik Singmann, Klaus Oberauer\*

University of Zurich, Switzerland

## ARTICLE INFO

## Keywords:

Working memory  
Attention  
Retro-cue  
Recency  
Focus of attention  
Speed-accuracy trade-off  
Hierarchical-Bayes

## ABSTRACT

In working memory research, individual items are sometimes said to be in the “focus of attention”. According to one view, this occurs for the last item in a sequentially presented list (last-item benefit). According to a second view, this occurs when items are externally cued during the retention interval (retro-cue benefit). We investigated both phenomena at the same time to determine whether both result from the same cognitive mechanisms. If that were the case, retro-cue benefits should be reduced when the retro-cue is directed to the item that already benefits from being presented last. We measured speed-accuracy-tradeoff functions with the response-deadline paradigm to measure retrieval dynamics in a short-term recognition task. Across three experiments, we found that retro-cues benefited the last item and other items to the same extent. The additivity of the last-item benefit and the retro-cue benefit points towards the co-existence of at least two distinct forms of attentional prioritization in working memory.

## 1. Introduction

Working memory is a system devoted to the selective maintenance of information in a highly accessible state to support cognitive activities such as reading, reasoning, and arithmetic calculations. Often, processing the contents of working memory requires selective access to a single element of the memory set – for instance when one element needs to be reported, updated, or used as input for a decision. Some theories of working memory therefore assume a focus of attention as part of the working memory system, which serves to select elements within working memory for processing (Cowan, 1998; Oberauer, 2003, 2009).

The term *focus of attention* is used in two different ways. In the embedded process model introduced by Cowan (1998), the “broad” focus of attention refers to a small number of about four items that are protected from forgetting through decay and interference, and thereby form the core of working memory. Here, we will focus on a more “narrow” focus of attention that serves to select representations within working memory – typically a single item – for use in an upcoming cognitive operation (Oberauer, 2003). We are concerned with two phenomena in which an item is thought to be brought into the focus of attention. First, it is theorized that the last item in a sequentially presented list remains in the focus of attention (McElree & Doshier, 1989; McElree, 2006); second, it is postulated that an item that is retro-cued

is brought into the focus of attention (Griffin & Nobre, 2003; for a review see Oberauer & Hein, 2012). The main aim of the present study is to test whether the last-item benefit and retro-cue benefit are brought about by the same cognitive mechanism. If they are, then there should be interactive effects in a situation where both effects are present. To this end, we merged the two paradigms, which allowed us to investigate both the retro-cue benefit and the last-item benefit at the same time.

Retro-cues are seen as a tool to direct attention to the cued item during the retention interval (Griffin & Nobre, 2003; Landman, Spekrijse, & Lamme, 2003; Souza & Oberauer, 2016). In a typical retro-cue experiment, after presentation of a memory array, a cue identifies the location of one item that is most likely to be tested in a subsequently following recognition test. A validly retro-cued item can be accessed faster and more accurately in comparison to conditions where no cue, or an uninformative cue, is provided (Griffin & Nobre, 2003; Niklaus, Nobre, & van Ede, 2017; Rerko & Oberauer, 2013; Souza & Oberauer, 2016; Souza, Rerko, & Oberauer, 2016, 2014; Souza, Rerko, Lin, & Oberauer, 2014; van Ede, Niklaus, & Nobre, 2016).

The last-item benefit refers to the finding that when items are presented in serial order, retrieval speed for the last item is faster than for any other item. This observation has motivated the assumption that the last item is held in the focus of attention (McElree, 2006). Retrieval speed has been gauged through the response-deadline method, which measures speed-accuracy trade-off (SAT) functions for retrieval of

\* Corresponding author at: Department of Psychology, Cognitive Psychology Unit, University of Zurich, Binzmühlestrasse 14/22, 8050 Zurich, Switzerland.  
E-mail address: [k.oberauer@psychologie.uzh.ch](mailto:k.oberauer@psychologie.uzh.ch) (K. Oberauer).

memory items in the classic Sternberg recognition task (Sternberg, 1966). In the response-deadline paradigm, participants are instructed to give a recognition response immediately when a response signal is given. By varying the point in time when the response signal is presented after probe onset (the response-deadline lag; e.g., from 100 to 1500 ms), accuracy can be measured as a function of time over the full time course of retrieval. The growth of accuracy over time that is derived from this procedure can be characterized by three periods. As the processing time before the deadline increases, an initial period of chance performance (1) is followed by a period of increasing accuracy (2) until an asymptote (3) is reached for the final period. In their seminal study, McElree and Doshier (1989) found that the rate at which the probability of correct recognition responses increases with available response time (during the second period) differed between serial positions. The rate was increased for the last item in comparison to all previously shown items, whose rates were statistically indistinguishable from each other. This finding supported the conclusion that the last item is held in the focus of attention by default (McElree, 2006). The rate of retrieval for an item in the focus of attention is increased because when the last item appears as a probe, it can be compared to its memory representation much faster than any other item that is not held in the focus of attention (McElree, 2006). The view that the last-item benefit reflects the focus of attention is further supported by the finding that this benefit disappeared when specific instructions directed rehearsal processes towards early list items (McElree, 2006).

The notion that the last-item benefit reflects a special state of the last item has been challenged (see Cowan, 2011). Donkin and Nosofsky (2012a) showed that the model-derived memory strength for serially presented items can be described by a power-law. Memory strength is high for the last item, drops drastically already for the second-to-last item, and then becomes (decreasingly) smaller with earlier serial positions. According to this proposition, the last item does not have a qualitatively different status from other items. Rather, the last-item benefit might simply reflect the extreme point of a continuous but steep power gradient on memory strength. Another finding questions the proposition that the last-item benefit reflects the same state as an item that is selected by retro-cues. Hu, Hitch, Baddeley, Zhang, and Allen (2014) showed that the last presented item is particularly vulnerable to interference from irrelevant visual material presented after it. In contrast, research with the retro-cue paradigm has shown that a retro-cued item is protected from different kinds of visual interference (Souza, Rerko, & Oberauer, 2016; van Moorselaar, Günseli, Theeuwes, & Olivers, 2014).

In the present study, we present three experiments that directly test whether the last-item-benefit and the retro-cue benefit are empirical manifestations of the same mechanism of a single-item focus of attention. We presented items in serial order and, after a brief retention interval, assessed participants' memory with a recognition probe. The task of the participants was either to decide whether the probe was presented in the study list (Experiment 1) or whether the probe was presented at a particular position in the study list (Experiments 2 and 3). Serial presentation is thought to leave the last-presented item in the focus of attention by default. We used the response-deadline method to measure SAT functions, which allowed us to decompose the data into separate measures of retrieval speed and memory availability. This method allowed us to measure the last-item benefit specifically for retrieval speed, as described by McElree and Doshier (1989). In half of the trials, we used retro-cues to direct attention to one of the list items during the retention interval. In this way we can measure the last-item benefit and the retro-cue benefit simultaneously. We used a hierarchical Bayesian model to assess the last-item benefit and the retro-cue benefit on parameters of the SAT model.

We tested the following predictions. If the increased retrieval speed found for the last item (i.e., the last-item benefit) and the retro-cue benefit reflect the same mechanism of the focus of attention, then a retro-cue directing the focus of attention to the last item should have a

minimal effect at best, because the last item is already in a prioritized state (i.e. it is “in” the focus of attention) regardless of the cue. Under this assumption, the retro-cue benefit should be attenuated for the last item compared to the retro-cue benefit for earlier list items. In contrast, if the retro-cue benefit is a manifestation of an attentional mechanism that is different from what drives the last-item benefit, we should find additive effects of retro-cue and serial position. In other words, the retro-cue benefit should be as large when the retro-cue is directed to the last item as when it is directed to earlier items. This prediction is based on the view that an item can only be either “in” or “outside” the focus of attention. This view follows McElree (2006) who argued that the single item that is held “in” the focus of attention does not have to be retrieved in order to be acted upon, whereas all other studied items are “outside” the focus of attention. In a similar vein, Souza, Rerko, and Oberauer (2014, 2016) showed that the retro-cue benefit emerges during a 300–500 ms interval between the retro-cue and the subsequent test probe. Longer intervals after the retro-cue do not further improve accuracy and only improve response time minimally. This is consistent with the binary “in” versus “outside” view of the focus of attention put forward by McElree (2006): Once an item is retrieved into the focus of attention following a retro-cue, its accessibility is not boosted further by continuing the retrieval process. Hence, if being the last item adds further accessibility to an item on top of the retro-cue benefit, then the two beneficial effects cannot reflect the same process. To preview our main result, we found the two effects to be additive, warranting the conclusion that the last-item benefit and retro-cue benefit are not driven by the same mechanism.

## 2. Measurement model

We now outline a hierarchical Bayesian measurement model that allows us to track changes in the time-course of retrieval of memory representations. We first describe the signal detection framework of the model. Then, we introduce the SAT function that captures the pattern of performance as a function of available processing time. Next, we describe how the model is embedded in a hierarchical-Bayesian framework. Finally, we discuss the advantages of this modeling technique in comparison to previously applied procedures.

### 2.1. Signal detection framework

Our memory task is a short-term recognition task in which participants are first presented with a list of (5 or 6) serially presented stimuli. Following a short retention interval, participants are presented with a single probe for which they have to make a memory decision. Participants have to accept positive probes (i.e., in Experiment 1 a stimulus that matches any of the items presented in the study list, and in Experiments 2 and 3 a stimulus that was presented in the same spatial location at test and in the study list) and reject negative probes. We denote accept responses to positive probes as *hits* and accept responses to negative probes as *false alarms*.

We use a signal detection framework (e.g. Kellen & Klauer, 2018; Macmillan, 2002) to relate *hits* and *false alarms* in a principled manner to obtain independent measures of memory performance and response bias. We assume that the presentation of positive and negative probes evoke memory signals whose distributions can be described by a normal (i.e., Gaussian) distribution with variance 1, and mean  $\mu_N$  for the negative probes, and mean  $\mu_P$  for the positive probes. At test, participants compare the memory signal of the probe against a fixed response criterion,  $c$ . If the memory signal of the current probe is larger than  $c$ , the probe is accepted, and rejected otherwise. In mathematical terms this corresponds to the following predictions:

$$P(\text{accept}|\text{positive probe}) = P(\text{hits}) = \int_c^\infty \mathcal{N}(\mu_P, 1) \quad (1)$$

$$P(\text{accept}|\text{negative probe}) = P(\text{false alarms}) = \int_c^\infty \mathcal{N}(\mu_N, 1) \quad (2)$$

where  $\mathcal{N}$  is the probability density function of the normal distribution. Given the properties of the normal distribution, this can be simplified to

$$P(\text{hits}) = \Phi(\mu_p + c) \quad (3)$$

$$P(\text{false alarms}) = \Phi(\mu_N + c) \quad (4)$$

where  $\Phi$  is the cumulative distribution function of the normal distribution. Above-chance performance is obtained if  $\mu_p > \mu_N$ . The distance between the two distributions,

$$d' = \mu_p - \mu_N \quad (5)$$

is a common measure of memory performance or *sensitivity* that is independent of response bias. Moreover, the parameterization of  $c$  is such that positive values indicate a response bias towards accepting a probe and negative values a response bias towards rejecting a probe.

## 2.2. Performance dynamics and the SAT function

To account for the full time course of retrieval as uncovered by the response deadline method, we describe the increase in sensitivity over time with the exponential SAT function with three parameters,

$$\left. \begin{aligned} \mu_p(t) \\ \mu_N(t) \end{aligned} \right\} = \lambda(1 - e^{-\beta(t-\delta)}), \quad t > \delta, \quad \text{else } 0, \quad (6a)$$

$$d'(t) = \lambda(1 - e^{-\beta(t-\delta)}), \quad t > \delta, \quad \text{else } 0, \quad (6b)$$

where the processing time  $t$  is the duration from probe onset until the response is recorded. For reasons explained below, in Experiment 1 we estimated a separate set of SAT parameters for both  $\mu_p$  and  $\mu_N$  (Eq. (6a)), and in Experiments 2 and 3 we estimated one set of SAT parameters and restricted  $\mu_p = \frac{d}{2}$ , and  $\mu_N = -\frac{d}{2}$  (Eq. (6b)). Parameter values were allowed to be negative: For example, the asymptote  $\lambda$  of negative probes was allowed to go below zero to allow for a decreasing SAT function that captures the fact that the propensity to accept a negative probe decreases with available processing time. With these parameterizations our model accounted for both hits and false alarms.

The SAT function has been shown to adequately summarize the retrieval dynamics in response-deadline tasks (McElree & Doshier, 1989; McElree, 2006; Wickelgren, Corbett, & Doshier, 1980). The parameters of the SAT function reflect the above mentioned three phases of retrieval, and our primary interest were the estimates of these parameters as a function of our experimental conditions (i.e., each of the three SAT parameters,  $\delta$ ,  $\beta$ , and  $\lambda$ , were allowed to vary across conditions). Initially, participants perform at chance level because at short processing times no information is available to them. The intercept  $\delta$  denotes the point in time where information first becomes available and performance departs from chance. During the second phase sensitivity grows with increasing processing time with rate  $\beta$ . The intercept and rate parameter jointly describe the speed of access to memory information, which according to McElree (2006) characterizes the focus of attention: The item in the focus of attention can be accessed immediately, which is reflected in a higher rate or an earlier intercept. These measures are independent of the probability of eventually recalling a memory representation, which is captured by the third parameter, the sensitivity asymptote  $\lambda$  reached in the last phase. According to McElree (2006), this parameter reflects the availability of memory representations and is therefore not relevant for distinguishing items in the focus of attention from other items. Accordingly, we restricted our hypotheses to effects on the rate and intercept parameter. However, in line with the work of McElree we could not make more specific predictions about which of the latter two should reflect the last-item benefit. To foreshadow our results, our manipulations targeting attentional prioritizations in working-memory affect the intercept but not the rate parameter.

## 2.3. Hierarchical Bayesian framework

The signal-detection SAT model was implemented in a hierarchical Bayesian framework (Gelman et al., 2013). In a Bayesian framework, one's information regarding the parameters is specified by probability distributions. The state of ignorance before any data is collected is represented via *prior distributions* (or *priors*). These priors are then updated in light of the data using Bayes' theorem. The resulting new state of knowledge, the *posterior distribution*, can be used for statistical inference. Here, we employed an efficient version of Hamiltonian Monte Carlo to obtain samples from posterior distributions (Carpenter et al., 2017).

For each experimental condition of interest,<sup>1</sup> we obtained posterior distributions for the three SAT parameters,  $\delta$ ,  $\beta$ , and  $\lambda$ . The posterior distributions represent the probabilities of the parameters conditional on data and model (where the latter includes the prior) and thereby directly allow statistical inference (Gelman, Carlin, Stern, & Rubin, 2014). To assess the difference between parameters of two conditions, we simply subtracted the posterior distributions of the to-be-compared conditions from each other to obtain a posterior distribution of their difference. For ease of interpretation, we always subtracted the distribution of the smaller parameter value from the distribution of the larger parameter value. In this way, between 0% and 50% of the posterior difference distribution lies below zero. The smaller the proportion below zero – or the larger the proportion above zero – the stronger the evidence for a difference between the two conditions. To gauge the strength of evidence for a difference, we calculated  $p_B$  as the proportion of the difference distribution below zero, multiplied by two. This makes  $p_B$  a statistic that ranges from zero to one, with values near zero denoting evidence for a difference, and values near one indicating that equal mass of the posterior difference distribution extended below and above zero. Values near one therefore provide some evidence against a difference.

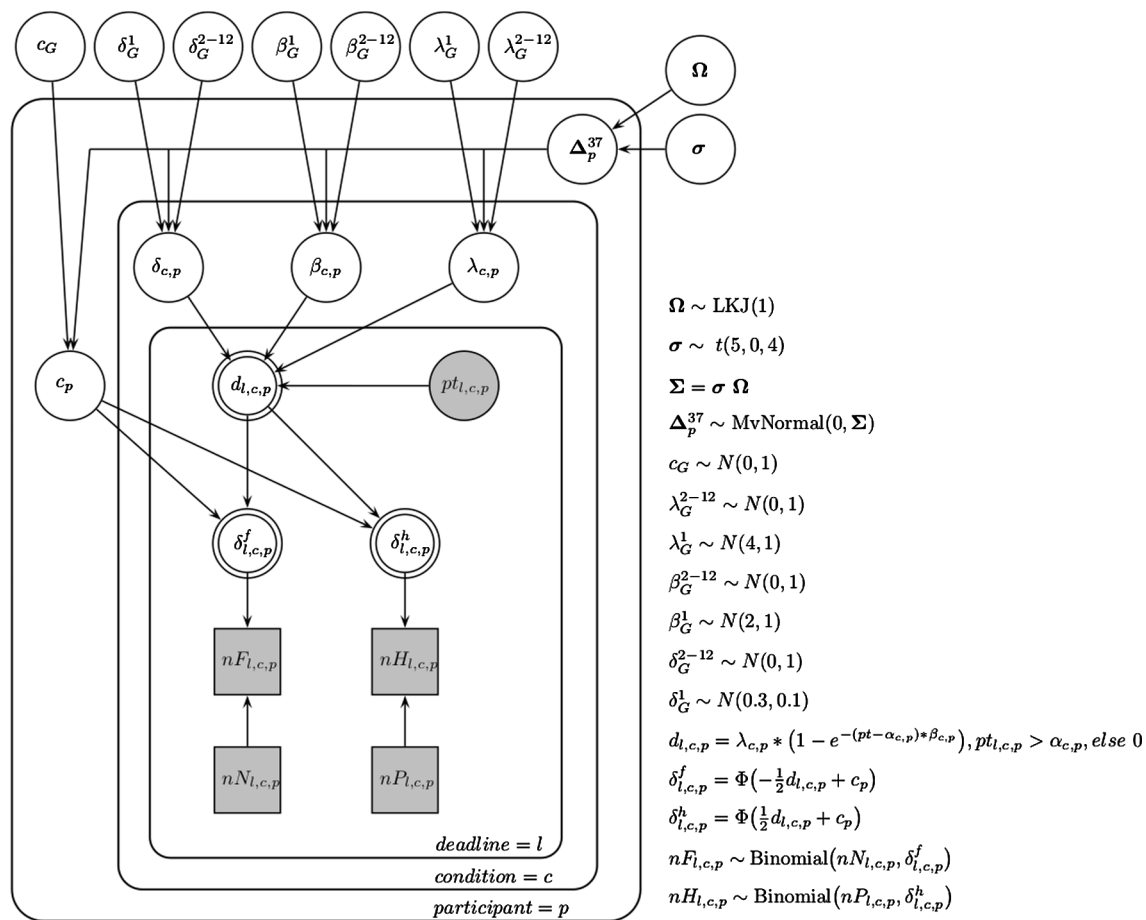
To adequately account for both inter-individual and intra-individual variability, we implemented the model in a hierarchical fashion using so-called *partial pooling* (Gelman & Hill, 2007). Individual parameters were assumed to come from group-level distributions. We assumed normal group-level distributions for all three SAT parameters, as well as for the signal detection criterion. In addition, we estimated the full variance-covariance matrix among all parameters (i.e., the group-level distribution was multivariate normal; Klauer, 2010). Note that all statistical inference was performed on the group-level parameters.

A graphical model of the hierarchical-Bayesian implementation of the signal-detection SAT model is depicted in Fig. 1. We used either non-informative priors (a so-called LKJ-prior with scale 1 for the correlation matrix of the multivariate normal distribution) or weakly informative priors with most of their mass on reasonable parameter values (following Gelman et al., 2014). To account for differences between conditions, we estimated separate SAT parameters for each experimental condition of interest, but only one overall signal-detection criterion  $c$  (i.e.,  $c$  did not vary across conditions).

### 2.3.1. Advantages of the present modeling approach

Previous SAT studies have commonly fitted the SAT function to estimates of  $d'$  of each individual participant, before individual parameter estimates or model performance indices were averaged. The crucial conclusions were then drawn from comparing a set of models using model performance indices (Liu & Smith, 2009; McElree & Doshier, 1989; McElree, 2006; Mızrak & Öztekin, 2016b, 2016a; Öztekin & McElree, 2010). For example, in the seminal study by McElree and Doshier (1989), the authors compared a model with the same rate for all

<sup>1</sup> Serial position and cue condition were considered conditions of interest. The response deadline lag, although experimentally manipulated, was not, as it was part of the SAT function and thereby already accounted for in the model.



**Fig. 1.** Graphical depiction of the hierarchical Bayesian SAT model. Observed variables are represented by shaded nodes. Discrete variables are displayed as squared nodes and continuous variables are displayed as circular nodes. Deterministic nodes have a double border. The direction of arrows indicates that the node at the end of the arrow depends on the node at the start of the arrow. Plates visualize the hierarchical structure in the data. Subscripts denote different conditions, superscripts denote the length or index of vectors.  $N$  is the probability density of the normal distribution.  $\text{MvNormal}$  is the multivariate extension of  $N$ .  $\Phi$  is the cumulative distribution function of the normal distribution.  $nF$  and  $nH$  are the number of false-alarms and hits, respectively.  $nN$  is the number of negative, and  $nP$  the number of positive probes.

serial positions, a model with two rates – one for the last position and one for all previous positions – and a model with a separate rate for each serial position. Among these three models, the second (with 2 rate parameters) provided the best account of the data (i.e., highest adjusted- $R^2$ ). The same pattern of results was found when effects of serial position were modeled on the intercept parameter. (The rate model was chosen as the winning model because its overall performance in terms of adjusted  $R^2$  was slightly better than when the effects were modeled on the intercept parameter.) In this model-selection procedure, only a limited set of models has been considered (but see Mızrak & Öztekin, 2016b). For instance, it was not tested whether a model that assigns the second-to-last and last item an increased rate would fare better than any of the other models. Such comparisons, however, are crucial to investigate the nature of the last-item benefit. If the last-item benefit is unique to the last item (McElree, 2006), retrieval rate should show a dichotomous pattern with only the last item showing an increased rate. In contrast, if the last-item benefit represents the peak of a steep power gradient of memory strength (Donkin & Nosofsky, 2012a), the second-to-last item could also show a higher rate than any earlier item.

The present approach improves on this procedure in several regards. First, in contrast to the classical approach in which  $d'$  is calculated in a first step and the SAT function is applied to the calculated  $d'$  in a second step, we estimate the signal-detection model and the SAT function in one step, which avoids the accumulation of estimation error. Second, our model avoids arbitrary corrections for hit-rates or false-alarm rates of zero or one, which are necessary to compute  $d'$  using the classical

approach (Stanislaw & Todorov, 1999). Third, a Bayesian statistical approach provides us with full posterior distributions, which allows us to assess the precision of the estimates in a direct manner. Fourth, a general property of hierarchical models is that individual and group-level parameters are estimated simultaneously using partial pooling. The estimation of each individual parameter is informed by the data of all participants because all data are used to estimate the group-level parameters, which at the same time provide soft constraints for the individual-level parameter estimates. The crucial benefit of this procedure is that both the group-level parameters and the individual-level parameters are estimated more precisely because unrealistic or extreme individual parameter estimates have been constrained (Katahira, 2016).

One further difference between the current approach and the previous approach is that we did not base our inference on model comparison, but on parameter estimates within one encompassing model in which all parameters were allowed to vary freely across conditions. The reason for this choice is two-fold. First, given the large space of all possible models (i.e., all possible partitions resulting from the number of serial positions times two for the cue conditions, for each of the 3 SAT parameters), an exhaustive exploration of the full model space is comparatively expensive. Second, penalized model fit indices that rely on counting the model parameters such as adjusted  $R^2$ , AIC, and BIC, were developed in the context of linear models (Burnham & Anderson, 2003) and assume that each parameter has an approximately equal influence on the complexity of the model. This assumption is at least questionable for a nonlinear model such as the SAT. A principled model

comparison within a Bayesian framework requires calculating the Bayes factor for each pairwise model comparison, which we found not to be feasible with current methods.

### 3. Experiment 1

In our first experiment we employed a Sternberg task merged with a retro-cue paradigm. In the study phase participants had to remember a list of six words. Presentation occurred in serial order in six spatial locations located along a virtual circle. After a brief period of time, participants were asked whether a centrally presented probe matches any of the words presented during the study phase. In half of the trials, a spatial retro-cue indicated the word that will be probed in a positive trial. Participants' processing times were manipulated using a response-deadline method, allowing us to track the full time course of retrieval. The data and the analysis scripts for all experiments can be accessed in the Open Science Framework (<https://osf.io/6apd9/>).

#### 3.1. Method

##### 3.1.1. Participants

We recruited 16 volunteers (13 females, mean age = 23) through the University of Zurich participant volunteer pool who participated in eight 1 h test sessions. All participants read and signed an informed consent form before participation. They completed one or two training sessions beforehand to acquaint themselves with the response-deadline method. No further practice trials were run in the test sessions. Participation was reimbursed with 15 Swiss Francs per session. Due to technical problems, we could not record data from one session of one participant. We excluded one participant from the analysis due to below-chance performance even at long response-deadline lags.<sup>2</sup>

##### 3.1.2. Materials

For each trial, we took a pseudo-random subset of words from a set of 639 one- or two-syllable German nouns consisting of four to five letters. The sampling algorithm ensured that neither one of the six study words, nor a negative probe word, had appeared in any of the previous three trials. The experiment was programmed and run in MATLAB using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997).

##### 3.1.3. Design

Each test session consisted of a total of 432 trials resulting from three complete permutations of twelve response-deadline lags, six serial positions of items matching the probe (for positive probes), and whether the probe had appeared in the study phase (positive probe) or not (negative probe). The trials were presented in random order. The presentation of a retro-cue (cue condition) was varied across odd and even sessions. Nine participants began with a session with retro-cues, and eight participants with a session without retro-cues. Across eight test sessions, this design yielded a total of 12 trials for each combination of response-deadline lag, serial position, and cue condition for a positive probe, and 72 trials per response-deadline lag and cue condition for negative probes (which cannot be associated with a serial position).

##### 3.1.4. Procedure

Fig. 2 displays the flow of events in Experiment 1. Each trial began with the presentation of six blue frames (visual angles with a viewing distance of 50 cm: width = 8.5°, height = 6.7°), equally distributed on a virtual circle (diameter = 23°) on grey background for 1000 ms. Frames remained on the screen for the entire trial. Then, six words were presented, one in each frame, with a presentation time of 450 ms and an inter-stimulus-interval of 50 ms. Our choice of the 500 ms stimulus

onset asynchrony follows the work of McElree (2006) and McElree and Doshier (1989) who found the last-item benefit with very similar presentation times. Presentation occurred in serial order along the virtual circle in clock-wise direction starting from the frame at the top of the screen. After presentation of the last word, the empty frames were shown for 500 ms.

The sequence of events that followed depended on the cue condition. In the retro-cue condition, an arrow was presented for 500 ms that indicated with certainty the word that would match the probe if the probe was positive. Then, after a 500 ms post-cue interval, the test probe was shown in the center of the screen. In the no-cue condition, the test probe was presented immediately, such that the retention interval matched the pre-cue interval in the retro-cue condition. These timings were chosen to rule out a decay-based explanation of the retro-cue benefit (Rerko & Oberauer, 2013).

Participants indicated whether the probe matched any word of the study set ("accept" responses were indicated with the "j" key, "reject" responses with the "f" key). The probe appeared for a variable length of time, depending on the response-deadline lag. At 100, 121, 164, 227, 312, 481, 545, 693, 864, 1054, 1267 or 1500 ms after onset of the probe, the probe disappeared from the screen, and participants were cued by a tone (the response signal; duration = 50 ms, frequency = 2000 Hz, played over headphones) to immediately respond.<sup>3</sup> Participants were instructed to respond within 270 ms regardless of their ability to make a correct response. They received visual latency feedback, which provided their response time as well as written feedback in the form of "schneller antworten" (respond more rapidly) for latencies above 270 ms, "Bitte antworten Sie erst nach dem Tonsignal" (please respond only after hearing the auditory cue) for anticipated responses with a response latency below 100 ms, and "Rechtzeitig" (in time) for responses within the accepted time window. Each trial was initiated by pressing the space bar.

#### 3.2. Results

##### 3.2.1. Response latencies

We considered trials with a response time below 100 ms as anticipations or motor errors. Trials above 500 ms were likely due to attentional lapses undermining the response-deadline method. We therefore excluded extreme trials with response latencies above 500 ms or below 100 ms from analyses (2.35%).

To verify that participants obeyed the deadline response instructions, we investigated their response times. After exclusion of extreme trials, and averaged across participants and experimental conditions, participants met the response-deadline criterion (< 270 ms) in 92.9% of all trials.

Response latencies have been reported to be longer for shorter lags than for longer lags (McElree & Doshier, 1989). To account for such differences across conditions and participants in our model-based analysis, we created a new variable called processing time by adding the mean reaction time per individual and condition cell to the response-deadline lag of this specific condition (McElree & Doshier, 1989; McElree, 2006; Mizrak & Öztekin, 2016b).

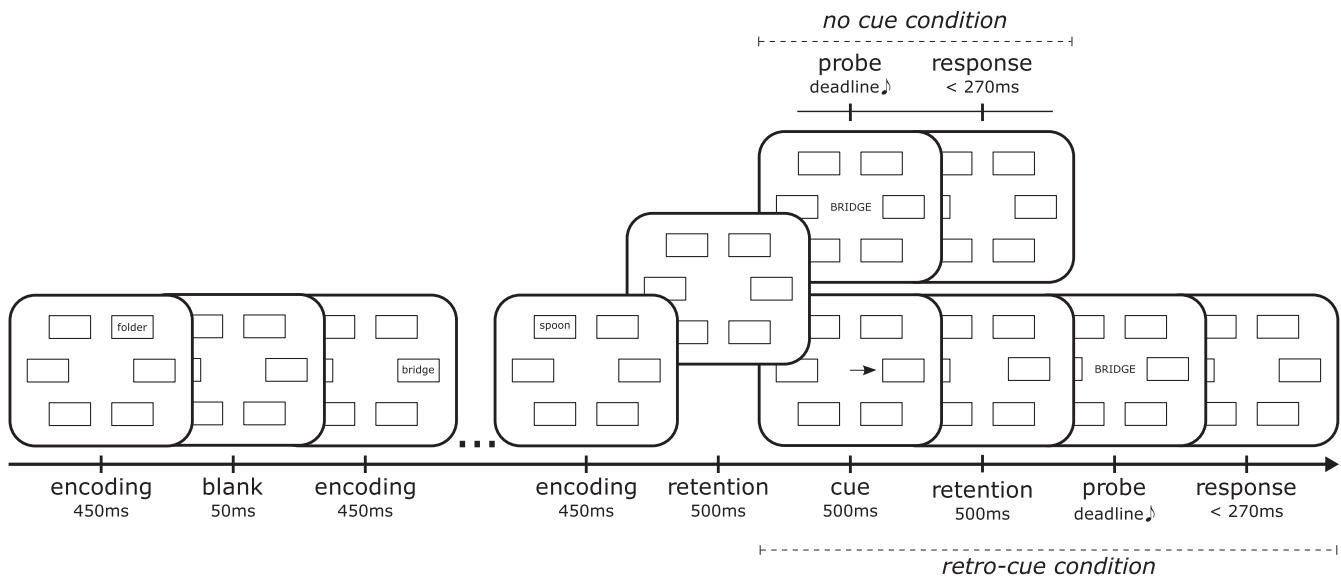
##### 3.2.2. Model-based analysis

Due to the central presentation of probes, non-cued new probes cannot be associated with any serial position. To account for this, we analyzed the data with a model that estimates the means of the signal distributions, and consequently the SAT parameters, for positive and negative probes separately.

$$\mu_p = \lambda_{pos}(1 - e^{-\beta_{pos}(t - \delta_{pos})}), \quad t > \delta_{pos}, \quad \text{else } 0, \quad (7)$$

<sup>2</sup> This participant probably confused the response keys. The pattern of results does not change when the data is recoded and included in analyses.

<sup>3</sup> Actual presentation times may have varied slightly due to the 60 Hz refresh rate of the monitor.



**Fig. 2.** The sequence of events in Experiment 1. Experiment 2 and Experiment 3 differed from this procedure as follows: Only five instead of six words (Experiment 2) or color patches (Experiment 3) were presented during encoding. Moreover, probes were presented in one of the locations of the study items. In Experiment 3, we varied the spatial position of the last item and extended the retention interval between the offset of the last item and the onset of the probe or the retro-cue from 500 ms to 1000 ms.

$$\mu_N = \lambda_{neg}(1 - e^{-\beta_{neg}(t - \delta_{neg})}), \quad t > \delta_{neg}, \quad \text{else } 0. \quad (8)$$

As a consequence, we will report credible differences of the three SAT parameters between experimental conditions of interest separately for positive and negative probe trials (results of negative probes are reported in Appendix A).

The model was estimated in Stan through R (R Core Team, 2014) using rstan (Carpenter et al., 2017). After discarding 1000 warmup samples, we retained a total of 1000 post-warmup samples for each of 4 independent chains, keeping every second sample. Convergence statistics indicated good mixing behavior with  $\hat{R} \leq 1.02$  for all estimated model parameters (Gelman & Rubin, 1992). Visual inspection of MCMC trace plots of the group-level parameters indicated the same. The number of effective samples was above 400 for all estimated model parameters.<sup>4</sup>

Model fits are depicted in Fig. 3, which compares the mean of the predicted proportions of “accept” responses, aggregated across participants and posterior samples (the lines), to the observed proportions of “accept” responses, aggregated across participants (the points), for positive and negative probes separately. Visual inspection of the model fit shows that the model accounted well for the retrieval dynamics of all experimental conditions.

**3.2.2.1. Positive trials – last item benefit.** The last-item access speed benefit is expected to be expressed in an advantage of serial position 6 in comparison to serial positions 1–5 in either the intercept or rate parameter, or both (McElree & Doshier, 1989; McElree, 2006). For the rate parameter we found no credible pairwise differences between any of the conditions, all  $p_B > 0.05$ . Furthermore, inspection of Fig. 4

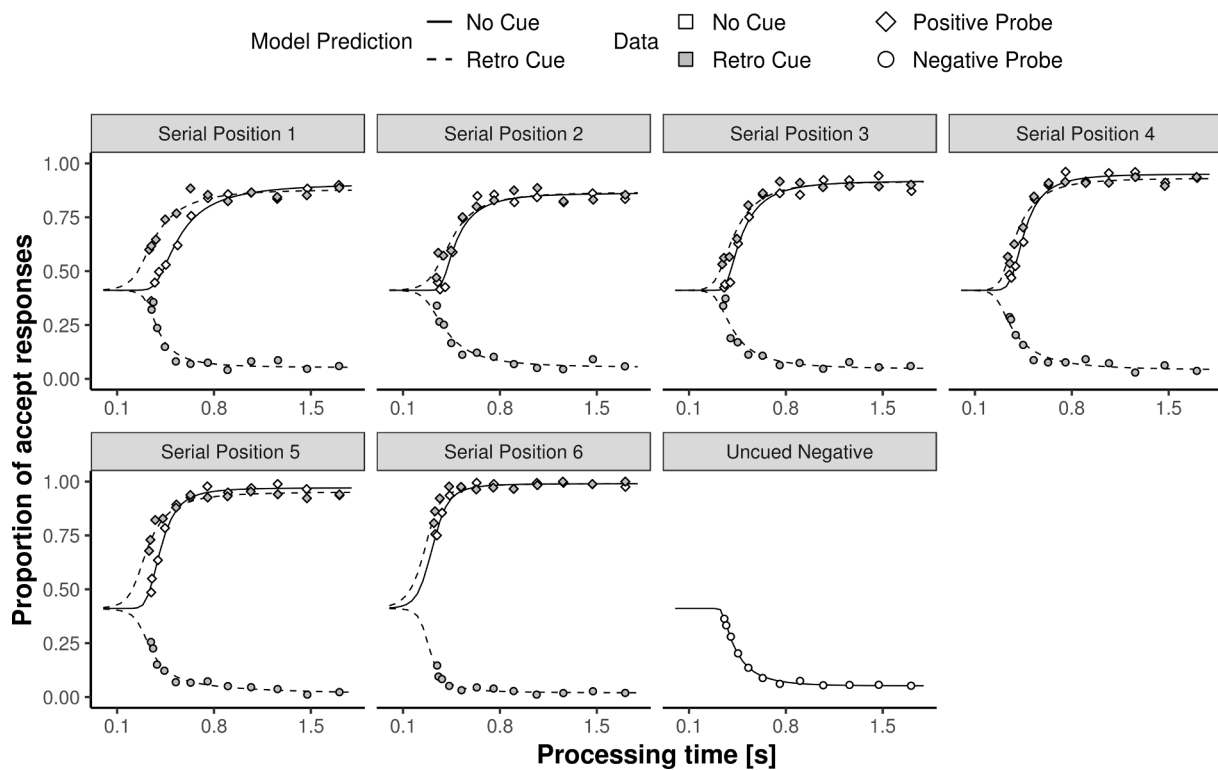
(middle panel) suggests no results pattern relevant to our research question. Consequently, we focused on the intercept parameter in the following (Fig. 4, top panel).

To test the last-item benefit, we compared the mean intercept for serial positions 1 to 5 with the mean intercept for serial position 6 across retro-cued and non-cued probes. This comparison revealed that serial position 6 had a credibly smaller intercept ( $p_B = 0.001$ , median benefit = 114.8 ms [95% credible interval (CI) = 64.0, 170.9]). The last-item benefit was also observed when we compared the mean intercept for serial positions 1–5 with the mean intercept for serial position 6 separately for non-cued probes ( $p_B < 0.001$ , 134.9 ms [79.9, 210.4]) and for retro-cued probes ( $p_B = 0.03$ , 91.8 ms [9.7, 186.9]). For comparison, in an experiment that presented only non-cued probes, McElree and Doshier (1989) reported a somewhat smaller last-item benefit of 74 ms on the intercept parameter.

We additionally investigated all pairwise comparisons between serial positions for each cue condition. For non-cued probes, the intercept for serial position 6 was smaller than each of the intercepts of serial positions 1–5 (all  $p_B < 0.01$ ), whereas the intercepts of serial positions 1–5 could not be credibly differentiated from each other, with the exception of serial position 5 which was credibly shorter than serial position 2 ( $p_B = 0.02$ ). For retro-cued probes, the intercept of serial position 6 was credibly smaller than the intercepts of serial positions 2 ( $p_B = 0.02$ ), 3 ( $p_B = 0.02$ ), and 4 ( $p_B = 0.02$ ), but there was no such evidence in comparison to the intercepts of serial positions 1 ( $p_B = 0.28$ ) and 5 ( $p_B = 0.56$ ).

**3.2.2.2. Positive trials – retro-cue benefit on intercept.** We next investigated the effects of displaying a retro-cue during the retention interval. As mentioned before, we found no credible pairwise differences for the rate parameter, and we therefore focused on the intercept parameter. Aggregated across all serial positions we found a smaller intercept for retro-cued in comparison to non-cued probes ( $p_B < 0.001$ , 90.2 ms [53.7, 127.8]). Fig. 5 shows the retro-cue benefits across serial positions. As can be seen, the 95% CIs do not include 0 for all but two serial positions, 2 and 6. However, even for those serial positions, the posterior median was very similar to the effects observed for serial positions 3 and 4. In addition, for serial position 2, 95.0% and for serial position 6, 85.6% of the posterior mass was above zero.

<sup>4</sup> The Stan algorithm reported several *divergent transitions* suggesting a non-smooth geometry of the bulk of the posterior probability mass, potentially questioning the validity of the approximation to the posterior distribution. Following Gabry, Simpson, Vehtari, Betancourt, and Gelman (2017), we performed a visual inspection of the coordinate plots for all 4350 parameters (see supplemental material). This inspection indicated no “particular structure” of the divergent transitions implying that these divergent transitions are false alarms and the approximation to the posterior distribution is accurate. Furthermore, the results in Experiments 2 and 3 were very similar to the ones for Experiment 1 and at the same time there were no divergent transitions in Experiments 2 and 3, further strengthening the validity of the present results.



**Fig. 3.** Observed (points) and predicted (lines) group-level proportions of accept responses for positive (diamonds) and negative (circles) probes for each serial position and cue condition as a function of processing times (response-deadline lag plus the individual mean response time per experimental condition) of Experiment 1. Filled objects connected through a dashed line depict retro-cued probes whereas non-filled objects connected through a solid line depict non-cued probes. Non-cued negative probes cannot be associated with a serial position and are depicted in their own panel.

The theoretically most important question is whether the retro-cue benefit is attenuated for the last item. To test this prediction, we compared the retro-cue benefit of serial position 6 against the mean retro-cue benefit of serial positions 1 to 5. This comparison provided no evidence for an attenuation of the retro-cue benefit ( $p_B = 0.41$ ). We also compared the retro-cue benefit of serial position 6 with that of each earlier serial position individually. Fig. 5 reports the  $p_B$  values for these comparisons and shows that none of these comparisons provides credible evidence for an attenuation of the retro-cue benefit (smallest  $p_B = 0.19$ ). In other words, despite the fact that the retro-cue benefit appears to be slightly smaller for serial position 6, there is no evidence for this reduction when comparing the retro-cue benefits at different serial positions directly.

**3.2.2.3. Positive trials – asymptote.** Fig. 4 (bottom panel) shows the estimates of the asymptote parameter of the SAT model. As is clear from the figure, we found no retro-cue benefits for the asymptote, neither when all non-cued probes were compared to all retro-cued probes ( $p_B = 0.67$ ), nor when the non-cued and retro-cued conditions were compared for each serial position individually (all  $p_B > 0.40$ ).

Fig. 4 hinted towards a primacy effect for the first position, together with an extended recency effect for serial positions 2–6. This matches the serial-position curve often observed with the Sternberg task using the deadline method (McElree & Doshier, 1989), as well as studies using free-response paradigms (Donkin & Nosofsky, 2012b; Monsell, 1978; Nosofsky & Donkin, 2016; Nosofsky, Little, Donkin, & Fific, 2011; Oberauer et al., 2018; Ratcliff, 1978). To test whether there was evidence for this pattern we formalized it as a set of order constraints (i.e.,  $SP1 > SP2 < SP3 < SP4 < SP5 < SP6$ ) and calculated the proportion of posterior samples for which exactly this pattern held. This proportion was 0.43 for the non-cued probes and 0.40 for the cued probes. Because there exist a total of  $6! = 720$  such orderings, the prior probability of obtaining this ordering is  $\frac{1}{720} = 0.001$ . Therefore, the

Bayes factor for this particular ordering is at least  $\frac{0.40}{0.001} = 288$ , providing considerable support for both a primacy and a recency effect on the asymptote parameter.

**3.2.2.4. Bias and correlations.** The median bias parameter was  $-0.26$  [ $CI = -0.53, 0.01$ ], indicating an overall bias to reject probes. We obtained no substantial correlation between individual-level parameters above  $r = 0.13$ , and all CIs included zero.

### 3.3. Discussion

We merged the classical Sternberg task with the retro-cue paradigm in order to investigate whether the retro-cue benefit is attenuated for the last item. This prediction can be derived from the proposition that both phenomena are reflections of the same mechanism of a single-item focus of attention. We measured the full time course of retrieval using a response-deadline method. Analysis of the data using a hierarchical Bayesian implementation of the SAT function showed that the retro-cue benefit for the last item, which already benefits from being presented last, was not credibly smaller than the retro-cue benefit at other serial positions. This supports the claim that the retro-cue benefit is not attenuated for the last item.

One potential criticism of the results of Experiment 1 is that the retro-cue benefit for the last item appears to be attenuated compared to the other items. However, a careful inspection of Fig. 5 shows that even descriptively this is only really the case in comparison to serial positions 1 and 5. In addition, an analysis of the posterior samples provides no evidence for this critique. Rather, the last-item benefit and the retro-cue benefit are additive effects of attentional prioritization in working memory.

A direct extension of the classical Sternberg task with the retro-cue paradigm as employed here results in two potential problems that limit the conclusions that can be drawn from this experiment. First, in this

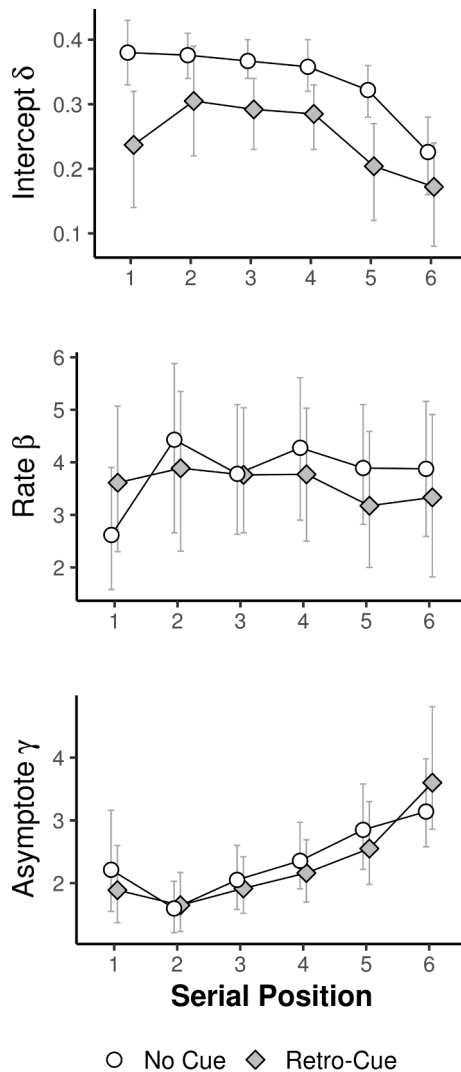


Fig. 4. Median parameter values for each SAT parameter for the positive probes of Experiment 1. Error bars depict the 2.5th and 97.5th percentile of the posterior distribution.

experiment we only used valid retro-cues. Thus, non-cued trials require the comparison of the probe stimuli to six other stimuli, whereas for retro-cued trials only one comparison is necessary. As a consequence, it remains a possibility that the retro-cue benefit could be solely driven by a reduction in the number of comparisons that need to be performed (Makovski, Sussman, & Jiang, 2008). Second, because negative probes were not associated with a serial position (at least for non-cued trials) we could not calculate  $d'$  independently for each serial position. Consequently, we had to calculate SAT parameters separately for positive and negative probes, which led to a model with extra parameters that were not of direct relevance to the research question (i.e., the SAT parameters for the negative probes). In the classical SAT approach (McElree & Doshier, 1989), this problem does not occur because  $d'$  is calculated from the observed data before fitting the SAT function, recycling the negative probes for each serial position. As our Bayesian approach required us to specify the likelihood of the data (whereas the classical approach simply minimizes squared deviations) such a recycling would have been mathematically inappropriate.

In order to address these shortcomings of Experiment 1, in Experiment 2 and Experiment 3 probes were presented in one of the locations where items were presented in the study phase. The task required participants to compare the probe to the study item that had been presented in the probe's location. This ensures that only one

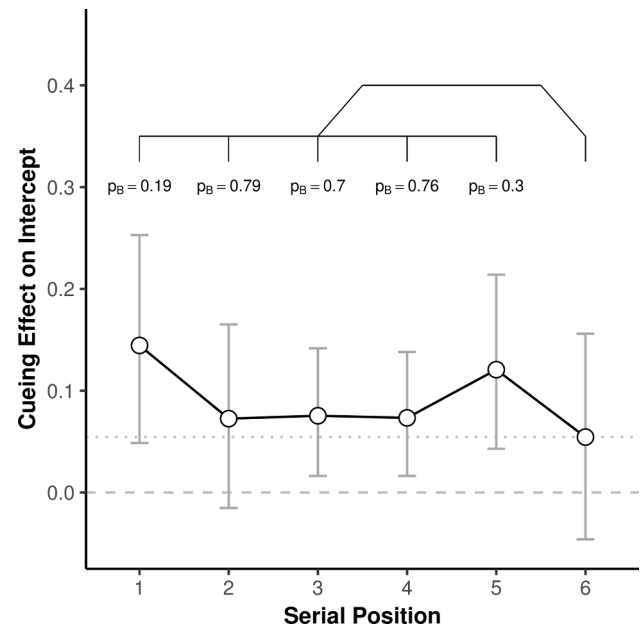


Fig. 5. Median group-level posterior estimates for the retro-cue benefit in seconds for each serial position based on positive probes of Experiment 1.  $p_B$  above serial positions 1 to 5 denotes the evidence for a difference of the cueing effect between this particular serial position and serial position 6. The dotted line depicts the median cueing effect for the last serial position. The dashed line indicates the absence of a cueing effect. Error bars depict 95% CI.

comparison is required for non-cued and retro-cued probes, and further allows non-cued negative probes to be associated with the serial position of the probed location. This then allows us to estimate a single set of SAT parameters for both negative and positive probes. With this procedure, we can more directly compare retro-cueing effects across serial position.

#### 4. Experiment 2

In our second experiment, we modified the procedure of Experiment 1 only slightly by presenting location-specific probes (see Fig. 2). The participants' task was to indicate whether the probe and the item at the same location in the study array matched. In half of the trials, we presented a retro-cue that validly indicated the spatial location of where the probe will appear.

##### 4.1. Method

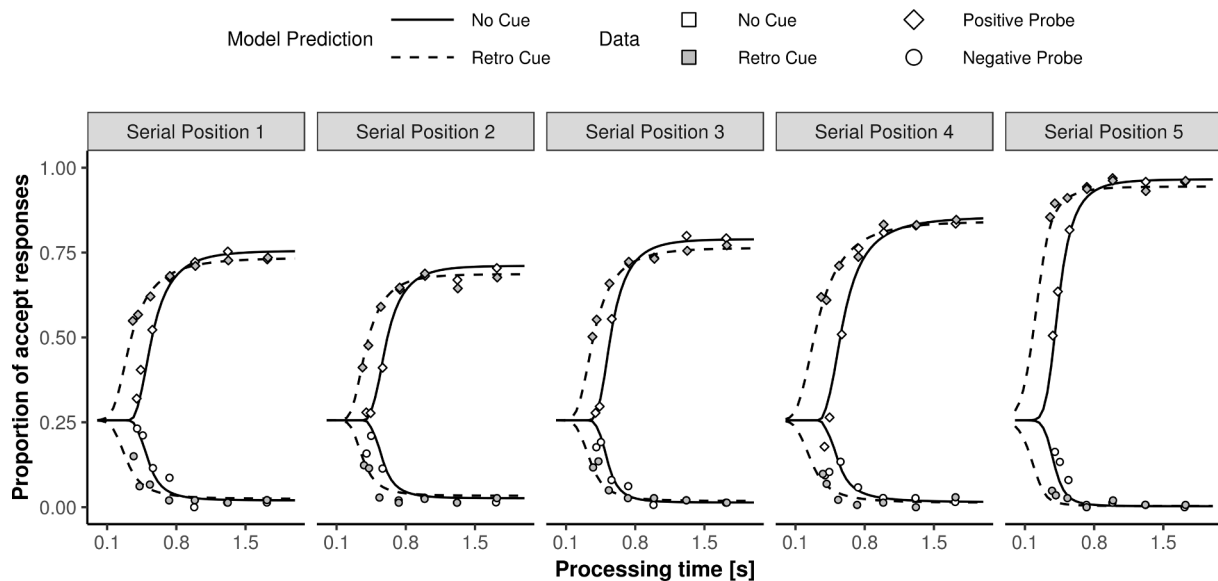
###### 4.1.1. Participants

Eleven volunteers (8 females, mean age = 25), recruited through the University of Zurich participant volunteer pool, participated in ten test sessions each lasting 1 h. 1–2 practice sessions were completed by each participant. All participants read and signed an informed consent form before participation. Due to technical problems, data from one session of one participant was not recorded.

###### 4.1.2. Materials and design

We used the same materials as in Experiment 1, with the slight modification that we only selected five words per trial.

In each session, participants completed 420 trials. In half of all trials we presented a positive probe, in a quarter of all trials a new probe which had not been part of the study set, and in the remaining quarter of trials we presented a probe that had been presented at a different serial position, a so called intrusion probe. Intrusion probes were equally likely to be chosen from all not-tested serial positions. Serial position, probe type, and response-deadline lag were permuted within



**Fig. 6.** Observed (points) and predicted (lines) group-level proportions of accept responses for positive (diamonds) and negative (circles) probes for each serial position and cue condition as a function of processing times of Experiment 2. Filled objects connected through a dashed line depict retro-cued trials whereas non-filled objects connected through a solid line depict non-cued trials.

each session and cue condition was varied across sessions. Five participants started with a session with retro-cues, and the remaining six started with a session without retro-cue presentation. This design yielded 30 positive trials for each combination of response-deadline lag, cue type, and serial position across the entire experiment. For negative trials, this design yielded 15 new and 15 intrusion probe trials for each experimental condition cell across the entire experiment.

#### 4.1.3. Procedure

The same procedure was applied as in Experiment 1 with the following changes: Only five words were presented in five boxes that were presented equidistantly on a virtual circle. Also, for the response-deadline paradigm, the number of response-deadline lags was reduced to seven. Participants were cued by a response signal to give an immediate response 100, 167, 300, 500, 767, 1100, or 1500 ms after probe onset.

## 4.2. Results

Our main analyses will be restricted to positive and new negative probes. Analyses with intrusion probes can be found in Appendix B. The analysis of the intrusion probes supports the same conclusions regarding the interaction of serial position and cue condition. However, visual inspection of model fits suggests that SAT curves for intrusion probes require a more substantive theory regarding the underlying processes in order to capture the recognition performance dynamics during early processing times (Göthe & Oberauer, 2008; Oberauer, 2008). Here, we chose to fit the descriptive SAT model introduced by McElree and Doshier (1989) in order to maintain comparability between experiments and analyses.

#### 4.2.1. Response latencies

We excluded 2.11% of all trials due to extreme response latencies. We then investigated participants' response latencies to verify that participants obeyed the response-deadline instructions. After exclusion of extreme trials and averaged across participants and experimental conditions, the response-deadline criterion was met in 80.52% of all trials. To account for different response latencies across experimental conditions, we again computed the processing time for each combination of participant, response-deadline lag, serial position and cue

condition and used these times in the model based analysis.

#### 4.2.2. Model based SAT analysis

After discarding 1000 warmup samples, we retained 1000 post-warmup samples for each of 4 independent chains, keeping every second sample. Convergence statistics indicated good mixing behavior with  $\hat{R} \leq 1.01$  for all estimated model parameters (Gelman & Rubin, 1992). Visual inspection of MCMC trace plots of the group-level parameters indicated the same. The number of effective samples was above 360 for all estimated model parameters.

Model fits are depicted in Fig. 6, which compares the mean of the predicted proportions of “accept” responses, aggregated across participants and posterior samples (the lines), to the observed proportions of “accept” responses, aggregated across participants (the points), for positive and negative probes separately. Visual inspection of the model fit shows that the model accounted well for the retrieval dynamics of all experimental conditions.

**4.2.2.1. Last-item benefit.** As in Experiment 1, we found no credible differences between the experimental conditions for the rate parameter. Likewise, Fig. 7 (middle panel) did not suggest any pattern relevant to our research question. Consequently, we again focused on the intercept parameter (Fig. 7, top panel).

To test the last-item benefit, we compared the mean intercept for serial positions 1 to 4 with the mean intercept for serial position 5 across retro-cued and non-cued probes. This comparison revealed a credibly smaller intercept for serial position 5 ( $p_B = 0.001$ , 124.7 ms [53.9, 222.6]). We also observed this last-item benefit for non-cued probes ( $p_B < 0.001$ , 130.9 ms [62.6, 202.8]) and to a slightly lesser degree for retro-cued probes ( $p_B = 0.05$ , 116.6 ms [−1.5, 302.3]).

We additionally investigated all pairwise comparisons for each combination of serial position and cue condition. For non-cued probes, the intercept for serial position 5 was smaller than the intercepts of serial positions 1–4 (all  $p_B < 0.01$ ), whereas the intercepts of serial positions 1–4 could not be credibly differentiated from each other. For retro-cued probes, the intercept of serial position 5 was smaller than the intercepts of serial positions 2 ( $p_B = 0.001$ ) and 3 ( $p_B = 0.02$ ), but there was no such evidence in comparison to the intercept of serial positions 1 ( $p_B = 0.22$ ) and 4 ( $p_B = 0.61$ ).

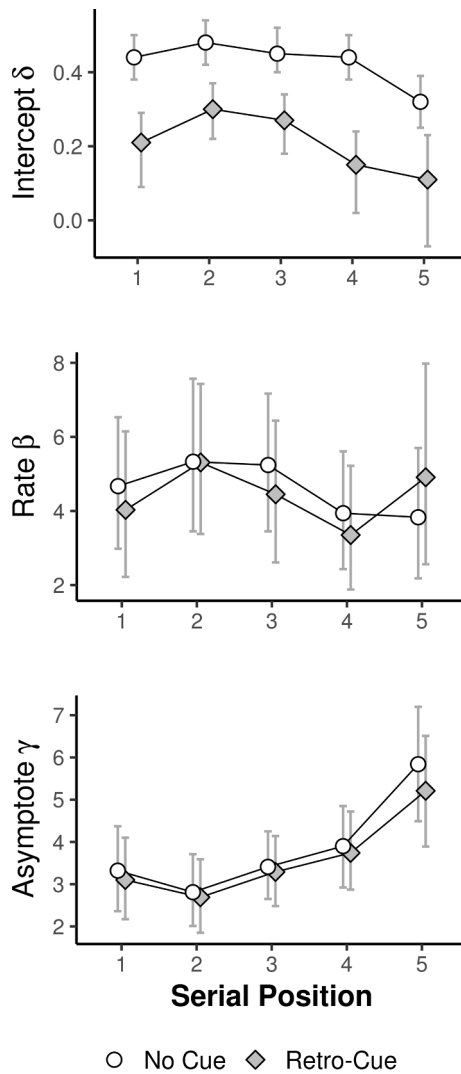


Fig. 7. Median parameter values for each SAT parameter of Experiment 2. Error bars depict the 2.5th and 97.5th percentile of the posterior distribution.

**4.2.2.2. Retro-cue benefit.** We next investigated the effects of presenting a retro-cue. Again, we focused on the intercept parameter, as we found no credible pairwise differences for the rate parameter. Aggregated across all serial positions, we found a smaller intercept for retro-cued in comparison to non-cued probes ( $p_B < 0.001$ , 222.8 ms [177.8, 285.2]). As can be seen in Fig. 8, the retro-cue benefit was observed for each individual serial position.

To investigate whether the retro-cue benefit is attenuated for the serial position 5, we compared the magnitude of the retro-cue benefit of serial position 5 against the mean retro-cue benefit of serial positions 1–4. Again, this comparison yielded no evidence for an attenuation of the retro-cue benefit ( $p_B = 0.87$ ). We also compared the retro-cue benefit of serial position 5 with each earlier serial position individually. Fig. 8 reports the  $p_B$  values for these comparisons and demonstrates that none of these comparisons provides credible evidence for an attenuation of the retro-cue benefit (all  $p_B > 0.41$ ). Descriptively the retro-cue benefit of serial position 5 was even larger than the retro-cue benefits of serial positions 2 and 3.

**4.2.2.3. Asymptote.** Fig. 7 (bottom panel) shows the estimates of the asymptote parameter of the SAT model. We found no retro-cue benefits, neither when all non-cued probes were compared to all retro-cued probes ( $p_B = 0.25$ ), nor when the non-cued and retro-cued conditions were compared individually for each serial position (all  $p_B > 0.65$ ).

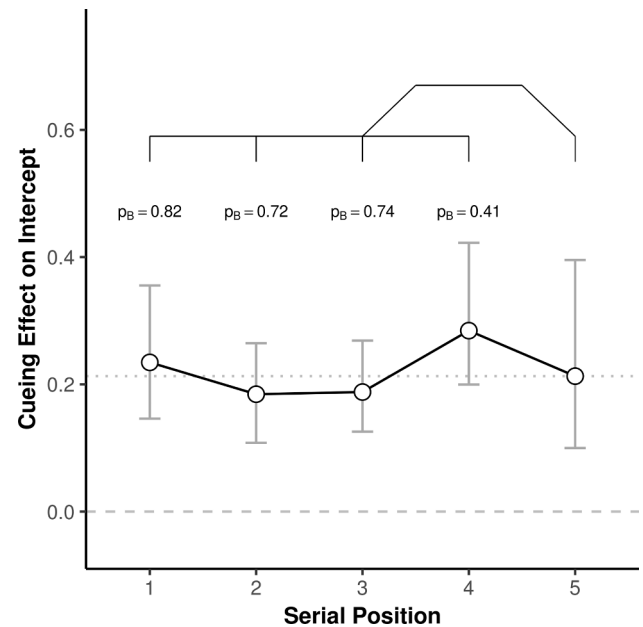


Fig. 8. Median group-level posterior estimates for the retro-cue benefit in seconds for each serial position of Experiment 2.  $p_B$  above serial positions 1–4 denotes the evidence for a difference of the cueing effect between this particular serial position and serial position 5. For serial positions 2 and 3, descriptively the cueing effect was smaller than for serial position 5, and  $p_B$  denotes the evidence for this direction. The dotted line depicts the median cueing effect for the last serial position. The dashed line indicates the absence of a cueing effect. Error bars depict 95% CI.

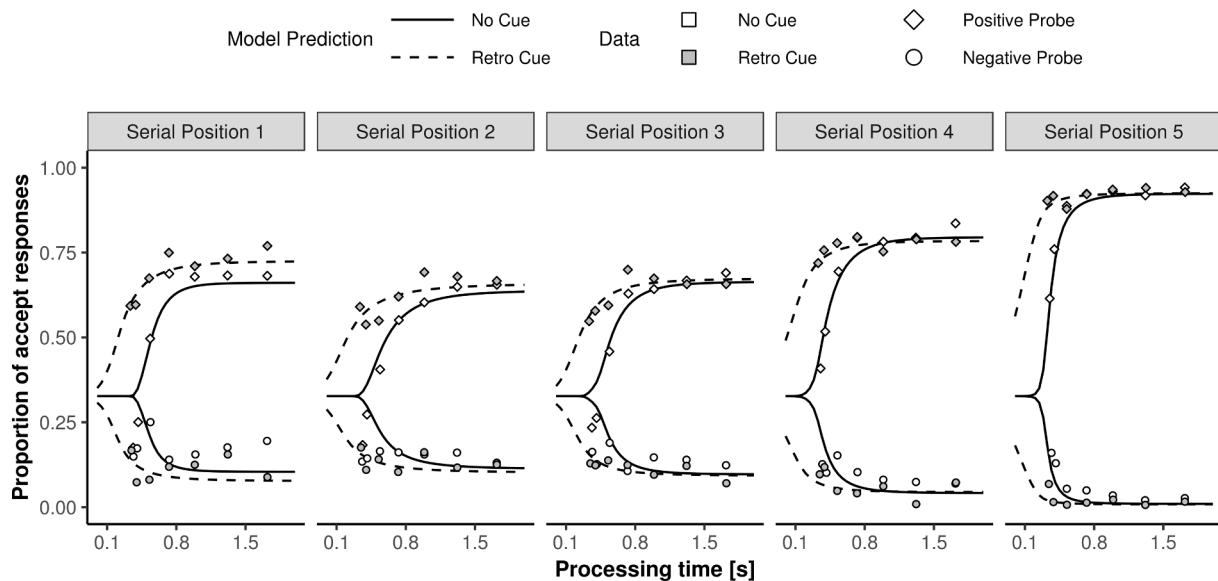
Fig. 7 again suggests the presence of both primacy and recency effects. We tested this assumption using the same analysis as in Experiment 1, that is, we calculated the proportion of posterior samples for which the pattern  $SP1 > SP2 < SP3 < SP4 < SP5$  holds. This proportion was 0.74 for the non-cued probes and 0.72 for the cued probes. Because there exist a total of  $5! = 120$  such orderings, the prior probability of obtaining this ordering is  $\frac{1}{120} = 0.008$ . Therefore, the Bayes factor for this particular ordering is at least  $\frac{0.72}{0.008} = 86$ , providing considerable support for both a primacy and a recency effect on the asymptote parameter.

**4.2.2.4. Bias and correlations.** The median bias parameter was  $-0.67$  [CI =  $-0.90, -0.43$ ] indicating an overall bias to reject probes. We obtained no substantial correlation between individual-level parameters above  $r = 0.17$ , and all CIs included zero.

### 4.3. Discussion

Experiment 2 addressed two concerns of Experiment 1. By presenting probes at the same locations where they had been presented during the study phase, we were first able to associate non-cued negative probes to a serial position, and second, limit the number of comparisons to one for both non-cued and retro-cued probes. For non-cued probes, we again found a last-item benefit, as indicated by a faster intercept for the last item. Moreover, retro-cueing benefits on the intercept parameter were found for all serial positions with no attenuation for the last item, which already benefits from being the last item. This experiment thus provides additional evidence against the proposition that the mechanisms responsible for the last-item benefit are identical to the mechanisms that drive the retro-cue benefit.

Whereas the last-item benefit has predominantly been investigated using verbal material (McElree & Doshier, 1989; McElree, 2006) the retro-cue benefit is most often studied using visual material such as colors and orientations (Souza & Oberauer, 2016). In order to generalize



**Fig. 9.** Observed (points) and predicted (lines) group-level proportions of accept responses for positive (diamonds) and negative (circles) probes for each serial position and cue condition as a function of processing times of Experiment 3. Filled objects connected through a dashed line depict retro-cued probes whereas non-filled objects connected through a solid line depict non-cued probes.

our results to visual working memory, Experiment 3 measured SAT functions for serially presented color patches with or without a retro-cue.

## 5. Experiment 3

In our third experiment, we replicated the procedure of Experiment 2 with colors instead of words as stimuli. Moreover, we addressed two experimental parameters that could possibly limit the generalization of our previous results. First, in Experiments 1 and 2, the last-presented item was always presented in the top left corner of the screen. Although we consider it as highly unlikely, the last-item benefit in these two experiments could be driven by a spatial preference of attention for items presented in the top left corner. Second, because no visual masks were used, after a retention interval of 500 ms some faint iconic traces could still be left that would support the last-item benefit without relying on attentional processes. To address these two issues, in Experiment 3, we varied the spatial position of the last item and extended the retention interval between the offset of the last item, and the onset of the probe or the retro-cue.

### 5.1. Method

#### 5.1.1. Participants

Ten volunteers (9 females, mean age = 26) recruited through the University of Zurich participant volunteer pool participated in ten test sessions each lasting around 1 h, after completion of a one-hour practice session. All participants read and signed an informed consent form before participation.

#### 5.1.2. Materials and procedure

Color patches (diameter = 5.9°) were filled with one of nine distinct colors (RGB codes in parentheses): Dark green (0,63,0), blue (0,0,255), green (0,255,0), yellow (255,255,0), pink (255,50,255), turquoise (90,160,255), orange (255,127,0), brown (127,45,0) and red (255,0,0).

Presentation of color patches occurred in serial order along the virtual circle in clock-wise direction starting at a randomly selected placeholder. The retention interval from the offset of the last color patch until either the onset of a probe (in the no-cue condition), or the onset of a retro-cue (in the retro-cue condition), was set to 1000 ms. All other experimental parameters were identical to Experiment 2.

## 5.2. Results

Analysis will be restricted to positive and new trials. We report all analyses with  $d'$  calculated by relating positive with intrusion probes in Appendix C. These analyses support the same conclusions regarding the interaction of serial position and cue condition. Visual inspection of intrusion trials suggests that a more complex model would have to be fitted in order to account for retrieval dynamics at early response-deadline lags.

### 5.2.1. Response latencies

We excluded 2.08% of all trials due to extreme response latencies. After exclusion of these trials, the response-deadline criterion was met in 91.5% of all trials, which indicates that participants obeyed the response-deadline instructions.

To account for different response latencies in the model-based analysis, we computed processing times for all combinations of participants, response-deadline lags, serial position and cue condition.

### 5.2.2. Model based SAT analysis

We fitted the same model to the data as in Experiment 2. After discarding 1000 warmup samples, we retained 2000 post-warmup samples for each of 4 independent chains, keeping every second sample. Convergence statistics indicated good mixing behavior with  $\hat{R} \leq 1.01$  for all estimated model parameters (Gelman & Rubin, 1992). Visual inspection of MCMC trace plots of the group-level parameters indicated the same. The number of effective samples was above 1000 for all estimated model parameters.

Model fits are depicted in Fig. 9, which compares the mean of the predicted proportions of “accept” responses, aggregated across participants and posterior samples (the lines), to the observed proportions of “accept” responses, aggregated across participants (the points), for positive and negative probes separately. Visual inspection of the model fit shows that whereas positive probes are well captured by the model, this is less the case for new probes. The observed pattern suggests that participants had a strong bias to reject probes, and accumulated evidence for accepting probes over time.

**5.2.2.1. Last item benefit.** We found no credible differences between the experimental conditions for the rate parameter. Furthermore, Fig. 10

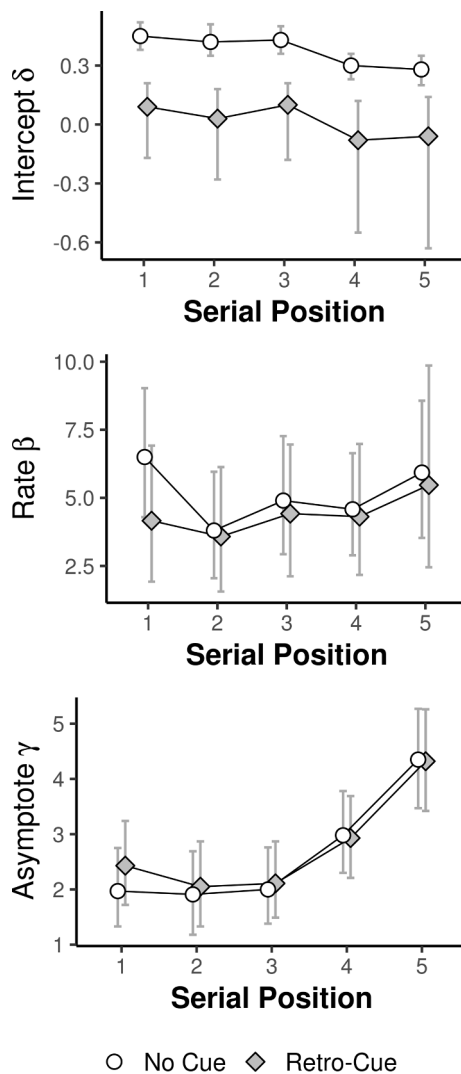


Fig. 10. Median parameter values for each SAT parameter of Experiment 3. Error bars depict the 2.5th and 97.5th percentile of the posterior distribution.

(middle panel) suggests no pattern relevant to our research question. Consequently, we focused on the intercept parameter (Fig. 10, top panel).

To test the last-item benefit, we compared the mean intercept for serial positions 1 to 4 with the mean intercept for serial position 5 across cued and non-cued probes. This comparison revealed no credible difference ( $p_B = 0.18$ , 99.9 ms [−43.4, 399.4]). When we analyzed the last-item benefit for each cue condition separately, we did observe a last-item benefit for non-cued probes ( $p_B = 0.009$ , 121.1 ms [31.6, 217.2]), but not credibly for retro-cued probes ( $p_B = 0.61$ , 79.5 ms [−193.4, 654.2]). Although the posterior medians of both last-item benefits were of similar magnitude, the considerably larger CI of the latter one led to this result.

We additionally investigated all pairwise comparisons for each combination of serial position and cue condition. For non-cued probes, the intercept for serial position 5 was smaller than the intercepts of serial positions 1–3 (all  $p_B < 0.01$ ), but not in comparison to the intercept of serial position 4 ( $p_B = 0.68$ ). The intercept for serial position 4 was credibly smaller than the intercepts for serial positions 1 to 3 (all  $p_B < 0.001$ ). For retro-cued probes, the intercept of serial position 5 was not smaller than the intercepts of serial positions 1 to 4 (all  $p_B > 0.35$ ).

**5.2.2.2. Retro-cue benefit.** We next investigated the effects of presenting a retro-cue. Again, we focused on the intercept parameter as we found

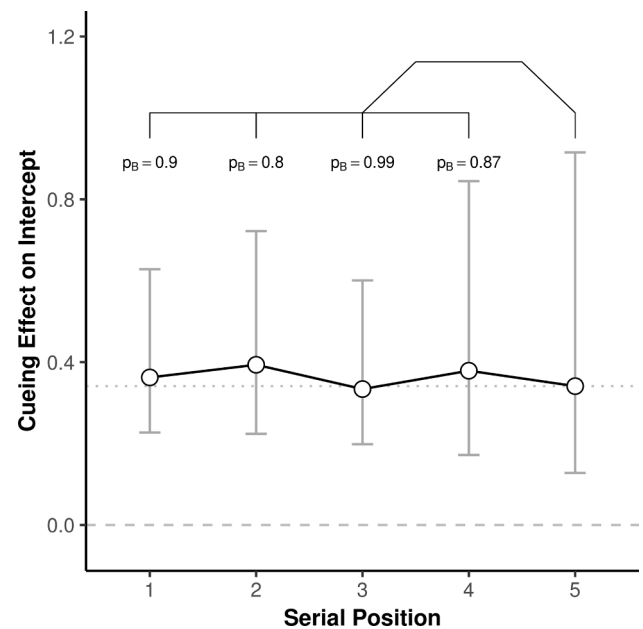


Fig. 11. Median group-level posterior estimates for the retro-cue benefit in seconds for each serial position of Experiment 3.  $p_B$  above serial positions 1–4 denotes the evidence for a difference of the cueing effect between this particular serial position and serial position 5. The dotted line depicts the median cueing effect for the last serial position. The dashed line indicates the absence of a cueing effect. Error bars depict 95% CI.

no credible pairwise difference for the rate parameter. Aggregated across all serial positions, we found a smaller intercept for retro-cued in comparison to non-cued probes ( $p_B < 0.001$ , 382.9 ms [268.8, 546.8]). As can be seen in Fig. 11, the retro-cue benefit was observed for each serial position individually.

To investigate whether the retro-cue benefit is attenuated for serial position 5, we compared the magnitude of the retro-cue benefit of serial position 5 against the mean retro-cue benefit of serial positions 1–4. Again, this comparison yielded no evidence for an attenuation of the retro-cue benefit ( $p_B = 0.83$ ). We also compared the retro-cue benefit of serial position 5 with each earlier serial position individually. Fig. 11 reports the  $p_B$  values for these comparisons and demonstrates that none of these comparisons provides credible evidence for an attenuation of the retro-cue benefit (all  $p_B > 0.80$ ).

**5.2.2.3. Asymptote.** Fig. 10 (bottom panel) shows the estimates of the asymptote parameter of the SAT model. We compared the mean asymptote for all non-cued against all retro-cued probes. We found no evidence for a difference ( $p_B = 0.32$ ). We also compared the non-cue against retro-cue condition for each serial position individually. We found a credible retro-cue benefit for serial position 1 ( $p_B = 0.01$ ), yet we found no such evidence for serial positions 2 to 5 (all  $p_B > 0.25$ ).

Fig. 10 again suggests the presence of both a primacy effect and a recency effect. The proportion of posterior samples for which this pattern (i.e.,  $SP1 > SP2 < SP3 < SP4 < SP5$ ) holds was 0.47 for the non-cued probes and 0.58 for the cued probes. The prior probability of obtaining this ordering is  $\frac{1}{120} = 0.008$ . Therefore, the Bayes factor for this particular ordering is at least  $\frac{0.47}{0.008} = 56$ , providing considerable support for both a primacy and a recency effect on the asymptote parameter.

**5.2.2.4. Bias and correlations.** The median bias parameter was  $-0.45$  [CI =  $-0.70$ ,  $-0.19$ ] indicating a trend towards rejecting the probe. Correlations among group-level parameters were generally low ( $r < 0.12$ ), and all CIs included zero.

### 5.3. Discussion

The purpose of Experiment 3 was to determine whether the conclusions drawn from Experiment 2 can be extended to visual working memory, a lengthened retention interval, and spatially varying locations of the last item. The results indeed closely mirrored those obtained with verbal stimuli in Experiment 2. The crucial comparison of the magnitude of the retro-cue benefit across serial positions clearly shows that the retro-cue benefit was not attenuated for the last item.

One specific aspect of the results in Experiment 3 worth noting is that the last-item benefit on retrieval speed for non-cued probes was found to be extended to the second-to-last item. [McElree \(1998\)](#) reported a similar finding, in which he showed that the retrieval speed benefit extended to three items when they could be semantically grouped with each other. Here, a quarter of all trials involved intrusion probes (e.g., a probe presented in the last item's position, but matching the next-to-last item). In order to correctly reject such probes, it would be fatal to group multiple items together in such a way that they are retrieved and compared to the probe together, rather than individually. We give a possible explanation of this finding when discussing the mechanisms of the last-item benefit in the General Discussion.

In comparison to Experiments 1 and 2, we did not observe a last-item intercept benefit for retro-cued probes. This is likely due to a floor effect. The intercept of all retro-cued conditions was close to, or even below, zero, which left no room for effects of serial position to be detected. In addition, the precision of the parameter estimate for the last two serial positions was extremely poor compared to all other intercept estimates in this manuscript. This further diminished our chances of finding a last-item benefit here. The median posterior estimate of the last-item benefit for cued items was very close to that for non-cued items, supporting our contention that there is no real difference in the size of the last-item benefit between the two cueing conditions.

## 6. General discussion

We set out to investigate whether the last-item benefit and the retro-cue benefit are driven by the same mechanism. If they are, the retro-cue benefit should be attenuated when the retro-cue is directed to the item which already benefits from being presented last. We presented items in serial order and assessed participants' memory with a central (Experiment 1) or location-specific (Experiments 2 and 3) recognition probe. While participants held studied items in working memory, in half of the trials we presented a retro-cue which indicated the item relevant for the subsequent comparison to the probe. To investigate retrieval speed, we measured SAT functions with the response-deadline method. Across three experiments, we found additive last-item benefits and retro-cue benefits on the SAT intercept, which allows us to conclude with confidence that the retro-cue benefit is not attenuated for the last item. Therefore, the retro-cue benefit and the last-item benefit are likely to be driven by different mechanisms.

Our results extend previous research providing indirect evidence for a dissociation between the prioritization of the last item and the retro-cue benefit. [Donkin and Nosofsky \(2012a\)](#) proposed that the last-item benefit reflects the extreme point of a continuous but steep power gradient on memory strength, rather than a special status of the last item. Moreover, [Hu et al. \(2014\)](#) showed that the last list item is especially vulnerable to interference by an irrelevant suffix, whereas retro-cued items are protected from different kinds of visual interference ([Souza et al., 2016](#); [van Moorselaar et al., 2014](#)). Together, these findings provide converging evidence for a distinction between at least two forms of attentional prioritization: Attentional selection through retro-cues and the prioritization by virtue of the last serial position involve different mechanisms. A recent finding by [Kalogeropoulou, Jagadeesh, Ohl, and Rolfs \(2016\)](#) even indicates a potential third form of attentional prioritization. These authors orthogonally manipulated the validity of pre-cues and retro-cues in a

delayed-estimation task of oriented gratings. They found no evidence for an attenuation of the retro-cue benefit when the retro-cued item had already been validly pre-cued. Therefore, attentional prioritization mechanisms involved in pre-cues may be differentiated from mechanisms involved in retro-cues as well.

The interpretation of our results is complicated by the finding that there was only weak evidence for a retro-cue benefit for the last serial position in Experiment 1. However, in support of our interpretation, the magnitude of the retro-cue benefit for the last item in Experiment 1 was very similar to earlier serial positions. Moreover, 85.6% of the posterior mass provide evidence for a retro-cue benefit, meaning that the data still speak more in favor than against such a benefit. Another weakness of our results is that, due to the intercept parameter being at floor, there was no credible last-item benefit for retro-cued items in Experiment 3. Nevertheless, there was still a pronounced retro-cue benefit in the last position, which showed no sign of being smaller than in preceding conditions.

### 6.1. Which retrieval speed parameter reflects the last-item and retro-cue benefits

Across all three experiments, we were able to replicate the last-item benefit reported by [McElree and Doshier \(1989\)](#), generalizing it to location-specific probes and to visual materials. These results are in line with the last-item benefits observed in a location-specific change detection study with the free-response paradigm by [Nosofsky and Donkin \(2016\)](#). Here, the response-deadline paradigm allowed us to study the retrieval dynamics of this benefit: For both, cued and non-cued items, the last item was shown to have a faster intercept parameter than any previous item, whose intercepts were found to be indistinguishable from each other. [McElree and Doshier \(1989\)](#) reported slightly better model performance when the last-item benefit was accounted for by a higher rate in comparison to a faster intercept parameter. In contrast, using more sophisticated modeling techniques, we here show that serial position effects on retrieval speed are best captured by the intercept parameter of the SAT function. Likewise, all observed retro-cue effects on retrieval speed were also consistently accounted for by the intercept, and not the rate parameter of the SAT function. In terms of the SAT function, these results imply that both the last-item benefit and the retro-cue benefit are driven by information being available sooner, rather than a faster accumulation of information once it is available.

In summary, the reported results provide consistent evidence for a distinction of the mechanisms of the last-item and retro-cue benefit on the intercept parameter of the SAT function. In what follows, we speculate as to what mechanisms could drive the last-item and retro-cue benefit.

### 6.2. Mechanisms of retro-cue benefit

Many hypotheses have been put forward to explain the retro-cue benefit (for a review, see [Souza & Oberauer, 2016](#)), including the propositions that retro-cues strengthen item-context bindings ([Rerko & Oberauer, 2013](#)), reduce interference from the test display ([Makovski et al., 2008](#); [Souza, Rerko, Lin, et al., 2014](#); [Souza et al., 2016](#)), and provide a head start of retrieval ([Souza et al., 2016](#)). Our results have implications for the plausibility of these explanations of the retro-cue benefit: Our finding that the retro-cue benefit reflects a shortened intercept parameter of the SAT function fits well with the head-start of retrieval hypothesis.

According to the head-start of retrieval hypothesis, retro-cues allow participants to start retrieving the retro-cued item ahead of the recognition decision-making period. As a consequence, when the probe appears, its comparison to the relevant item in memory can start sooner, and finish sooner. [Shepherdson, Oberauer, and Souza \(2017\)](#) fleshed this hypothesis out and proposed a two-stage model of short-term recognition: During the first stage, one item is retrieved from

working memory, and in the second stage, that item is compared to the probe to arrive at a recognition decision. Support for this two-stage model came from an analysis of response-time distributions with the diffusion model (Ratcliff & McKoon, 2008; Ratcliff, 1978). Retro-cues were found to decrease the model's non-decision time parameter, which reflects the time that is required for non-decisional processes, including the time for retrieving an item from working memory. In addition, retro-cues increased the drift rate, which reflects the quality of information that enters the decision process. In line with the head-start of retrieval hypothesis, Shepherdson and colleagues argued that the retro-cue effect on the non-decision time parameter reflects the retrieval of an item into the focus of attention before the probe is presented (Ratcliff & McKoon, 2008; Sewell, Lilburn, & Smith, 2016). Our model-based analysis of the SAT curves converges with the analysis of response time distributions by Shepherdson et al. (2017): In SAT curves, the intercept reflects the duration of any process preceding the decision process, because during that time no evidence in favor of either response accrues. In contrast, the rate parameter reflects the rate at which evidence in favor of one or the other response accumulates over time. Therefore, the finding of a retro-cue benefit on the intercept confirms the conclusion of Shepherdson et al. (2017) that a retro-cue shortens the duration of a pre-decision process, arguably the retrieval of the relevant item from working memory.

Shepherdson et al. (2017) explained the retro-cue effect on drift rate as reflecting the protection of the cued item against interference by the probe, or other visual information at test (Makovski et al., 2008; Souza et al., 2016). Less interference implies that the comparison of the cued item to the probe provides better information, resulting in a higher rate of evidence accumulation towards the correct response, and as a consequence, faster and more accurate responses. Here we found no evidence that the retro-cue accelerated the rate of accumulation of evidence towards a response, and no evidence that it increased asymptotic accuracy. This renders protection from visual interference a less attractive explanation of the retro-cue benefit in our experiments.

The strengthening hypothesis states that a retro-cue strengthens the retro-cued item and the binding to its context (Rerko & Oberauer, 2013). Strengthened bindings improve access to representations, which is compatible with our findings of retro-cue benefits on retrieval speed. However, such strengthened bindings should also increase the quality of the information retrieved from working memory, and by implication, increase the rate of evidence accumulation, and improve performance at asymptotic levels. Yet, we found no evidence for retro-cue benefits on the rate or the asymptote parameter, which makes the strengthening hypothesis less plausible as an explanation of the retro-cue benefit.

The lack of a retro-cue benefit on asymptotic accuracy in our experiments contrasts with the common finding that retro-cue benefits improve accuracy (in addition to speed) in change-detection experiments (for a review see Souza & Oberauer, 2016). It could be that in regular change-detection experiments, when there is no deadline and participants decide when to respond, participants choose to respond at a point in time where they have not reached their asymptotic level of evidence accumulation. Against this possibility, one experiment by Souza et al. (2016) found that forcing participants to delay their response by one second did not improve change-detection.

Retro-cues allowed participants to direct their eyes to the location of where the probe will appear, whereas in non-cued trials participants could do so only after probe onset. Although this could to some extent explain the retro-cue benefits in Experiments 2 and 3, it cannot explain the retro-cue benefit in Experiment 1, where probes were presented centrally and the eye could fixate the probe location ahead of time regardless of the retro-cue condition. Moreover, Griffin and Nobre (2003) showed in a task similar to ours that retro-cue benefits are obtained even when participants' gaze is held in the center of the screen while probes are presented peripherally. Taken together, even though eye movements were not controlled in our experiments, we are confident that they play at best a minor role in explaining our results.

To conclude, our finding that, consistently across three experiments, the retro-cue only shortened the intercept parameter of the SAT function is best compatible with the assumption that the retro-cue enables a head start for retrieval of the relevant item, thereby shortening a processing stage preceding the decision stage.

### 6.3. Mechanisms of last-item benefit

Can a head-start of retrieval mechanism also account for the last-item benefit? In line with such an explanation, McElree (2006) argues that the last-presented item does not have to be retrieved because it still is in the focus of attention. Therefore, the comparison of the probe to that item, which yields evidence towards one or the other decision, can commence immediately once the probe is presented. However, if indeed both the last-item benefit and retro-cue benefit arose from the same mechanism, we should have observed an attenuated retro-cue benefit for the last item: On this assumption, the last item is already in the focus of attention whether or not a retro-cue points to it, so there is nothing the retro-cue could contribute in addition. Our results rule out this scenario.

An alternative explanation of the last-item benefit is that it reflects the extreme point of a steep power gradient on memory strength (Donkin & Nosofsky, 2012a). Due to the rapid fall of strength, the last item seems to have a special status, when in fact memory strengths for these items simply reflect the power gradient. This proposition can accommodate the finding of Experiment 3 in which we also found a faster intercept for the second-to-last item. If the slope of the power gradient is not as steep between the last two serial positions, the strength of the second-to-last item may still lie well in the non-asymptotic part of the power function.

The power law merely describes the pattern of memory strength with serial position. Possible causes for its pattern involve temporal distinctiveness and retro-active interference. According to a temporal distinctiveness account, retrieval of an item is driven by the uniqueness of its temporal context. The probability of successfully retrieving an item is a function of the distance from all other studied items along a temporal dimension. The last item can be distinguished easiest from all other memory items, because the retention interval following this item renders it more distinct (Brown, Neath, & Chater, 2007). Moreover, the last item benefits from the absence of retro-active interference. All items but the last are interfered with by the presentation of subsequent memory items. The finding of Hu et al. (2014) that the recency effect on accuracy is diminished by a subsequent visual stimulus supports this notion.

### 6.4. Multiple mechanisms of prioritization

Additive benefits of the last list position and of retro-cues indicate that there are at least two forms of attentional prioritization of individual items in working memory. We propose that the last-item benefit reflects a recency gradient on memory strength, which arises from sequential encoding regardless of task demand. In contrast, the retro-cue benefit reflects the fact that the retro-cued item is retrieved, which is a selective and controlled process. This distinction is comparable to the difference between controlled and automatic mechanisms in working memory and perceptual attention. On the one hand, stored memory representations can be prioritized according to task demands in a selective, controlled manner. On the other hand, recency effects reflect an automatic updating process, which occurs regardless of task requirements (Rac-Lubashevsky & Kessler, 2016). The postulation of these multiple mechanisms of prioritization in working memory parallels the distinction of bottom-up and top-down mechanisms of prioritization in perceptual attention (Egeth & Yantis, 1997; though see Awh, Belopolsky, & Theeuwes, 2012, arguing for a third category arising from the person's learning history). In a similar vein, the retro-cue may reflect "top down" attentional selection (i.e., driven by the person's goal and goal-relevant information from the cue), whereas the

last-item benefit reflects a more “bottom-up” prioritization (i.e., to some extent independent of the person's goals, driven by the event sequence in the environment).

In summary, our proposition to distinguish controlled (retro-cues) and automatic (last-item) prioritization in working memory converges with a more global distinction between automatic and controlled processes operating on working-memory contents.

The implications of our results for the concept of a focus of attention in working memory depend on the theoretical perspective taken. In McElree’s (2006) view, the last-item benefit reflects the focus of attention. Accordingly, the retro-cue benefit is not driven by the focus of attention, but instead is a manifestation of different processes. In contrast, in the view of Oberauer (2009), the focus of attention is a selection device for picking out one item from the current set in working memory. Accordingly, the retro-cue benefit reflects the selection of the cued item into the focus of attention, whereas the last-item benefit reflects different processes.

6.5. Conclusion

The present data provide evidence for two forms of attentional

prioritization of single items in working memory. We propose that the retro-cue benefit may reflect the operation of a goal-driven selection mechanism in working memory, whereas the last-item benefit may be a result of the unequal distribution of memory strength over list positions, resulting from the updating of working memory.

Acknowledgements

We thank Laura Rerko for her help in the planning of this study. We are grateful to Peter Shepherdson, Jacob Feldman, and three anonymous reviewers for valuable comments on an earlier draft. We also thank Sonja Peteranderl and Danielle Pessach for helping with data collection.

Parts of this paper were presented at the 7th European Working Memory Symposium in Edinburgh, UK, in September 2014, at the Annual Meeting of the Psychonomic Society in Boston, USA, in November 2016, and at the Conference of Experimental Psychologists in Dresden, Germany, in March 2017.

This research was supported by a grant from the Swiss National Foundation (143333) to K. Oberauer.

Appendix A. Negative probes experiment 1

Median parameter values and their 95% CI for each parameter of the SAT function for the negative probes are depicted in Fig. A.1.

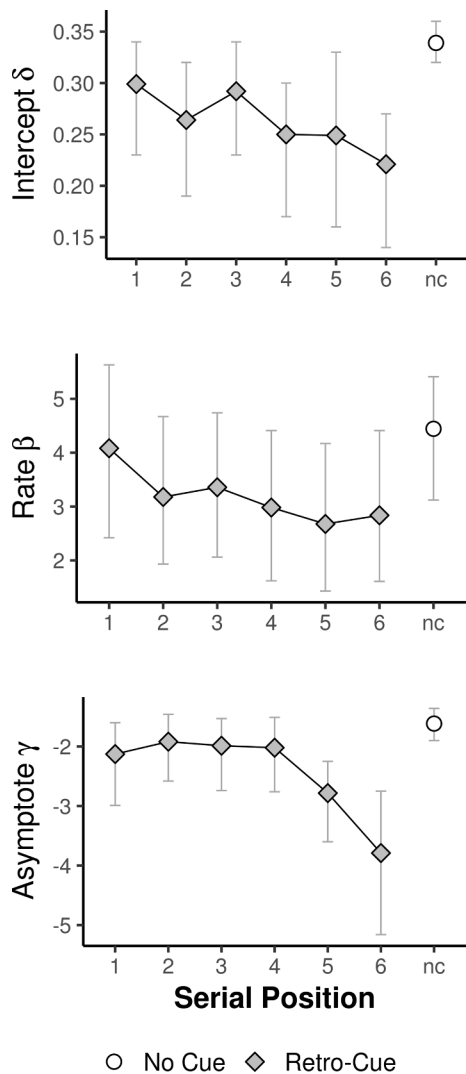


Fig. A.1. Median parameter values for each SAT parameter for negative probes of Experiment 1. Error bars depict the 2.5th and 97.5th percentile of the posterior distribution.

### A.1. Negative probes – last item benefit

We obtained no credible differences between the experimental conditions for the rate parameter. Furthermore, Fig. A.1 (middle panel) suggests no pattern relevant to our research question. Consequently, we focused on the intercept parameter (Fig. A.1, top panel).

New probes that were presented in the center of the screen cannot be associated with any serial position if no retro-cue was presented. Hence, no last-item benefit can be computed for these probes. To test the last-item benefit for retro-cued new probes, we compared the mean intercept for serial positions 1–5 with the mean intercept for serial position 6 for retro-cued probes. This comparison indicated that the retro-cued serial position 6 did not have a smaller intercept than serial positions 1–5 ( $p_B = 0.12$ , 48.0 ms [−11.5, 138.2]).

Moreover, aggregated across all cued serial positions, we found a smaller intercept for cued than for non-cued probes ( $p_B < 0.001$ , 78.4 ms [47.7, 111.8]).

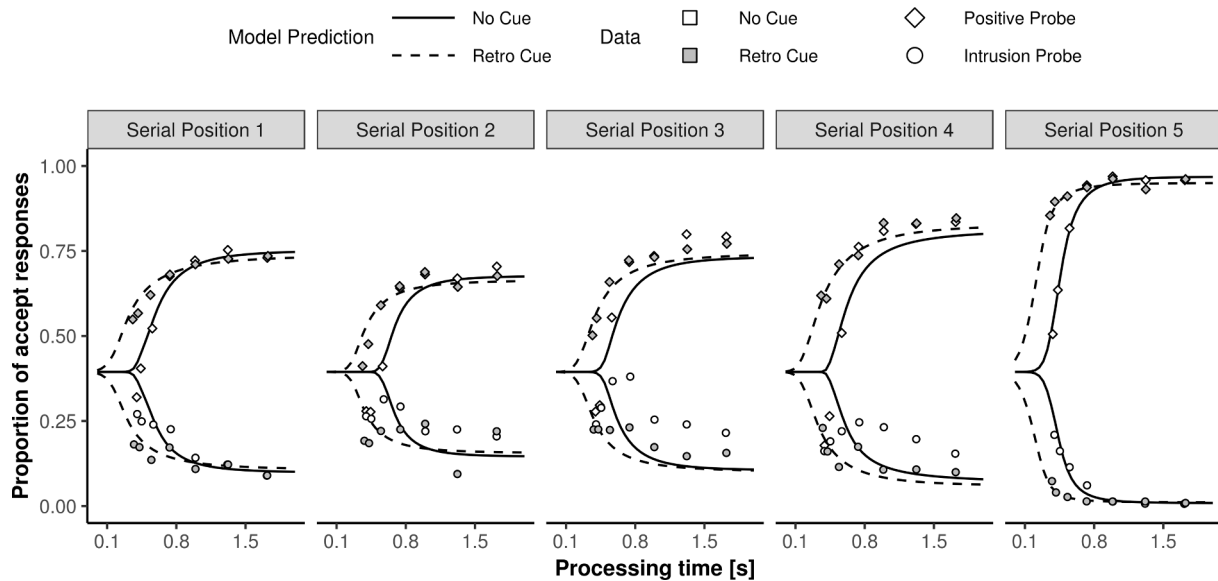
### A.2. Negative probes – asymptote

Fig. A.1 (bottom panel) shows the estimates of the asymptote parameter for negative probes of the SAT model. We found credible retro-cue benefits for the asymptote, when all non-cued probes were compared to all retro-cued probes ( $p_B < 0.001$ , 0.85 [0.46, 1.25]). When all retro-cued probes were compared to all non-cued probes separately, we found credible retro-cue benefits for serial positions 5 ( $p_B = 0.001$ ) and 6 ( $p_B < 0.001$ ), but not for serial positions 1–4 (all  $p_B > 0.095$ ).

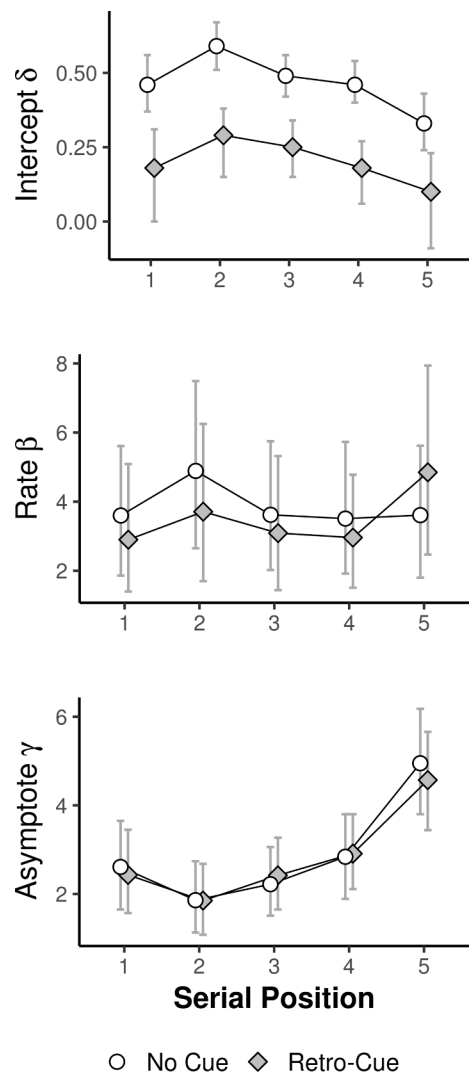
Fig. A.1 suggests the presence of both primacy and recency effects. We tested this assumption using the same analysis as for positive probes in Experiment 1, that is we calculated the proportion of posterior samples for which the pattern  $SP1 < SP2 > SP3 > SP4 > SP5 > SP6$  holds. This proportion was 0.17 for the cued probes. Therefore, the Bayes factor for this particular ordering is at least  $\frac{0.17}{0.001} = 122$ , providing considerable support for a recency effect on the asymptote parameter.

## Appendix B. Intrusion probes experiment 2

After discarding 1000 warmup samples, we retained 1000 post-warmup samples for each of 4 independent chains, keeping every second sample. Convergence statistics indicated good mixing behavior with  $\hat{R} \leq 1.05$  for all estimated model parameters (Gelman & Rubin, 1992). Visual inspection of MCMC trace plots of the group-level parameters indicated the same. The number of effective samples was above 100 for all estimated model parameters. Model fits are depicted in Fig. B.1, which compares the mean of the predicted proportions of “accept” responses, aggregated across participants and posterior samples (the lines), to the observed proportions of “accept” responses, aggregated across participants (the points), for positive and intrusion probes separately. Visual inspection of the model fit shows that the model struggled to account for the retrieval dynamics of intrusion probes. Median parameter values and their 95% CIs are displayed in Fig. B.2. Finally, we compared the retro-cue benefit of serial position 5 with each earlier serial position individually. Fig. B.3 reports the  $p_B$  values for these comparisons and shows that none of these comparisons provides credible evidence for an attenuation of the retro-cue benefit (smallest  $p_B = 0.48$ ).



**Fig. B.1.** Observed (points) and predicted (lines) group-level proportion of accept responses for positive (diamonds) and intrusion (circles) probes for each serial position and cue condition as a function of processing times (response-deadline lag plus the individual mean response time per experimental condition) of Experiment 2. Filled objects connected through a dashed line depict retro-cued probes whereas non-filled objects connected through a solid line depict non-cued probes.



**Fig. B.2.** Median parameter values for each SAT parameter of Experiment 2 (positive vs. intrusion probes). Error bars depict the 2.5th and 97.5th percentile of the posterior distribution.

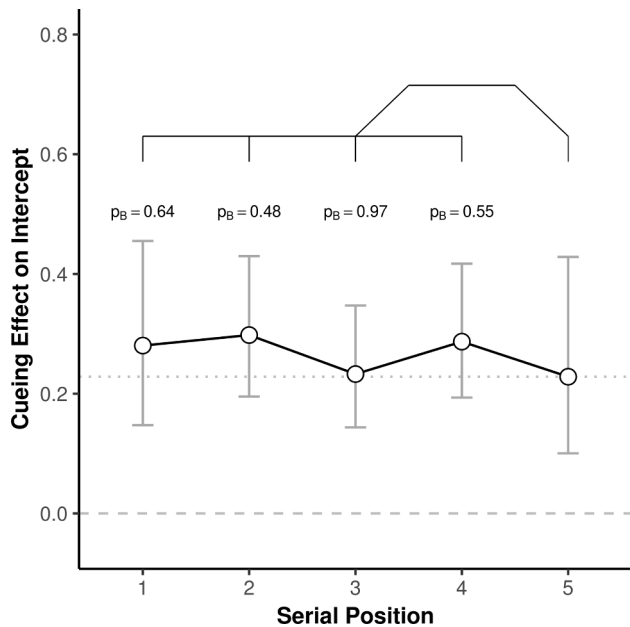


Fig. B.3. Median group-level posterior estimates for the retro-cue benefit in seconds for each serial position of Experiment 2 (positive vs. intrusion probes).  $p_B$  above serial positions 1–4 denotes the evidence for a difference of the cueing effect between this particular serial position and serial position 5. The dotted line depicts the median cueing effect for the last serial position. The dashed line indicates the absence of a cueing effect. Error bars depict 95% CI.

Appendix C. Intrusion probes experiment 3

After discarding 1000 warmup samples, we retained 1000 post-warmup samples for each of 4 independent chains, keeping every second sample. Convergence statistics indicated good mixing behavior with  $\hat{R} \leq 1.01$  for all estimated model parameters (Gelman & Rubin, 1992). Visual inspection of MCMC trace plots of the group-level parameters indicated the same. The number of effective samples was above 180 for all estimated model parameters. Model fits are depicted in Fig. C.1, which compares the mean of the predicted proportions of “accept” responses, aggregated across participants and posterior samples (the lines), to the observed proportions of “accept” responses, aggregated across participants (the points), for positive and intrusion probes separately. Visual inspection of the model fit shows that the model struggled to account for the retrieval dynamics of uncued intrusion probes. Median parameter values and 95% CIs are depicted in Fig. C.2. Finally, we compared the retro-cue benefit of serial position 5 with each earlier serial position individually. Fig. C.3 reports the  $p_B$  values for these comparisons and shows that none of these comparisons provides credible evidence for an attenuation of the retro-cue benefit (smallest  $p_B = 0.49$ ).

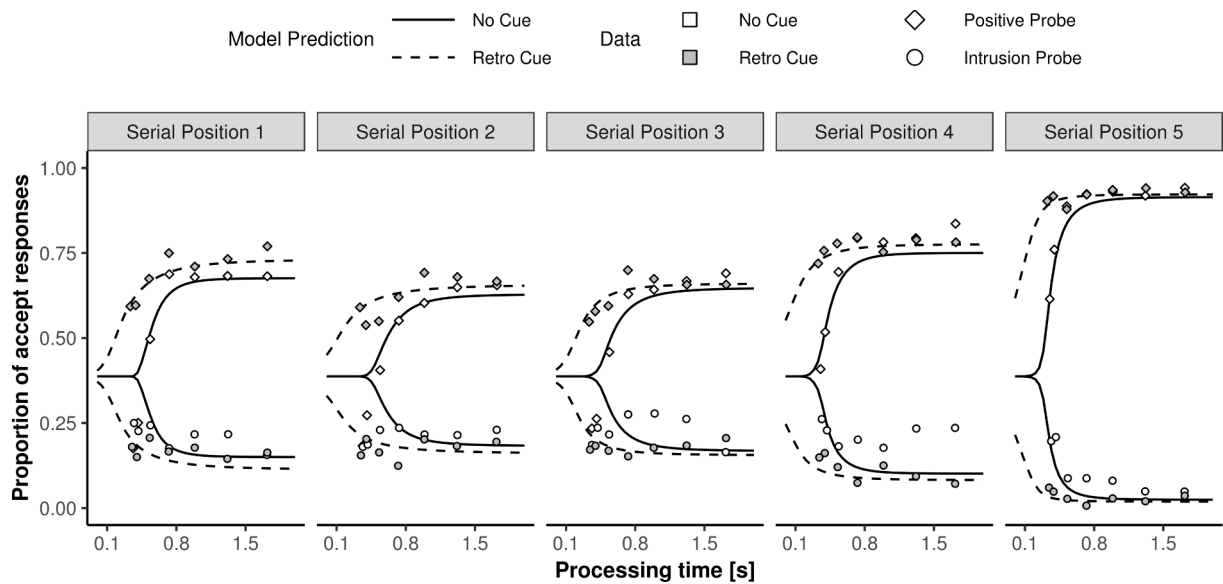
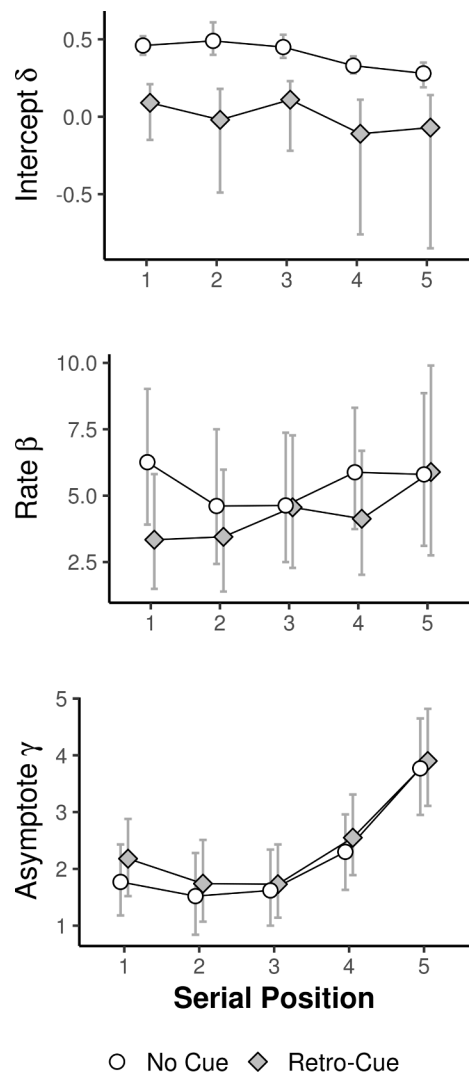
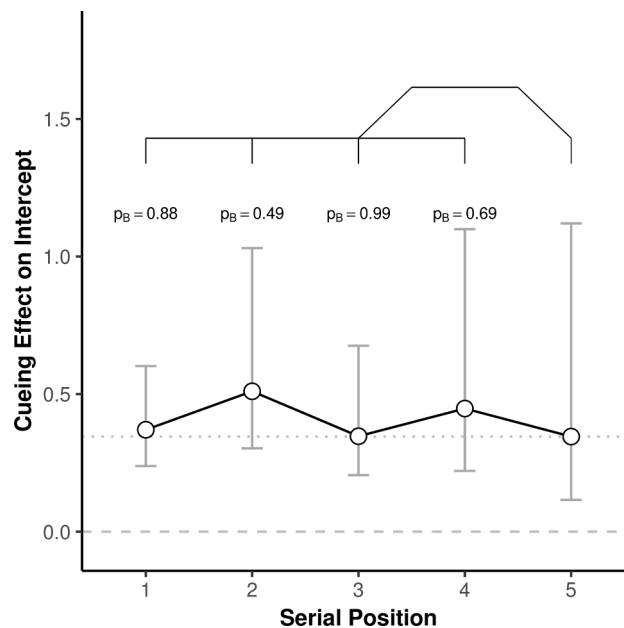


Fig. C.1. Observed (points) and predicted (lines) group-level proportion of accept responses for positive (diamonds) and intrusion (circles) probes for each serial position and cue condition as a function of processing times (response-deadline lag plus the individual mean response time per experimental condition) of Experiment 3. Filled objects connected through a dashed line depict retro-cued probes whereas non-filled objects connected through a solid line depict non-cued probes.



**Fig. C.2.** Median parameter values for each SAT parameter of Experiment 3 (positive vs. intrusion probes). Error bars depict the 2.5th and 97.5th percentile of the posterior distribution.



**Fig. C.3.** Median group-level posterior estimates for the retro-cue benefit in seconds for each serial position of Experiment 3 (positive vs. intrusion probes).  $p_B$  above serial positions 1–4 denotes the evidence for a difference of the cueing effect between this particular serial position and serial position 5. The dotted line depicts the median cueing effect for the last serial position. The dashed line indicates the absence of a cueing effect. Error bars depict 95% CI.

## Appendix D. Supplementary material

The raw trial-by-trial data, Stan model codes, stanfit objects, and the R analysis scripts for all experiments can be accessed in the Open Science Framework: <https://osf.io/6apd9/>.

## References

- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, 16(8), 437–443.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Brown, G. D., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, 114(3), 539.
- Burnham, K. P., & Anderson, D. R. (2003). *Model selection and multimodel inference: A practical information-theoretic approach*. Springer Science & Business Media.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 20.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1).
- Cowan, N. (1998). *Attention and memory: An integrated framework*. Oxford University Press.
- Cowan, N. (2011). The focus of attention as observed in visual working memory tasks: Making sense of competing claims. *Neuropsychologia*, 49(6), 1401–1406.
- 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(3), 363–394.
- Egeth, H. E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Review of Psychology*, 48(1), 269–297.
- Gabry, J., Simpson, D., Vehtari, A., Betancourt, M., & Gelman, A. (2017). Visualization in Bayesian workflow. ArXiv Preprint ArXiv:1709.01449.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2014). *Bayesian data analysis* (Vol. 2). FL, USA: Chapman & Hall/CRC Boca Raton.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models* (Vol. 1). New York, NY, USA: Cambridge University Press.
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 457–472.
- Gelman, A., Stern, H. S., Carlin, J. B., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. Chapman and Hall/CRC.
- Göthe, K., & Oberauer, K. (2008). The integration of familiarity and recollection information in short-term recognition: Modeling speed-accuracy trade-off functions. *Psychological Research*, 72(3), 289–303.
- Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience*, 15(8), 1176–1194.
- Hu, Y., Hitch, G. J., Baddeley, A. D., Zhang, M., & Allen, R. J. (2014). Executive and perceptual attention play different roles in visual working memory: Evidence from suffix and strategy effects. *Journal of Experimental Psychology: Human Perception and Performance*, 40(4), 1665.
- Kalogeropoulou, Z., Jagadeesh, A. V., Ohl, S., & Rolfs, M. (2016). Setting and changing feature priorities in visual short-term memory. *Psychonomic Bulletin & Review*, 1–6.
- Katahira, K. (2016). How hierarchical models improve point estimates of model parameters at the individual level. *Journal of Mathematical Psychology*, 73, 37–58.
- Kellen, D., & Klauer, K. C. (2018). Elementary signal detection and threshold theory. *Stevens' handbook of experimental psychology and cognitive neuroscience* (pp. 1–39). New York: John Wiley & Sons Inc.
- Klauer, K. C. (2010). Hierarchical multinomial processing tree models: A latent-trait approach. *Psychometrika*, 75(1), 70–98.
- Landman, R., Spekreijse, H., & Lamme, V. A. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, 43(2), 149–164.
- Liu, C. C., & Smith, P. L. (2009). Comparing time-accuracy curves: Beyond goodness-of-fit measures. *Psychonomic Bulletin & Review*, 16(1), 190–203.
- Macmillan, N. A. (2002). Signal detection theory. In J. Wixted, & H. Pashler (Eds.), *Stevens' handbook of experimental psychology: Vol. 4*. New York: John Wiley.
- Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(2), 369.
- McElree, B. (1998). Attended and non-attended states in working memory: Accessing categorized structures. *Journal of Memory and Language*, 38(2), 225–252.
- McElree, B. (2006). Accessing recent events. *Psychology of Learning and Motivation*, 46, 155.
- McElree, B., & Doshier, B. A. (1989). Serial position and set size in short-term memory: The time course of recognition. *Journal of Experimental Psychology: General*, 118(4), 346.
- Mızrak, E., & Öztekin, I. (2016a). Relationship between emotion and forgetting. *Emotion*, 16(1), 33.
- Mızrak, E., & Öztekin, I. (2016b). Working memory capacity and controlled serial memory search. *Cognition*, 153, 52–62. <https://doi.org/10.1016/j.cognition.2016.04.007>.
- Monsell, S. (1978). Recency, immediate recognition memory, and reaction time. *Cognitive Psychology*, 10(4), 465–501.
- Niklaus, M., Nobre, A. C., & van Ede, F. (2017). Feature-based attentional weighting and spreading in visual working memory. *Scientific Reports*, 7, 42384.
- Nosofsky, R. M., & Donkin, C. (2016). Response-time evidence for mixed memory states in a sequential-presentation change-detection task. *Cognitive Psychology*, 84, 31–62.
- Nosofsky, R. M., Little, D. R., Donkin, C., & Fific, M. (2011). Short-term memory scanning viewed as exemplar-based categorization. *Psychological Review*, 118(2), 280.
- Oberauer, K. (2003). Selective attention to elements in working memory. *Experimental*

- Psychology*, 50(4), 257.
- Oberauer, K. (2008). How to say no: Single- and dual-process theories of short-term recognition tested on negative probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(3), 439.
- Oberauer, K. (2009). Design for a working memory. *Psychology of Learning and Motivation*, 51, 45–100.
- Oberauer, K., & Hein, L. (2012). Attention to information in working memory. *Current Directions in Psychological Science*, 21(3), 164–169.
- Oberauer, K., Lewandowsky, S., Awh, E., Brown, G. D. A., Conway, A., Cowan, N., ... Ward, G. (2018). Benchmarks for models of short term and working memory. *Psychological Bulletin*.
- Öztekin, I., & McElree, B. (2010). Relationship between measures of working memory capacity and the time course of short-term memory retrieval and interference resolution. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(2), 383.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- R Core Team (2014). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Rac-Lubashevsky, R., & Kessler, Y. (2016). Dissociating working memory updating and automatic updating: The reference-back paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(6), 951.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922.
- Rerko, L., & Oberauer, K. (2013). Focused, unfocused, and defocused information in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(4), 1075.
- Sewell, D. K., Lilburn, S. D., & Smith, P. L. (2016). Object selection costs in visual working memory: A diffusion model analysis of the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(11), 1673.
- Shepherdson, P., Oberauer, K., & Souza, A. S. (2017). Working memory load and the retro-cue effect: A diffusion model account. *Journal of Experimental Psychology: Human Perception and Performance*.
- Souza, A. S., & Oberauer, K. (2016). In search of the focus of attention in working memory: 13 years of the retro-cue effect. *Attention, Perception, & Psychophysics*, 78(7), 1839–1860.
- Souza, A. S., Rerko, L., Lin, H.-Y., & Oberauer, K. (2014). Focused attention improves working memory: Implications for flexible-resource and discrete-capacity models. *Attention, Perception, & Psychophysics*, 1–23.
- Souza, A. S., Rerko, L., & Oberauer, K. (2014). Unloading and reloading working memory: Attending to one item frees capacity. *Journal of Experimental Psychology: Human Perception and Performance*, 40, 1237–1256.
- Souza, A. S., Rerko, L., & Oberauer, K. (2016). Getting more from visual working memory: Retro-cues enhance retrieval and protect from visual interference. *Journal of Experimental Psychology: Human Perception and Performance*, 42(6), 890.
- Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers*, 31(1), 137–149.
- Sternberg, S. (1966). High-speed scanning in human memory. *Science*, 153(3736), 652–654. <https://doi.org/10.1126/science.153.3736.652>.
- van Ede, F., Niklaus, M., & Nobre, A. C. (2016). Temporal expectations guide dynamic prioritization in visual working memory through attenuated alpha oscillations. *Journal of Neuroscience*, 2272–2316.
- van Moorselaar, D., Gonseli, E., Theeuwes, J., & Olivers, C. N. (2014). The time course of protecting a visual memory representation from perceptual interference. *Frontiers in Human Neuroscience*, 8.
- Wickelgren, W. A., Corbett, A. T., & Doshier, B. A. (1980). Priming and retrieval from short-term memory: A speed accuracy trade-off analysis. *Journal of Verbal Learning and Verbal Behavior*, 19(4), 387–404.