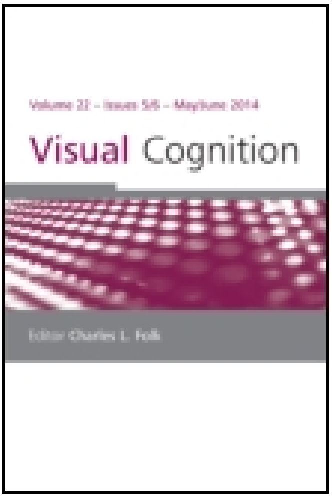

Visual Cognition
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/pvis20

Quality and accessibility of visual working memory during cognitive control of attentional guidance: A Bayesian model comparison approach

Emma Wu Dowd[a], Anastasia Kiyonaga[a], Jeffrey M. Beck[a] & Tobias Egner[a]

[a] Department of Psychology and Neuroscience, Center for Cognitive Neuroscience, Duke University, Durham, NC, USA

Published online: 13 Feb 2015.

Visual Cognition, 2015
Vol. 23, No. 3, 337–356, http://dx.doi.org/10.1080/13506285.2014.1003631

Quality and accessibility of visual working memory during cognitive control of attentional guidance: A Bayesian model comparison approach

Emma Wu Dowd*, Anastasia Kiyonaga, Jeffrey M. Beck, and Tobias Egner

Department of Psychology and Neuroscience, Center for Cognitive Neuroscience, Duke University, Durham, NC, USA

(*Received 15 December 2014; accepted 26 December 2014*)

Working memory (WM) can guide visual attention toward memory-matching objects. This influence of WM on attention can be modulated by cognitive control, such that attentional guidance is strategically suppressed or enhanced depending on whether WM contents are reliably hurtful or helpful for the current task. Cognitive control over memory-based guidance has been hypothesized to operate via some modulation of the WM representation itself, but it is unclear whether this modulation affects representational quality (i.e., how precise is it?) or accessibility of WM content (i.e., how easily is it remembered or forgotten?). Using probabilistic model fitting and Bayesian model comparison techniques, we show that cognitive control over memory-based guidance impacts the probability of remembering, but not the precision of, items in WM. These findings suggest that the WM-attention interaction may depend on distinct functional states in WM, which are in turn characterized by how easily an item is remembered or forgotten.

***Keywords**:* Visual attention; Working memory; Cognitive control; Bayesian model comparison.

The contents of working memory (WM) tend to guide visual attention toward memory-matching objects (see Olivers, Peters, Houtkamp, & Roelfsema, 2011; Soto, Hodsoll, Rotshtein, & Humphreys, 2008). In a typical dual-task paradigm assessing the relationship between WM and attention, participants are asked to remember an item (e.g., a coloured shape) while performing an intervening

Please address all correspondence to Emma Wu Dowd, Center for Cognitive Neuroscience, Duke University, Box 90999, Durham, NC 27708, USA. E-mail: emma.wudowd@duke.edu
We thank Ada Aka for collecting data.

visual search task. Critically, the memory item might reappear in the search display at the same location as the search target (valid) or as an irrelevant distractor (invalid). Compared to a neutral condition, in which the memory item does not reappear in the array, visual search performance reveals benefits (i.e., faster search times) for valid trials and costs (i.e., slower search times) for invalid trials, indicating that items matching the contents of WM attract attention (e.g., Downing, 2000; Olivers, Meijer, & Theeuwes, 2006; Soto, Heinke, Humphreys, & Blanco, 2005; but see Downing & Dodds, 2004; Houtkamp & Roelfsema, 2006; Woodman & Luck, 2007).

Intriguingly, recent studies have found that the WM-attention linkage can be strategically suppressed or enhanced, depending on whether WM contents are reliably hurtful or helpful for attentional task performance (Carlisle & Woodman, 2011; Kiyonaga, Egner, & Soto, 2012). By manipulating the probabilities of trial conditions, these studies demonstrated that participants could amplify the benefits of valid trials when the memory item reliably predicted the location of a target (i.e., even faster responses) and dampen the costs of invalid trials when the memory item was reliably associated with distractors (i.e., less slowed responses).

Cognitive control over attentional guidance by WM has been hypothesized to operate via some modulation of the WM representation itself (Kiyonaga et al., 2012); however, the exact way in which the maintenance of the memory item is altered to produce these strategic effects is presently unknown. Recent models of visual WM suggest that a limited pool of cognitive resources (or a limited number of resource slots) can be distributed to items in memory, such that WM representations might vary in quality, or precision (Bays & Husain, 2008; Zhang & Luck, 2008). By testing memory with continuous report tasks (Wilken & Ma, 2004), in which participants recall a target from a continuous spectrum of responses, memory responses can be fit with a probabilistic mixture model (see Bays, Catalao, & Husain, 2009; Zhang & Luck, 2008) that dissociates performance into two main components: (1) the representational quality, or precision, with which a target is recalled, given that it is remembered, and (2) the accessibility of that representation, or the probability of remembering that target at all (as opposed to randomly guessing or incorrectly reporting another item from the memory array).

Based on this influential distinction between precision and accessibility, the current study investigated how strategic enhancement or suppression of the interaction between WM and attention affects the WM representation, examining whether cognitive control of memory-based guidance impacts the quality of WM, the accessibility of WM, or both. Specifically, cognitive control over how WM influences attention could alter the precision of WM representations, which may in turn determine performance. For example, memory-based capture could depend on the specificity of the match between WM contents and stimuli in the visual field; strategically enhancing or dampening representational quality could affect WM-attention interactions, such that a more precise WM representation might mean better performance in a matching task, whereas a less precise representation might lead to more behavioural variability in the task. Alternatively, cognitive control over WM could change the accessibility of WM. For example, an "enhanced" (prioritized) WM item might be more easily accessed by WM retrieval processes and thus more likely to influence attention, whereas a suppressed WM item might be stored in a manner that limits its retrieval from WM, and thus its impact on attention (e.g., Olivers et al., 2011). These processes could also work conjointly.

A number of recent studies speak to the question of how knowledge about the task-relevance of a memory item affects the quality of that memory representation. One study used a "directed-forgetting" paradigm, in which certain items in WM were retrospectively cued to be forgotten (i.e., task-irrelevant) (Williams, Hong, Kang, Carlisle, & Woodman, 2013). The remaining items in WM (i.e., those that were not cued to be forgotten) had higher precision relative to conditions in which no items were directed to be forgotten, suggesting that the quality of memory representations could be flexibly and strategically controlled (see also Pertzov, Bays, Joseph, & Husain, 2013). In contrast, other studies found no differences in precision between retrospectively cued and uncued items, but rather only advantages in the probability of remembering the cued item (Hollingworth & Hwang, 2013; Murray, Nobre, Clark, Cravo, & Stokes, 2013; Souza, Rerko, Lin, & Oberauer, 2014).

Two studies have also found that the precision of a memory representation is not correlated to attentional capture by that matching item (Hollingworth & Hwang, 2013; van Moorselaar, Theeuwes, & Olivers, 2014). However, these previous studies did not manipulate the relevance of the memory item in relation to the subsequent attentional task. Hollingworth and Hwang (2013, Experiment 3) used a retrospective cue to indicate memory relevance before search, and the subsequent search task could contain either a singleton distractor that matched the uncued (memory-irrelevant) item or a distractor that was unrelated to any feature in WM, but never the cued (memory-relevant) item. While relatively precise representations of the uncued item were *not* sufficient to drive attentional capture by uncued-matching distractors, the study design prevented the testing of whether precision for the cued item was related to attentional capture by cued-matching distractors. Van Moorselaar et al. (2014, Experiment 3) also found no differences in precision for memory items that did or did not capture attention; however, this study presented the retro-cue *after* the search task, such that there was no manipulation of memory-relevance before search. Furthermore, neither study analyzed the probability of remembering the memory item in relation to attentional capture. In the current study, we manipulate the relevance of the memory item for the subsequent search task, in which the memory-matching item guides attention toward or away from the target, allowing us to examine how strategic modulation of that interaction affects the WM representation itself, both in terms of quality and accessibility.

Based on these findings from previous studies, it remains unclear how cognitive control over memory-based attentional guidance would affect the WM representation itself. The disagreement over whether the task-relevance of a WM representation affects its quality and/or accessibility may be explained in several ways. For example, the data could simply be too variable across different task paradigms and different samples. There also may not have been a well-grounded, uniform statistical test applied to appropriately analyze the dataset. Specifically, these previous studies shared a common analytic approach: fitting probabilistic models to the data from each participant, then combining parameter estimates across participants and comparing group means with *t*-tests or ANOVAs. While this seems like a straightforward way to analyze parameter estimates, this approach fails to account for the reliability of and correlations between parameter estimates of each subject (Kruschke, 2010; Oberfeld & Franke, 2013; Suchow, Brady, Fougnie, & Alvarez, 2013). Moreover, an ANOVA is only capable of identifying a change in parameter values of a given model and is not capable of determining if an entirely different model (perhaps one with more parameters) provides a better description of the behaviour in different task conditions. This is particularly problematic in behavioural studies as different subjects may employ a wide range of strategies, each of which might have a unique parameterization, making any simple comparison impossible.

To overcome these limitations, the current study employed Bayesian model comparison. While more computationally intensive, this approach allowed us to compare a large number of possible behavioural models to determine strategies employed by subjects and identify the various sources of suboptimal behaviour (i.e., diminished quality vs. accessibility of the WM representation). In Bayesian statistics, for a particular model M with parameter vector θ_M, the marginal likelihood of the data from subject k, D^k, is calculated by integrating over the parameters of model M:

$$p\left(D^k|M\right) = \int p\left(D^k|\theta_M, M\right)p(\theta_M|M)d\theta_M \tag{1}$$

Bayesian model comparison relies on these marginal data likelihoods (i.e., the model evidences) to compute how likely a particular model is compared to other models. For example, given either model $M = 1$ or $M = 2$, the posterior probability that $M = 1$ is the model used by subject k is:

$$p\left(M = 1|D^k\right) = \frac{p\left(D^k|M = 1\right)}{p\left(D^k|M = 1\right)p(M = 1) + p\left(D^k|M = 2\right)p(M = 2)} \tag{2}$$

where $p(M = i)$ is the prior probability of model i. At the individual subject level, this approach allows us to compute the probability that a given subject adopts any modelled strategy in response to a given change in task conditions. This approach has the added advantage of automatically taking into account parameter uncertainty and model complexity (Myung & Pitt, 1997). Moreover, consideration of a hierarchical model in which the prior probabilities over models, $p_i = p(M = i)$, are themselves sampled from a Dirichlet distribution, provides us with a natural way to combine data across multiple subjects. This is because inferring a distribution over p_i given data from all of the subjects provides an estimate of the probability that a new (or randomly chosen) subject's behaviour will be generated by each of the given models.

Here, we adopted a hypothesis-driven variant of factorial model comparison (see van den Berg, Awh, & Ma, 2014), by generating 20 different models, each of which assumes that a change in task conditions will have a particular effect. The 20 models differ only in the specification of prior probability over parameters used in each task condition. For example, one model may assume that the probability of remembering an item in WM is unaffected by task condition, while another model may assume that it increases in one condition and decreases in another. In this manner, the current study employed model comparison techniques to investigate how cognitive control over the interaction between WM and attention impacts the contents of WM—specifically, we examined how enhancing or suppressing memory-based guidance is related to differences in quality and/or accessibility of WM. In two experiments, we used slight variations of a dual-task paradigm that combined WM and visual search. The main measure of interest, WM performance, was measured with a continuous colour recall task, in which participants recalled the memory colour value from a spectrum of responses. Memory errors were fitted with a probabilistic model to estimate both the precision of the memory representation, as well as the probability that the item was retained in memory at all (see Zhang & Luck, 2008). Using factorial model comparison, we showed that strategic modulation of WM-attention interactions impacts WM accessibility rather than WM quality.

GENERAL METHODS

Both experiments in this study used a similar paradigm, detailed here. Any differences from this paradigm are noted for each experiment.

Apparatus

Both experiments were conducted on a Dell Optiplex 960 computer, running Windows XP, and were programmed in Matlab using Psychophysics Toolbox, Version 3.0 (Brainard, 1997). Participants viewed the experimental displays on a

LCD monitor with a refresh rate of 60 Hz and screen resolution of 1280 × 1024 pixels at an approximate distance of 60 cm.

Procedure

Figure 1 depicts an example trial sequence. Each trial began with the presentation of a central fixation cross (black, stroke width = .1°, subtending .5° × .5°) for 1000 ms. The cross was followed by a blank screen for 500 ms, then by the memory cue item, a single coloured circle (stroke width = .3°, subtending 2.8° in diameter), which was presented at the centre of the screen for 250 ms. Participants were instructed to remember the colour of the memory item for a subsequent test. The memory cue was followed by a blank delay for 2000 ms, and then by the search array for 100 ms. The search array consisted of three coloured circles (stroke width = .3°, each 2.8° diameter) at the corners of an imaginary triangle, with each corner approximately 5° from central fixation. Each circle contained a line of 1.2° length. Two of the lines (i.e., search distractors) were vertical, while one line (i.e., the search target) was tilted 15° to the left or right. Each target location and orientation occurred equally often and in a randomized order. Participants were instructed to indicate the orientation of the line via keypress to respond "left" or "right" tilted. The search array was followed by a blank screen until a keypress was recorded. After the search

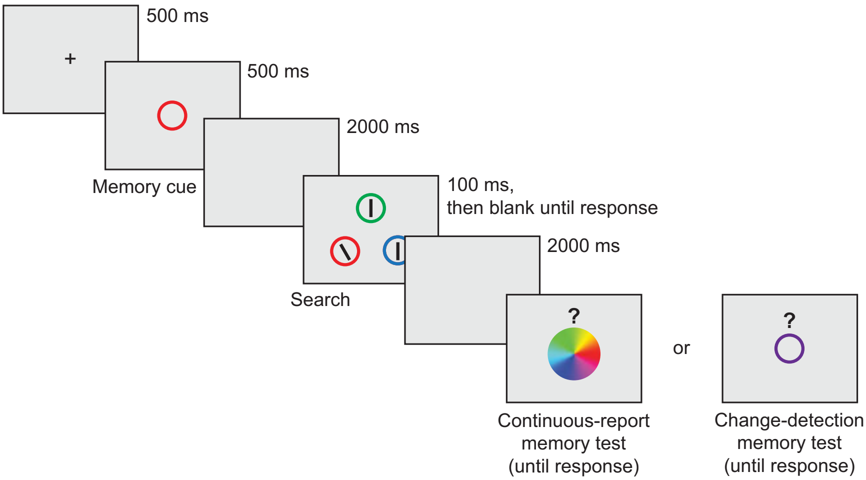

Figure 1. Example trial sequence. Participants were to remember a cue colour while performing an intervening visual search task, in which they reported the direction of the tilted line. Then participants were given one memory task: either (1) a continuous-report task, in which they recalled the memory cue from a continuous spectrum of responses, or (2) a change-detection task, in which they responded whether the memory probe was same or different from the memory cue.

response or search timeout, a blank delay of 2000 ms was presented, followed by a memory probe item that appeared in the centre of the screen until response.

Two types of memory tests were used (Figure 1). In a memory recall test, a continuous-report probe appeared as a multicoloured disc (2.8° diameter) with a question mark above it. Participants were instructed to adjust the colour of the probe to match the initial memory cue colour. The probe's colour was adjusted using an input dial (PowerMate USB Multimedia controller, Griffin Technology, USA). Turning the dial caused the probe's colour to cycle through the space of possible colours. Accuracy was stressed, and there was no time limit. In a memory recognition test, a change-detection probe appeared as a single coloured circle (stroke width = .3°, 2.8° diameter) with a question mark above it. Participants were instructed to indicate whether the probe was identical to the initial memory cue colour by using designated keypresses to respond "same" or "different." Match and non-match probes occurred equally often, in a randomized order; non-match probes were approximately 20° different from the initial memory cue. Both accuracy and speed were stressed for these trials.

Design

Trials were classified by validity (i.e., neutral, valid, or invalid), depending on the relationship between the memory cue and the search array. In valid trials, the memory cue matched a circle within the search display, such that the target tilted line was contained within this match. In invalid trials, the memory cue matched a circle within the search display, but the target tilted line was never contained within this match. Memory item colours were pseudorandomly drawn to sample the entire range of circular HSV colour space with constant value and saturation. Search circle colours were set equidistantly from each other and from the memory cue colour.

The critical manipulation was that trial conditions were blocked. Blocks could be composed of 100% valid trials ("Valid"), 100% invalid trials ("Invalid"), or 50% valid and 50% invalid trials ("Mixed"). Participants completed a practice session to learn about the different trial conditions and blocks; during the experimental phase, each block began with explicit instructions informing participants of block type and condition percentages.

Data analysis

Our primary analysis focused on the continuous-report memory task. For each trial, a measure of recall error was calculated as the angular deviation between the memory response and the initial memory cue (range −180° to 180°). Overall memory accuracy was calculated as the root mean squared error (RMSE; see Fan & Turk-Browne, 2013).

To quantify the different components (i.e., quality and accessibility) of memory performance, we employed a mixture modelling technique. For each subject and for each block type, the distribution of recall errors was fit with a standard mixture model (Zhang & Luck, 2008) that attributes memory errors to either noisy recall of the memory cue (a Gaussian, or Von Mises, distribution centred on the cue colour value) or to random guessing (a uniform distribution across all colours). This mixture model produces two parameters reflecting precision and accessibility, respectively: *sd* is the standard deviation of the Gaussian distribution, or the inverse of precision, while *g* is the probability of random guessing (i.e., the probability of forgetting an item).

As in previous studies, we first compared average precision and average probability of random guessing across block types. However, we extended this analysis by using a hypothesis-driven variant of factorial model comparison (see van den Berg & Ma, 2014), which allows us to obtain an approximate probability that each subject uses a particular strategy or model. We employed a total of 20 standard mixture models that varied in their application of prior probability distributions ("priors") on *sd* and *g* across block types (Table 1): Model 1, our baseline model, represented a standard mixture model with constant precision and constant guessing across block types; Models 2–10 represented models with constant precision, but variable or ordered guessing; Models 11–17 represented models with variable or ordered precision, but constant guessing; and finally, Models 18–20 represented models with variable precision and variable guessing. For these final models, only a few variants of variable precision and ordered guessing were included, based on the best fitting models within Models 2–17.

As previously mentioned, proper model comparison requires determining exact marginal likelihoods, which is computationally difficult for models with greater numbers of parameters. Instead we obtained approximate marginal likelihoods by computing a related quantity, the deviance information criterion (DIC; Spiegelhalter, Best, Carlin, & van der Linde, 2002), a relative measure of model fit and complexity that can be approximated from samples of the posterior distribution over parameters. These DIC values were calculated for each model and each subject using the MemToolbox (Suchow et al., 2013). We also employed a Bayesian model selection approach (see Stephan, Penny, Daunizeau, Moran, & Friston, 2009) that treats models as random effects that can vary across subjects. We submitted approximate log model evidences to the spm_BMS routine of the SPM8 software suite (http://www.fil.ion.ucl.ac.uk/spm/software/spm8). With this method, we approximated the expected posterior probability that a randomly chosen subject will use each model, as well as the exceedance probability P^* of a model being more likely than any other model. While the exceedance probability of a model should not be interpreted as a "Bayesian *p*-value," it does express confidence about that model's favourability (Rigoux, Stephan, Friston, & Daunizeau, 2014).

TABLE 1
Descriptions of model families and model prior probability distributions for each parameter across block types

Model number	*Inverse precision* (sd)	*Probability of guessing* (g)
Family 1: Constant precision, constant probability of guessing		
Model 1	Valid = Mixed = Invalid	Valid = Mixed = Invalid
Family 2: Constant precision, variable/ordered probability of guessing		
Model 2	Valid = Mixed = Invalid	{Valid},{Mixed},{Invalid}*
Model 3	Valid = Mixed = Invalid	Mixed > Invalid & Mixed > Valid
Model 4	Valid = Mixed = Invalid	Invalid > Mixed & Invalid > Valid
Model 5	Valid = Mixed = Invalid	Valid < Mixed & Valid < Invalid
Model 6	Valid = Mixed = Invalid	Mixed > Invalid > Valid
Model 7	Valid = Mixed = Invalid	Invalid > Mixed > Valid
Model 8	Valid = Mixed = Invalid	Mixed > Invalid & Invalid = Valid
Model 9	Valid = Mixed = Invalid	Invalid > Mixed & Mixed = Valid
Model 10	Valid = Mixed = Invalid	Valid < Invalid & Invalid = Mixed
Family 3: Constant probability of guessing, variable/ordered precision**		
Model 11	{Valid},{Mixed},{Invalid}	Valid = Mixed = Invalid
Model 12	Mixed > Invalid & Mixed > Valid	Valid = Mixed = Invalid
Model 13	Invalid > Mixed & Invalid > Valid	Valid = Mixed = Invalid
Model 14	Valid < Mixed & Valid < Invalid	Valid = Mixed = Invalid
Model 15	Mixed > Invalid & Invalid = Valid	Valid = Mixed = Invalid
Model 16	Invalid > Mixed & Mixed = Valid	Valid = Mixed = Invalid
Model 17	Valid < Invalid & Invalid = Mixed	Valid = Mixed = Invalid
Family 4: Variable precision, variable/ordered probability of guessing		
Model 18	{Valid},{Mixed},{Invalid}	{Valid},{Mixed},{Invalid}
Model 19	{Valid},{Mixed},{Invalid}	Mixed > Invalid > Valid
Model 20	{Valid},{Mixed},{Invalid}	Invalid > Mixed > Valid

*Curly brackets indicate variable precision or variable probability of guessing, such that the parameter value for each block type is independent.

**Note that *sd* refers to the standard deviation of the Gaussian distribution centre on the target, meaning that it is actually the inverse of precision. Thus, greater *sd* implies less precision, and vice versa.

EXPERIMENT 1

The present experiment was designed to replicate the finding that cognitive control can modulate memory-based attentional guidance, as reflected in both search performance and memory performance (Kiyonaga et al., 2012). Here, the memory task could probe recognition (change-detection) or recall (continuous-report). Including continuous-report responses allowed us to assess how cognitive control over memory-based guidance affects the contents of WM.

Twenty volunteers (three men, ages 18–21 years) participated in exchange for course credit or a payment of US$10 per hour and signed informed consent. Data from two additional volunteers were removed for low accuracy (<80%) on the search task. Participants completed two blocks of each condition (Valid, Invalid, Mixed), presented in random order. In each type of block, 40% of trials had a change-detection memory probe, while 60% of trials had a continuous-report memory probe. Each block had 42 trials, for a total of 50 continuous-report responses per block type and 252 trials total.

Behavioural results

Visual search task. Search accuracy was near ceiling (mean = 96.13%, SD = 2.47%). We analyzed response times for correct visual search responses according to block and trial conditions (Table 2). A two-way repeated-measures ANOVA across factors of validity (i.e., valid or invalid) and predictability (i.e., 100%-predictable or Mixed) revealed that search was overall faster on valid (vs. invalid) trials, $F(1,19) = 168.16$, $p < .001$, and nearly significantly faster on 100%-predictable (vs. Mixed) blocks, $F(1,19) = 4.25$, $p = .053$. In other words, valid memory cues led to search benefits, while invalid memory cues led to search costs; furthermore, benefits tended to be enhanced, and costs attenuated, by prior knowledge of whether the memory cues were reliably helpful or hurtful, respectively. The relative impact of predictability was not modulated by validity, $F(1,19) = .05$, $p = .827$.

Change-detection memory task. Mean accuracy for the change-detection memory task was 73.82% (SD = 5.37%), confirming that the task was challenging, yet participants were generally able to maintain cues in WM. However, neither accuracy nor response times for the change-detection memory task differed significantly across validity or predictability (Table 2).

Continuous-report memory task. Mean RMSE (i.e., error) on the continuous-report memory task was 24° (SD = 5.5°; Table 2). A two-way repeated-measures

TABLE 2
Mean accuracy and response times for visual search and memory tasks for each trial type

	Visual search task		*Change-detection task*		*Continuous-report task*
	Accuracy (%)	*RT (ms)*	*Accuracy (%)*	*RT (ms)*	*Accuracy (°, RMSE)*
Experiment 1, $n = 20$					
100% valid	97.0 (2.3)	590 (137)	75.9 (7.0)	996 (182)	19.5 (4.4)
50% valid	96.3 (5.0)	624 (133)	73.8 (10.0)	987 (234)	21.5 (9.0)
100% invalid	96.0 (3.2)	717 (131)	71.6 (6.9)	962 (204)	26.7 (9.1)
50% invalid	94.5 (4.1)	755 (177)	74.2 (12.6)	997 (218)	25.5 (8.9)
Experiment 2, $n = 20$					
100% valid	95.6 (4.2)	590 (118)	–	–	18.2 (3.3)
50% valid	96.4 (2.6)	637 (117)	–	–	18.8 (6.4)
100% invalid	94.0 (4.2)	728 (151)	–	–	24.6 (8.0)
50% invalid	93.3 (5.1)	756 (155)	–	–	23.0 (11.5)

Standard deviations are presented in parentheses.

ANOVA revealed that memory error was significantly greater for invalid (vs. valid) trials, $F(1,19) = 10.98$, $p = .004$, but there was no significant effect of predictability and no significant interaction, *p*s > .2. Similarly, a one-way repeated-measures ANOVA across block type (Valid, Invalid, Mixed) showed that memory error was greatest for Invalid blocks and lowest for Valid blocks, $F(2,38) = 7.15$, $p = .002$.

Model fitting and comparison

Recall errors on the continuous-report memory task were fit with a standard mixture model (see Zhang & Luck, 2008) to obtain average estimates of precision and guessing (Table 3). Because memory load included only a single item, parameter estimates for *sd* ($M = 17°$, $SD = 2°$) and *g* ($M = 4\%$, $SD = 3\%$) were relatively high; however, these ranges are comparable to those in previous studies with low memory load (e.g., Hollingworth & Hwang, 2013; van Moorselaar et al., 2014; also see Bays et al., 2009; Zhang & Luck, 2008). Comparing these parameter estimates across block types showed no significant differences in *sd*, $F(2,38) = .41$, $p = .668$, but significant differences in *g*, $F(2,38) = 4.54$, $p = .029$ (Greenhouse-Geisser corrected). In other words, a comparison of group means indicated that while the precision of memory recall did not change across block types, the probability of forgetting that memory item did.

For Bayesian model comparison, recall errors were fit with the 20 different mixture models (Table 1). Figure 2a depicts, for each subject, the DIC value of each model relative to the DIC value of the most likely model for that subject (higher values represent worse fits). We then computed the expected posterior probability that a randomly chosen subject would use each model, as well as the exceedance probability P^* of that model being more likely than any other model.

TABLE 3
Mean *sd* and g values for continuous-report memory task for each block type

	Continuous-report task	
	sd (°)	*g (%)*
Experiment 1, *n* = *20*		
Valid	17.0 (2.8)	1.3 (2.1)
Mixed	17.7 (3.4)	3.8 (4.1)
Invalid	17.0 (3.0)	5.6 (7.0)
Experiment 2, *n* = *20*		
Valid	16.3 (2.2)	1.3 (1.6)
Mixed	15.7 (2.4)	3.7 (5.0)
Invalid	15.4 (2.3)	5.2 (4.7)

Standard deviations are presented in parentheses.

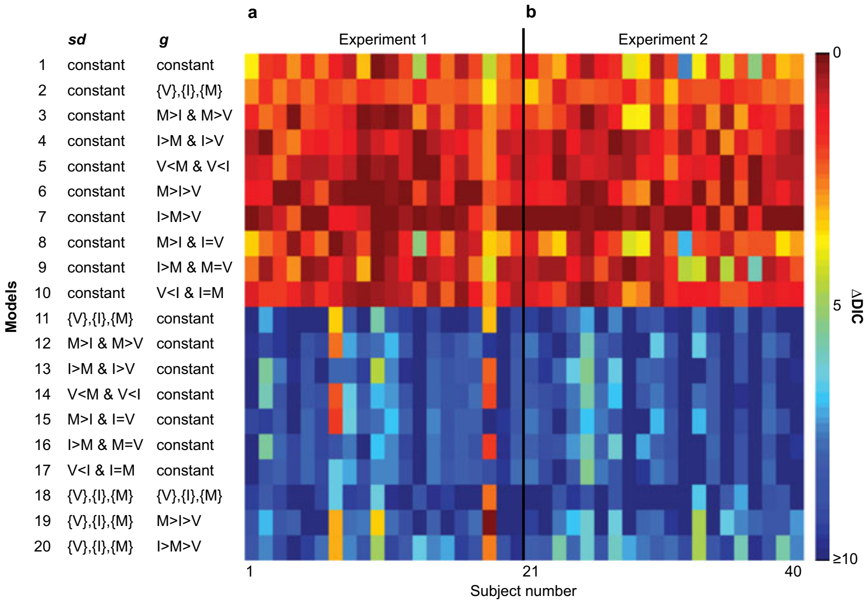

Figure 2. Bayesian model comparison between individual models for each participant. Each column represents a subject, divided by Experiment 1 (**a**) and Experiment 2 (**b**), and each row represents a mixture model as identified by model number (see Table 1) and prior probability distributions over two parameters: the standard deviation of the Gaussian component (*sd*) and the probability of drawing from the uniform component (i.e., guessing; *g*). Prior probabilities were applied across three block types: Valid (V), Invalid (I), and Mixed (M). Curly brackets indicate variable precision or guessing across the three block types. Cell colour indicates a model's deviance information criterion (DIC) relative to that of the subject's best-fitting model (a higher value means a worse fit). In general, the best-fitting models (redder cells) specify constant precision, whereas the worst-fitting models (bluer cells) specify variable or ordered precision.

Across subjects, the highest-ranking model specified a standard mixture model with constant *sd* but ordered *g* across block types, such that the probability of random guessing was greatest in Invalid blocks, followed by Mixed blocks, and least in Valid blocks (posterior = .26, P^* = .848). Critically, all of the models that specified variable or ordered *sd* were highly disfavoured, with the summed posterior probability of all 10 of these models not exceeding .20 (summed posterior = .18, summed P^* = .002).

Discussion

Experiment 1 indicated that cognitive control, as driven by strategic knowledge about whether memory cues would be reliably helpful or hurtful for the subsequent search task, can modulate WM-attention interactions, replicating the results of Kiyonaga et al. (2012, Experiment 1). These effects did not extend to memory performance in the change-detection memory task, although change-detection accuracy was numerically higher for valid compared to invalid trials. However, in the more difficult continuous-report task, overall memory accuracy was impacted by validity, with higher accuracy for valid versus invalid trials. These results support the idea that strategic enhancement or suppression of the memory cue during the delay (i.e., across the search task) affected memory performance. Model fitting and comparisons revealed that participants' behaviour was best described by a strategy in which the probability of forgetting, as opposed to the precision of WM items, varied across block types. In other words, the relative enhancement or suppression of the WM item was not associated with differences in precision of that item, but rather with differences in the item's accessibility in WM, as reflected in the greater probability of randomly guessing during Invalid blocks.

EXPERIMENT 2

To strengthen the model comparison results of Experiment 1, Experiment 2 was designed as a replication with enhanced power, by increasing the number of continuous-report response data points per block type, from 50 to 120 (see Anderson & Awh, 2012; Bays et al., 2009). Thus, in Experiment 2, all trials contained recall tests, using continuous-report memory probes.

Twenty volunteers (five men, age 19–29 years) participated in exchange for course credit or a payment of US$10 per hour and signed informed consent. The data from three additional participants were removed, one for low search accuracy (<80%), one for search RTs that exceeded 2.5 standard deviations above the group mean, and one for continuous-report memory accuracy (as measured by RMSE) that exceeded 2.5 standard deviations above the group mean. Participants completed two blocks of each condition (Valid, Invalid, Mixed), in random order. Each block had 60 trials, for a total of 120 continuous-report responses per block type and 360 trials total.

Behavioural results

Visual search task. Search accuracy was near ceiling (mean = 94.93%, SD = 3.53%). A two-way repeated-measures ANOVA across factors of validity and predictability revealed that search was significantly faster on valid (vs. invalid) trials, $F(1,19) = 86.02$, $p < .001$, and significantly faster on 100%-predictable (vs. Mixed) blocks, $F(1,19) = 6.46$, $p = .020$, with no interaction, $F(1,19) = 1.40$, $p = .252$ (Table 2). Thus, as in Experiment 1, search benefits and costs were strategically modulated based on prior knowledge of whether the memory cues were reliably helpful or hurtful.

Continuous-report memory task. Mean RMSE (i.e., error) on the continuous-report memory task was 22° (SD = 6.0°; Table 2). A two-way repeated-measures ANOVA revealed that memory error was significantly greater for invalid (vs. valid) trials, $F(1,19) = 14.27$, $p = .001$, but there was no significant effect of predictability and no significant interaction, ps > .5. Similarly, a one-way repeated-measures ANOVA across block type (Valid, Invalid, Mixed) showed that memory error was greatest for Invalid blocks and lowest for Valid blocks, $F(2,38) = 7.89$, $p = .001$.

Model fitting and comparison

Mean parameter estimates of sd (M = 16°, SD = 2°) and g (M = 3%, SD = 3%) were comparable to those in Experiment 1 (Table 3). Comparing parameter estimates of precision and probability of random guessing across block types showed no significant differences in sd, $F(2,38) = 1.57$, $p = .221$, but significant differences in g, $F(2,38) = 7.30$, $p = .002$ (Table 3). As in Experiment 1, the precision of memory recall did not change across block types, but the probability of forgetting did.

Figure 2b depicts for each subject, the DIC value of each model relative to the DIC value of the most likely model for that subject (higher values represent worse fits). Across subjects, the highest-ranking model was the same as in Experiment 1, specifying a standard mixture model with constant sd but ordered g across block types, such that the probability of random guessing was greatest in Invalid blocks, followed by Mixed blocks, and least in Valid blocks (posterior = .40, P^* = .999). As in the previous experiment, all of the models that specified variable or ordered sd were highly disfavoured (summed posterior = .18, summed $P^* < .001$).

Discussion

Experiment 2 again demonstrated that cognitive control can strategically modulate memory-based attentional guidance. Furthermore, this strategic modulation seems to be associated with changes in the accessibility, not the quality, of the WM item. Model fitting was performed on over twice as many continuous-report response data points per block, increasing the reliability of parameter estimates. Results from Bayesian model selection highlighted a specific winning strategy in which the probability of guessing changed across block types, whereas precision remained constant, with even stronger evidence for this particular model in Experiment 2 compared to Experiment 1 (P^* = .999 versus P^* = .848, respectively).

COMPARISON OF MODEL FAMILIES

For further Bayesian model comparisons, we collapsed the continuous-report memory task data from both Experiments 1 and 2. As reported above, the highest-ranking and best-fitting model across all 40 participants specified a standard mixture model with constant *sd* but ordered *g* across block types (posterior = .50, $P^* > .999$), such that the probability of random guessing was greatest in Invalid blocks, followed by Mixed blocks, and lowest in Valid blocks (Figure 3a).

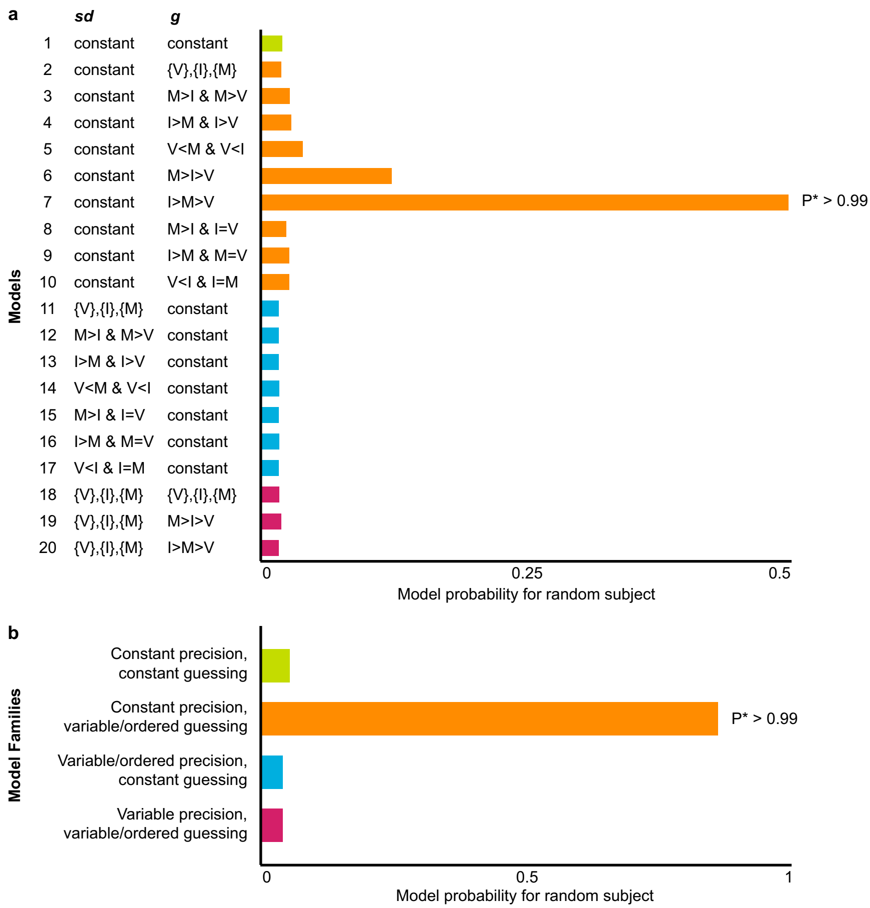

Figure 3. Comparison of individual models (**a**) and of four major families of models (**b**). Bars indicate the probability that a given model family generated the data of a randomly chosen subject. P^* values indicate exceedance probability (i.e., posterior probability that a given model is more likely than any other model). The best-fitting model (Model 7) specifies constant precision but ordered guessing (Invalid > Mixed > Valid), and the best-fitting model family specifies constant precision but variable/ordered guessing across block types.

To directly ask whether strategic modulation of memory-based guidance affects either the quality or the accessibility of WM, we performed Bayesian model comparison of model families. Here “model family” refers to the subset of all models that share a particular level of a particular parameter. From the previous sample of 20 models, four families emerged: constant *sd* and constant *g* across block types (Model 1); constant *sd*, but variable *g* (Models 2–10); variable *sd*, but constant *g* (Models 11–17); and variable *sd* and variable *g* (Models 18–20). For each model family, we computed the expected posterior probability that a randomly chosen subject will use that model family, as well as the exceedance probability P^* of that model family being more likely than any other model family. Figure 3b depicts the comparison of model families across all 40 participants. The entire family of models with variable or ordered *g* across block types was ranked highest (posterior = .92, $P^* = 1$), whereas the entire family of models with variable or ordered *sd* was ranked lowest (posterior = .02, $P^* = 0$). These results provide further evidence that cognitive control over the interaction between WM and attention does not seem to impact the precision of the WM item.

GENERAL DISCUSSION

Across two experiments, we assessed the manner in which a WM representation is modulated when the guidance of attention by WM is strategically enhanced or suppressed. In line with previous studies, valid memory cues benefited search performance, whereas invalid memory cues hindered search performance; prior knowledge about whether WM items would be reliably predictive (or anti-predictive) of the location of a search target allowed for strategic enhancement of valid memory cues and strategic suppression of invalid memory cues (e.g., Carlisle & Woodman, 2011; Kiyonaga et al., 2012). Importantly, our model comparison analysis showed that cognitive control over the impact of WM on visual attention was reflected in visual WM as a change in the probability of random guessing, but not as a change in precision. Thus, strategic modulation over memory-based attentional guidance appears to be instantiated via changes in the accessibility of that WM representation. Specifically, an enhanced WM item could be more easily accessed by WM retrieval processes and thus more likely to impact attention, whereas a suppressed WM item could be stored in a manner that limits its impact on attention, as well as its ease of subsequent retrieval, thus increasing the likelihood of forgetting. Importantly, cognitive control over WM does not seem to modulate representational quality.

The current study provides strong evidence in favour of a growing literature in which the task-relevance of a memory item affects its accessibility rather than its representational quality (e.g., Hollingworth & Hwang, 2013; Murray et al., 2013; Souza et al., 2014; but see Pertzov et al., 2013; Williams et al., 2013). While these previous studies indicated memory-relevance via retrospective cues, the current study indicated memory-relevance in relation to the subsequent attentional task. Here, memory-relevance was not based on whether the WM item would be tested, but was instead based on the likelihood that WM content would benefit visual search, thus recruiting strategic control over the interaction between WM and attention. Moreover, instead of simply evaluating changes in model parameter estimates, the current analysis employed a more powerful method, Bayesian model comparison, to determine the likeliest model of behaviour for the observed data (Stephan et al., 2009). Thus, by overcoming key limitations of previous studies, we here resolved a central ambiguity in the previous literature, showing that the mechanism by which cognitive control modulates the linkage between WM contents and attention seems to operate on the accessibility rather than the quality of memoranda.

Why would cognitive control over WM affect accessibility instead of quality? Previous studies have found that WM representations are maintained with intact precision for several seconds before terminating suddenly and completely (Donkin, Nosofsky, Gold, & Shiffrin, 2013; Zhang & Luck, 2009). This "sudden death" of items in WM is thought to reflect a thresholded forgetting process in which an item gradually decays (not in terms of representational quality, but perhaps in terms of feature-location bindings or distinctiveness) to a certain threshold, then is suddenly dropped from memory (Zhang & Luck, 2009). Along these lines, strategic suppression of memory cues that were reliably unhelpful for subsequent search may have operated via a similar thresholded decay, such that the "suppressed" maintenance of that WM item led to a higher likelihood of being forgotten altogether. An alternative theory suggests that WM accessibility is impacted by the interference and overwriting from simultaneous (Bays et al., 2009) or subsequent items (Gorgoraptis, Catalao, Bays, & Husain, 2011; Makovski, 2012; Souza & Oberauer, 2014). Thus, cognitive control over WM could enact an active suppression process that removes task-irrelevant or outdated information from WM (e.g., Ecker, Oberauer, & Lewandowsky, 2014). In the case of the current study, WM items in invalid trials were irrelevant for the search task, but not for the subsequent memory task; active suppression of the WM item during the search task, however, could spill over into the memory task, resulting in reduced accessibility at retrieval.

This idea of active suppression dovetails well with an influential theory of how WM interacts with attention, which proposes multiple, functionally-different states of WM: a more active type of representation that functions as an attentional template and can directly affect perception, and a less active "accessory" representation that does not influence the deployment of attention (Olivers et al., 2011). Another prominent theory of multiple-state WM posits a more distant state of activated long-term memory (Oberauer, 2002), in which WM representations are functionally partitioned away from the direct focus of (internal) attention (i.e., the most active state of WM), and as such might be more prone to retrieval failure. This form of maintenance in less active long-term memory, however, would not necessarily reduce the precision of the WM representation; recent studies have shown that visual long-term memory can store and retrieve thousands of objects and scenes with very high fidelity (e.g., Brady, Konkle, Gill, Oliva, & Alvarez, 2013). Within these frameworks, cognitive control over WM-attention interactions could be instantiated by strategically assigning different memory representations to different functional states in WM, based on their assumed utility for the forthcoming task (Kiyonaga et al., 2012). In the context of the present study, reliably invalid memory cues could be actively relegated to an accessory or long-term-like state, which may edge closer to a forgetting threshold, thus increasing the probability of forgetting the item and increasing the probability of guessing during a continuous-report memory task.

Note that the current study does not differentiate between quantized or continuous models of visual WM capacity (i.e., how cognitive resources are distributed between representations in visual WM), particularly because only a single item was maintained during each trial. However, the present results can be reconciled with both discrete (e.g., Zhang & Luck, 2008) and flexible (e.g., Bays & Husain, 2008; Ma, Husain, & Bays, 2014) models of visual WM. Both camps allow for individual representations to be maintained with some variability and imperfection, such that individual items might gradually decay along some dimension that impacts accessibility. Thus, both models could be subject to a thresholded, lower bound of memory retrieval, which could be impacted by cognitive control.

The current design also does not speak to whether changes in memory quality and/or accessibility are sufficient to drive attentional guidance (but see Hollingworth & Hwang, 2013; van Moorselaar et al., 2014); instead, our interpretation is focused on how cognitive control over the interaction between WM and attention may be instantiated. We propose that cognitive control in this paradigm is enacted through changes to WM accessibility; however, an alternative explanation is that participants were strategically attending to the memory-matching item during search, in order to refresh their memory representation of that colour for the subsequent probe (see Woodman & Luck, 2007). We do not believe that resampling would have specifically impacted WM accessibility across blocks because the memory-matching colour reappeared in every single search display, valid or invalid. Thus, the opportunity for memory refreshing was available in all blocks.

Using a powerful Bayesian model comparison technique, the present study demonstrates how strategic modulation of memory-based attentional guidance impacts the probability of remembering, but not the precision, of items in WM. These findings suggest that different functional states in WM may be characterized by differing retrieval accessibility rather than representational quality, such that cognitive control over the interaction between WM and attention operates by changing how likely an item is to be remembered or forgotten.

REFERENCES

Anderson, D. E., & Awh, E. (2012). The plateau in mnemonic resolution across large set sizes indicates discrete resource limits in visual working memory. *Attentino, Perception, & Psychophysics*, *74*, 891–910. doi:10.3758/s13414-012-0292-1

Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7. doi:10.1167/9.10.7

Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*, 851–854. doi:10.1126/science.1158023

Brady, T. F., Konkle, T., Gill, J., Oliva, A., & Alvarez, G. A. (2013). Visual long-term memory has the same limit on fidelity as visual working memory. *Psychological Science*, *24*, 981–990. doi:10.1177/0956797612465439

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436. doi:10.1163/156856897X00357

Carlisle, N. B., & Woodman, G. F. (2011). Automatic and strategic effects in the guidance of attention by working memory representations. *Acta Psychologica*, *137*, 217–225. doi:10.1016/j.actpsy.2010.06.012

Donkin, C., Nosofsky, R. M., Gold, J. M., & Shiffrin, R. M. (2013). Discrete-slots models of visual working-memory response times. *Psychological Review*, *120*, 873–902. doi:10.1037/a0034247

Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science*, *11*, 467–473. doi:10.1111/1467-9280.00290

Downing, P., & Dodds, C. (2004). Competition in visual working memory for control of search. *Visual Cognition*, *11*, 689–703. doi:10.1080/13506280344000446

Ecker, U. K. H., Oberauer, K., & Lewandowsky, S. (2014). Working memory updating involves item-specific removal. *Journal of Memory and Language*, *74*, 1–15. doi:10.1016/j.jml.2014.03.006

Fan, J. E., & Turk-Browne, N. B. (2013). Internal attention to features in visual short-term memory guides object learning. *Cognition*, *129*, 292–308. doi:10.1016/j.cognition.2013.06.007

Gorgoraptis, N., Catalao, R. F., Bays, P. M., & Husain, M. (2011). Dynamic updating of working memory resources for visual objects. *The Journal of Neuroscience*, *31*, 8502–8511.

Hollingworth, A., & Hwang, S. (2013). The relationship between visual working memory and attention: Retention of precise colour information in the absence of effects on perceptual selection. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *368*(1628), 1–9.

Houtkamp, R., & Roelfsema, P. R. (2006). The effect of items in working memory on the deployment of attention and the eyes during visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *32*, 423–442. doi:10.1037/0096-1523.32.2.423

Kiyonaga, A., Egner, T., & Soto, D. (2012). Cognitive control over working memory biases of selection. *Psychonomic Bulletin & Review*, *19*, 639–646. doi:10.3758/s13423-012-0253-7

Kruschke, J. K. (2010). What to believe: Bayesian methods for data analysis. *Trends in Cognitive Sciences*, *14*, 293–300. doi:10.1016/j.tics.2010.05.001

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, *17*, 347–356. doi:10.1038/nn.3655

Makovski, T. (2012). Are multiple visual short-term memory storages necessary to explain the retro-cue effect? *Psychonomic Bulletin & Review*, *19*, 470–476. doi:10.3758/s13423-012-0235-9

Murray, A. M., Nobre, A. C., Clark, I. A., Cravo, A. M., & Stokes, M. G. (2013). Attention restores discrete items to visual short-term memory. *Psychological Science*, *24*, 550–556. doi:10.1177/0956797612457782

Myung, I. J., & Pitt, M. A. (1997). Applying Occam's razor in modeling cognition: A Bayesian approach. *Psychonomic Bulletin & Review*, *4*(1), 79–95. doi:10.3758/BF03210778

Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 411–421.

Oberfeld, D., & Franke, T. (2013). Evaluating the robustness of repeated measures analyses: The case of small sample sizes and nonnormal data. *Behavior Research Methods*, *45*(3), 792–812. doi:10.3758/s13428-012-0281-2

Olivers, C. N., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: Visual working memory content affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, *32*, 1243.

Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, *15*, 327–334.

Pertzov, Y., Bays, P. M., Joseph, S., & Husain, M. (2013). Rapid forgetting prevented by retrospective attention cues. *Journal of Experimental Psychology: Human Perception and Performance*, *39*, 1224.

Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies—Revisited. *Neuroimage*, *84*, 971–985. doi:10.1016/j.neuroimage.2013.08.065

Soto, D., Heinke, D., Humphreys, G. W., & Blanco, M. J. (2005). Early, involuntary top-down guidance of attention from working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 248–261. doi:10.1037/0096-1523.31.2.248

Soto, D., Hodsoll, J., Rotshtein, P., & Humphreys, G. W. (2008). Automatic guidance of attention from working memory. *Trends in Cognitive Sciences*, *12*, 342–348. doi:10.1016/j.tics.2008.05.007

Souza, A. S., & Oberauer, K. (2014). Time-based forgetting in visual working memory reflects temporal distinctiveness, not decay. *Psychonomic Bulletin & Review*. [epub ahead of print]. doi:10.3758/s13423-014-0652-z

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*. *76*, 2080–2102. doi:10.3758/s13414-014-0687-2

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology*), *64*, 583–639.

Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *Neuroimage*, *46*, 1004–1017. doi:10.1016/j.neuroimage.2009.03.025

Suchow, J. W., Brady, T. F., Fougnie, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision*, *13*(10), 9.

van den Berg, R., & Ma, W. J. (2014). "Plateau"-related summary statistics are uninformative for comparing working memory models. *Attention, Perception, & Psychophysics*, *76*, 2117–2135. doi:10.3758/s13414-013-0618-7

van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological Review*, *121*(1), 124–149. doi:10.1037/a0035234

Van Moorselaar, D., Theeuwes, J., & Olivers, C. N. L. (2014). In competition for the attentional template: Can multiple items within visual working memory guide attention? *Journal of Experimental Psychology: Human Perception and Performance*, *40*, 1450–1464. doi:10.1037/a0036229

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, *4*(12), 11.

Williams, M., Hong, S. W., Kang, M.-S., Carlisle, N. B., & Woodman, G. F. (2013). The benefit of forgetting. *Psychonomic Bulletin & Review*, *20*, 348–355. doi:10.3758/s13423-012-0354-3

Woodman, G. F., & Luck, S. J. (2007). Do the contents of visual working memory automatically influence attentional selection during visual search? *Journal of Experimental Psychology: Human Perception and Performance*, *33*, 363–377. doi:10.1037/0096-1523.33.2.363

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*, 233–235. doi:10.1038/nature06860

Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, *20*, 423–428. doi:10.1111/j.1467-9280.2009.02322.x